

Article

Sources of Cryptocurrency Return Predictability: Econometric and Machine Learning Evidence Across Market States

Jie Xu^{1*}

¹Management School, University of Sheffield, Sheffield, South Yorkshire, United Kingdom.

Keywords:

Cryptocurrency Returns
Machine Learning
Cross-Sectional Prediction
Market Volatility
Asset Pricing
Forecasting

ARTICLE INFO

ABSTRACT

This study examines whether nonlinear machine-learning methods improve cross-sectional cryptocurrency return prediction beyond parsimonious linear benchmarks. Using a dynamic point-in-time cryptocurrency universe, we construct weekly characteristics capturing size, past returns, liquidity, risk, and return distributions and compare OLS, Elastic Net, Random Forest, XGBoost, a shallow neural network, and a forecast combination under strictly out-of-sample evaluation. Return levels remain difficult to forecast: every model produces a negative out-of-sample R^2 , although Random Forest is least negative at -0.29%. Ranking performance is more informative, with Random

Forest delivering the strongest Spearman information coefficient (0.116) and the largest weekly high-minus-low portfolio spread (0.70%), but the latter is statistically imprecise (HAC $t=1.18$). SHAP evidence concentrates predictive importance in risk, return-distribution, and past-return characteristics. State-dependent tests further show that nonlinear and combined forecasts improve relative to OLS as the probability of a high-volatility regime rises; the combination forecast displays the clearest relation ($p=0.034$). Overall, the evidence indicates that cryptocurrency return predictability is concentrated, time varying, and multi-dimensional: model flexibility has conditional rather than universal value, and improvements in return-level accuracy do not automatically translate into stronger rankings or statistically reliable portfolio profits.

1. Introduction

Cryptocurrency markets present an unusual asset-pricing problem. Crypto assets are traded in highly liquid global markets and increasingly resemble conventional financial assets in their investment use, yet their return dynamics remain difficult to reconcile with standard asset-pricing explanations. Early evidence shows that cryptocurrency returns have limited exposure to conventional equity, currency, commodity, and macroeconomic factors. Instead, market-specific variables such as momentum and investor attention appear to contain substantial predictive information (Liu & Tsyvinski, 2021). This distinction raises a broader question. If the returns on crypto assets are not well described by the factors that price traditional assets, what information does the market use to differentiate expected returns across individual tokens?

*Corresponding author
Email: jxu171@Sheffield.ac.uk

Online ISSN: 3136-4578

DOI: <https://doi.org/10.68471/fbai.2026.006>

Copyright: © 2026 The Author(s). Published by Global Frontiers Press Limited. This article is open access under the Creative Commons Attribution 4.0 International license (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/>.

Cross-sectional evidence provides a partial answer. [Liu et al. \(2022\)](#) show that several characteristics familiar from equity markets also predict cryptocurrency returns. Among a broad set of price- and market-based characteristics, size, momentum, volume, and related variables generate significant return spreads, while a parsimonious three-factor model based on the cryptocurrency market, size, and momentum absorbs much of this cross-sectional variation. Subsequent research has substantially expanded this set of candidate predictors to include liquidity, volatility, reversal, tail risk, network activity, and other crypto-specific characteristics. The resulting literature therefore creates a problem similar to the “factor zoo” in equity markets: many predictors may contain overlapping information, some may proxy for the same underlying economic mechanism, and their relationships with future returns need not be linear or stable over time.

Machine learning provides a natural way to address this problem, but its role requires careful interpretation. In conventional equity markets, [Gu et al. \(2020\)](#) show that machine-learning methods can improve the measurement of expected returns by accommodating large predictor sets, nonlinearities, and interactions that standard linear specifications may miss. Their strongest gains arise from tree-based methods and neural networks, and the improvement is traced particularly to nonlinear predictor interactions. At the same time, they emphasize an important distinction: better prediction improves the measurement of expected returns, but prediction alone does not identify the economic mechanism generating those returns. Machine learning should therefore be evaluated not simply by whether a sophisticated algorithm produces a lower forecast error, but by whether it identifies incremental predictive information and whether that information can be given a coherent economic interpretation.

Evidence from cryptocurrency markets makes this issue more complicated. [Cakici et al. \(2024\)](#) apply a broad set of machine-learning methods to more than 500 cryptocurrencies and find substantial cross-sectional return predictability, but only limited gains from model complexity. Simple characteristics—including price, past alpha, momentum, and illiquidity—account for much of the predictive information, while relatively simple models can perform as well as, or better than, more sophisticated tree and neural-network specifications. This result contrasts with the nonlinear advantage documented for equities by [Gu et al. \(2020\)](#). It suggests that the central question for cryptocurrency asset pricing is not whether machine learning can be applied, but when and why model complexity adds information beyond parsimonious return predictors.

This tension motivates the present study. The analysis asks three related questions. First, which market, liquidity, momentum, volatility, attention, and crypto-specific characteristics contain incremental information about future token returns after competing predictors are considered jointly? Second, does the relationship between these characteristics and expected returns remain stable across different market conditions, or does predictability change when the cryptocurrency market moves between relatively calm and more volatile states? Third, do nonlinear machine-learning methods provide economically meaningful out-of-sample improvements over simpler linear and regularized benchmarks? These questions distinguish three concepts that are often combined in forecasting studies: return predictability, model complexity, and economic value.

The empirical design follows this distinction. We first establish a conventional benchmark using economically motivated token characteristics and parsimonious linear specifications. The purpose of this stage is to determine which return patterns survive joint controls rather than to assume that a large predictor set necessarily improves forecasting. We then introduce regularized and nonlinear models, including tree-based methods, to test whether interactions and nonlinearities contain incremental predictive information. Model evaluation is conducted strictly out of sample using temporally ordered training, validation, and testing periods. Forecast accuracy is evaluated alongside cross-sectional ranking ability because a model may estimate the level of future returns imprecisely while still distinguishing relative winners from losers. Finally, predictive signals are examined through variable-importance and explainability methods and are translated into portfolio outcomes to determine whether statistical predictability survives economically relevant constraints.

The study makes three contributions to the literature on cryptocurrency asset pricing and machine-learning forecasting. First, it reframes the model-comparison problem from a horse race among algorithms to a test of whether additional functional flexibility produces incremental information beyond parsimonious economic predictors. This distinction is important because predictive gains can arise from regularization, nonlinear interactions, or adaptation to changing market states, and these mechanisms have different implications for asset-pricing interpretation. Second, the

study introduces a multi-layer evaluation architecture that separates return-level accuracy, cross-sectional ranking, and economic portfolio performance. Treating these outcomes separately avoids the common inference that a model with lower forecast error must necessarily generate superior asset rankings or tradable abnormal returns. Third, the analysis embeds model comparison in a time-varying market environment by combining a dynamic point-in-time universe with volatility-regime conditioning and SHAP-based explainability. This allows the paper to identify a boundary condition for model complexity: flexibility appears to be useful primarily when the mapping from characteristics to returns becomes less stable, rather than as a universally superior forecasting strategy.

The broader objective is therefore not to construct the most complex forecasting model. It is to identify the simplest model capable of capturing economically meaningful variation in token returns, and to determine the conditions under which additional model complexity becomes informative. In this sense, machine learning is used as a tool for testing the structure of cryptocurrency return predictability rather than as a substitute for economic reasoning.

2. Literature Review and Research Development

2.1 *Cryptocurrency Returns as a Distinct Asset-Pricing Problem*

Early cryptocurrency research raised a fundamental asset-pricing question: whether crypto returns could be understood using the same risk factors that explain conventional assets. This distinction matters because the economic interpretation of return predictability depends on what cryptocurrencies represent. If they behave primarily as currencies, their returns should exhibit systematic exposure to currency-market factors. If they function as stores of value, commodity-related risks should matter. If they represent claims on technological growth, conventional equity factors may capture part of their return variation.

[Liu and Tsyvinski \(2021\)](#) provide an important empirical benchmark for this question. Examining Bitcoin, Ripple, and Ethereum, they find little consistent exposure to common stock-market, currency, commodity, or macroeconomic factors. Conventional asset-pricing variables therefore leave much of cryptocurrency return variation unexplained. In contrast, cryptocurrency-specific information performs considerably better. Short-horizon momentum predicts subsequent returns, while measures of investor attention, including Google search activity, contain significant forecasting information. The contrast is economically important: cryptocurrency returns are not simply unusually volatile versions of traditional asset returns; their predictability appears to arise from information generated within the cryptocurrency market itself.

This evidence established two principles that remain central to subsequent research. First, predictors should be evaluated against conventional benchmarks rather than interpreted in isolation. A significant association is economically informative only if it captures information not already represented by established risk factors. Second, cryptocurrency-specific predictors need not correspond to traditional notions of fundamental value. Momentum and attention can contain information even when price-to-fundamental proxies, realized volatility, and supply-side variables show limited forecasting ability. The implication is that understanding cryptocurrency returns requires moving beyond the question of which traditional asset class cryptocurrencies most closely resemble and toward the structure of information within the crypto market itself.

[Baur et al. \(2018\)](#) distinguish Bitcoin's use as an exchange medium from its role as a speculative asset. This distinction supports evaluating return predictability as an asset-pricing question rather than inferring predictive content from payment adoption alone.

2.2 *From Individual Predictors to Common Cryptocurrency Factors*

The next stage of the literature shifted the analysis from the time series of a few major cryptocurrencies to the cross-section of a much broader universe of coins. This change is important. Time-series predictability asks whether the past or current state of one asset predicts its own future return. Cross-sectional predictability instead asks why some cryptocurrencies subsequently outperform others at the same point in time. The latter question is more closely connected to asset pricing because it allows researchers to identify common characteristics associated with differences in expected returns.

[Liu et al. \(2022\)](#) provide the main benchmark for this approach. Rather than selecting a small number of cryptocurrency-specific variables *ex ante*, they begin with characteristics already established in equity research and

construct their cryptocurrency counterparts using information that is actually available in crypto markets. Their candidate variables fall broadly into size, momentum, volume, and volatility categories. The empirical procedure then narrows this broad predictor set through portfolio sorts. Of 25 candidate characteristics, nine generate statistically significant long-short return spreads, including measures related to market capitalization, price, short-horizon momentum, dollar volume, and dollar-volume volatility.

The more important contribution, however, is not the discovery of nine individual return patterns. It is the attempt to determine whether these patterns represent independent sources of expected returns. Liu et al. first show that a cryptocurrency market factor alone cannot account for the identified return spreads. Adding a size factor explains several price-, capitalization-, and volume-related strategies but leaves momentum returns largely unexplained. Conversely, a market-plus-momentum specification captures the momentum strategies but not the size-related patterns. A three-factor model combining the cryptocurrency market, size, and momentum eliminates the significant alphas of all nine strategies.

This progression introduces a second principle for cryptocurrency asset pricing: the objective is not to maximize the number of statistically significant predictors, but to identify a parsimonious structure that captures apparently different return patterns. Several observed characteristics may therefore be different manifestations of the same underlying dimension. Low-volume cryptocurrencies, for example, also tend to be smaller, making an apparent volume premium partly a manifestation of size. The distinction matters increasingly as the cryptocurrency literature documents more return predictors. Without joint evaluation, a growing list of anomalies may overstate the true dimensionality of cryptocurrency expected returns.

2.3 The Expansion of the Cryptocurrency Predictor Set

More recent research has considerably expanded the information set used to explain cryptocurrency returns. In addition to size and momentum, studies have considered liquidity, turnover, bid-ask spreads, realized and idiosyncratic volatility, return skewness and kurtosis, reversals, extreme returns, network adoption, on-chain activity, and other market-specific signals. [Cakici et al. \(2024\)](#), for example, synthesize this literature into 40 cryptocurrency characteristics covering on-chain activity, liquidity, risk and volatility, past returns, return distributions, and additional market-based variables.

This expansion creates an empirical problem. A large predictor set does not necessarily imply a high-dimensional return-generating process. Many variables are correlated, some measure closely related economic concepts, and others may only appear significant because they proxy for more fundamental predictors. The evidence reported by Cakici et al. illustrates this distinction. Although their analysis includes 40 characteristics, most predictive information ultimately concentrates in a small group of comparatively simple variables, particularly price, past alpha, momentum, and illiquidity. When characteristics are grouped by economic function, past returns and liquidity account for the largest shares of predictive information, while more complex on-chain measures play a relatively modest role.

This creates a cryptocurrency version of the broader asset-pricing “factor zoo” problem. The methodological challenge is therefore no longer simply to discover additional variables correlated with future returns. It is to determine which variables remain informative when many competing predictors are considered simultaneously, whether their information is redundant or complementary, and whether their relationships with expected returns are sufficiently nonlinear to justify more flexible models.

2.4 Machine Learning and the Value of Model Complexity

Machine learning enters asset pricing precisely because of this problem. Conventional linear models become increasingly restrictive when the number of predictors grows and when expected returns depend on nonlinearities or interactions between characteristics. [Gu et al. \(2020\)](#) formalize this argument in the equity market. Using a large cross-section of stocks and a high-dimensional predictor set, they compare linear, dimension-reduction, tree-based, and neural-network approaches. Their evidence shows that flexible models, particularly tree methods and neural networks, can generate substantial out-of-sample gains. Importantly, these gains are traced not simply to using more predictors but to the models' ability to capture nonlinear interactions among them.

The interpretation of this result requires care. Machine learning can reveal whether a more flexible conditional-expectation function improves forecasts, but a successful forecast is not itself an economic mechanism. Predictor

importance, nonlinear response functions, and interactions provide evidence about where predictive information resides; they do not by themselves establish why a variable causes expected returns to change. This distinction makes a strong linear benchmark essential. Without it, superior performance by a complex model cannot be attributed specifically to nonlinearities or interactions, because the improvement may instead arise from regularization, variable selection, or differences in model estimation.

Cryptocurrency evidence provides a particularly useful test of this argument because it does not reproduce the equity result mechanically. [Cakici et al. \(2024\)](#) compare OLS, partial least squares, LASSO, elastic net, random forests, gradient-boosted trees, a feed-forward neural network, and forecast combinations. Despite the flexibility of the more sophisticated models, they find only limited benefits from model complexity. OLS performs surprisingly well, while the neural-network specification produces relatively weak statistical forecasts. Tree models improve some measures of ranking ability, and forecast combinations perform competitively, but no general monotonic relationship exists between computational complexity and forecasting performance.

This result creates an important tension with equity-market evidence. If nonlinear methods substantially improve expected-return prediction in stocks but offer only modest gains in cryptocurrencies, at least three interpretations are possible. The cross-section of cryptocurrency returns may be genuinely lower-dimensional; the dominant relationships may be approximately linear; or the shorter and substantially noisier history of cryptocurrencies may make complex functions difficult to estimate reliably. Cakici et al. explicitly consider these possibilities, noting that the volatile and rapidly changing cryptocurrency universe may provide insufficient stable information for sophisticated models to learn persistent nonlinear patterns.

The unresolved issue is therefore not whether machine learning is technically applicable to cryptocurrency returns. It clearly is. The relevant question is whether greater model flexibility reveals economically meaningful information that is absent from parsimonious specifications.

The emphasis on out-of-sample benchmarks follows a broader forecasting lesson: strong in-sample associations need not improve forecasts against a simple historical-mean benchmark ([Welch & Goyal, 2008](#)).

2.5 Statistical Predictability, Economic Value, and Market Frictions

A second gap concerns how forecasting success should be evaluated. Standard statistical measures assess how closely predicted returns match realized return levels. Yet cross-sectional investment strategies often depend more on ranking assets correctly than on forecasting their exact returns. Cakici et al. show this distinction clearly. Their machine-learning models can produce weak or even negative out-of-sample R^2 values while still generating positive cross-sectional correlations between predicted and realized returns. The same models can therefore identify relative winners and losers even when absolute forecast errors remain large.

Economic evaluation consequently provides information that statistical accuracy alone cannot capture. Sorting cryptocurrencies by predicted returns allows the forecasts to be translated into portfolio outcomes, while factor-adjusted alpha, Sharpe ratios, turnover, and transaction costs reveal whether predictability is economically exploitable. This step also helps identify where apparent predictive gains originate. Cakici et al. find that abnormal returns are concentrated disproportionately in small, illiquid, volatile, and difficult-to-trade cryptocurrencies. Moreover, trading costs absorb a substantial share of gross strategy returns, and profitability is sensitive to extreme observations.

These findings introduce an important distinction between predictability and arbitrageability. A return pattern may be statistically persistent precisely because the assets in which it occurs are difficult or costly to trade. In this case, machine-learning profitability does not necessarily indicate a freely exploitable market inefficiency. Instead, it may reflect limits to arbitrage that prevent investors from eliminating the underlying mispricing. Liquidity, volatility, and market conditions therefore have two possible roles: they may directly contain information about expected returns, and they may determine when other predictive relationships can survive.

Market segmentation also constrains implementation. [Makarov and Schoar \(2020\)](#) document persistent cross-exchange price deviations, indicating that a forecast-based spread should not be equated with an immediately executable arbitrage opportunity.

2.6 Research Gap and Empirical Development

Taken together, the literature establishes several robust facts. Cryptocurrency returns are not fully explained by conventional asset-market factors; market-specific characteristics, particularly momentum and attention, contain predictive information. The cross-section of cryptocurrency returns can nevertheless be summarized to a considerable extent by relatively parsimonious factors such as market, size, and momentum. At the same time, the set of documented predictors has expanded substantially, creating a high-dimensional environment in which nonlinear models are potentially useful. Yet existing cryptocurrency evidence suggests that increasing model complexity does not automatically improve forecasting performance.

What remains less clear is when nonlinear complexity becomes useful and what economic conditions account for its incremental value. Existing factor studies emphasize parsimonious unconditional relationships, while machine-learning studies primarily evaluate average forecasting performance across the sample. These approaches leave open the possibility that apparently conflicting findings are state dependent. A simple model may be sufficient under normal market conditions but fail when volatility, liquidity, investor attention, or other market characteristics move into unusual states. Conversely, apparent gains from complex models may be concentrated in market segments where extreme returns and trading frictions are strongest.

The present study addresses this gap by treating model complexity as an empirical question rather than a methodological assumption. The analysis first identifies which token characteristics contain incremental information under parsimonious specifications. It then evaluates whether regularized and nonlinear models improve strictly out-of-sample forecasts. Finally, the analysis examines whether any improvement varies across market conditions and whether the corresponding signals remain economically meaningful after accounting for liquidity, turnover, and trading constraints. This sequence allows the study to distinguish three questions that should not be conflated: whether token returns are predictable, whether nonlinear models improve that predictability, and whether the resulting forecasts represent economically exploitable information.

2.7 Research Questions and Testable Expectations

The preceding literature motivates three research questions. RQ1 asks whether nonlinear models deliver incremental out-of-sample information beyond linear and regularized benchmarks. RQ2 asks whether the source of predictive information is concentrated in a small set of economically interpretable characteristics or dispersed across the broader predictor space. RQ3 asks whether the relative value of model flexibility changes systematically with market conditions, particularly aggregate volatility.

Rather than imposing strong directional hypotheses on an unstable and rapidly evolving market, the empirical design adopts benchmark-based expectations. A nonlinear model is considered informative only if it improves at least one out-of-sample performance dimension without relying on in-sample fit. Predictor-importance results are interpreted as evidence of conditional predictive relevance rather than causal or equilibrium pricing effects. Regime-dependent gains are interpreted as boundary conditions on model usefulness rather than proof that volatility causes nonlinear predictability.

This research-question structure is deliberately conservative. It aligns the statistical claims with what the design can identify and separates prediction, explanation, and economic implementation.




Figure 1. Empirical Architecture and Inference Map of Cryptocurrency Return Predictability

3. Data and Empirical Design

3.1 Data and Sample

We collect daily cryptocurrency market data from 1 January 2018 to 30 April 2026. The principal market variables are closing price, market capitalization, and trading volume denominated in U.S. dollars. The sample is constructed as a dynamic universe rather than from a fixed list of cryptocurrencies observed at the end of the sample. Thus, a token enters the sample when it first satisfies our screening criteria and exits when valid market observations are no longer available.

We apply several filters to obtain a tradeable sample. We exclude stablecoins, wrapped tokens, and assets whose prices are explicitly designed to track fiat currencies, commodities, or other underlying instruments. We require positive observations for price, market capitalization, and trading volume. We further require a market capitalization of at least USD 1 million and a minimum trading history of 26 weeks. All continuous token characteristics are winsorized at the 1st and 99th percentiles within each formation week. We consider alternative capitalization thresholds of USD 2 million, USD 5 million, and USD 10 million in the robustness analysis.

We use daily observations to construct weekly token returns and characteristics. Cryptocurrency markets trade continuously, so each prediction period consists of seven calendar days. The dependent variable is the one-week-ahead excess return,

$$r_{i,t+1}^e = \ln(P_{i,t+1}/P_{i,t}) - r_{f,t+1}, \quad (1)$$

where $P_{i,t}$ is the closing price of token i at the end of week t . We impose a one-day gap between the information used to construct the predictors and the beginning of the subsequent holding period. Tokens with insufficient historical observations for a particular characteristic remain excluded from that prediction week.

The sample period contains several distinct cryptocurrency market episodes. It begins after the 2017 expansion, includes the 2018 market contraction, the 2020–2021 appreciation cycle, the sharp market adjustment of 2022, and the subsequent recovery. This variation is central to our analysis because the predictive relation between token characteristics and future returns need not remain constant across these market conditions.

In addition to token-level information, we collect a set of common market variables. The cryptocurrency market return is constructed as the value-weighted return of all eligible tokens. We measure aggregate crypto volatility from daily market returns. The external market variables consist of the CBOE VIX, the effective federal funds rate, and the broad U.S. dollar index. Investor attention is measured using cryptocurrency-related Google search interest. These variables are observed before each forecast is formed.

3.2 Cryptocurrency Characteristics

We next construct a set of characteristics that describe differences across tokens at each point in time. Table 1 summarizes their definitions. The variables fall into five broad groups: size and price, past returns, liquidity, risk, and the distribution of recent returns.

Table 1. Cryptocurrency Characteristics

Category	Variable	Definition
Size	SIZE	Log market capitalization
Size	PRICE	Log token price
Size	AGE	Number of weeks since first valid observation
Past returns	MOM1	Previous 1-week return
Past returns	MOM2	Previous 2-week return
Past returns	MOM4	Previous 4-week return
Past returns	MOM8	Previous 8-week return
Past returns	MOM12	Previous 12-week return
Past returns	REV	Short-term reversal
Liquidity	DVOL	Log dollar trading volume
Liquidity	TURN	Dollar volume divided by market capitalization
Liquidity	ILLIQ	Average absolute return divided by dollar volume
Liquidity	VOLVOL	Standard deviation of trading volume
Risk	RV4	Realized volatility over the previous 4 weeks
Risk	RV12	Realized volatility over the previous 12 weeks
Risk	BETA	Beta relative to the cryptocurrency market
Risk	IVOL	Idiosyncratic volatility
Distribution	SKEW	Skewness of daily returns
Distribution	KURT	Kurtosis of daily returns
Distribution	MAXRET	Maximum daily return
Distribution	MINRET	Minimum daily return

Momentum over k weeks is constructed from cumulative past log returns,

$$MOM_{i,t}^{(k)} = \ln(P_{i,t}/P_{i,t-k}). \quad (2)$$

Our liquidity measures capture different aspects of trading activity. Dollar volume measures the scale of trading, turnover scales this activity by market capitalization, and the [Amihud \(2002\)](#) illiquidity measure captures the price movement associated with a given amount of trading:

$$ILLIQ_{i,t} = \frac{1}{D} \sum_{d=1}^D \frac{|r_{i,d}|}{DollarVolume_{i,d}}. \quad (3)$$

For risk characteristics, we calculate realized volatility from daily returns over four- and twelve-week windows. We estimate market beta using daily token returns and the value-weighted cryptocurrency market return. Idiosyncratic volatility is the standard deviation of the corresponding regression residuals. Finally, skewness, kurtosis, and extreme daily returns capture information in the shape and tails of the recent return distribution that is not contained in average returns or volatility.

Before estimation, characteristics are standardized cross-sectionally each week. Consequently, the prediction models use information about the relative position of a token within the contemporaneous cryptocurrency market rather than differences in the units or scale of the raw variables.

3.3 Common Factors and Market States

We first construct three common cryptocurrency factors. The market factor, $CMKT_t$, is the value-weighted excess return of all eligible tokens. For the size factor, we rank tokens each week by market capitalization and form portfolios from the bottom and top 30% of the distribution. The cryptocurrency size factor is

$$CSIZE_t = R_{Small,t} - R_{Large,t}. \quad (4)$$

We construct the momentum factor from tokens with high and low recent momentum:

$$CMOM_t = R_{HighMom,t} - R_{LowMom,t}. \quad (5)$$

These three factors provide the baseline return structure against which the additional token characteristics are evaluated. The first stage of the analysis therefore examines whether the broader characteristic set contains information that is not absorbed by market, size, and momentum.

We next allow the predictive relationships to vary with market conditions. We use two complementary measures of market states. The first consists of observable market variables: cryptocurrency-market volatility, VIX, the federal funds rate, the U.S. dollar index, and investor attention. Because these variables are common to all tokens in a given week, we interact them with token characteristics rather than treating them as independent sources of cross-sectional variation. For example, interactions between illiquidity and market uncertainty allow the return associated with illiquid tokens to change when aggregate risk rises. Similarly, interactions between momentum and cryptocurrency volatility allow momentum predictability to strengthen or weaken in turbulent markets.

The second measure is obtained from a two-state Markov-switching model (Hamilton, 1989) applied to log realized cryptocurrency-market volatility. The model identifies persistent low- and high-volatility states. We retain the posterior probability of the high-volatility state for each week and use it in the conditional forecasting analysis.

$$R_{m,t} = \mu_{S_t} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_{S_t}^2), \quad (6)$$

where $S_t \in \{1, 2\}$. We identify the states by their estimated conditional volatility and refer to them as the low-volatility and high-volatility regimes. We retain the posterior probability of each state for every week and use these probabilities in the conditional forecasting analysis.

This produces two related tests. The observable-state specification examines which market conditions change the relation between individual token characteristics and subsequent returns. The regime specification asks whether the relative forecasting performance of simple and nonlinear models differs systematically between tranquil and volatile cryptocurrency markets.

3.4 Return Prediction

We now turn to return prediction. Let $z_{i,t}$ denote the vector of characteristics of token i observed at time t . Expected one-week-ahead returns are written generally as

$$E_t(r_{i,t+1}^e) = g(z_{i,t}), \quad (7)$$

where the form of $g(\cdot)$ varies across models.

We consider models of increasing flexibility. Our first specification is ordinary least squares (OLS), which provides a linear benchmark using the full characteristic set. We next estimate Elastic Net (Zou & Hastie, 2005), allowing correlated predictors to be regularized and selected jointly. We then consider Random Forest (Breiman, 2001) and XGBoost (Chen & Guestrin, 2016) as two tree-based specifications. Finally, we estimate a shallow feed-forward neural network. We also construct a combination forecast as the equal-weighted average of the five individual predictions.

The sequence allows us to distinguish different forms of predictive improvement. OLS measures the information available from a common linear specification. Elastic Net allows the same information to be selected and regularized. Random Forest and XGBoost additionally allow nonlinearities and interactions, while the neural network provides a more flexible functional approximation. We do not impose an ordering on their expected performance.

We estimate all models using only information available before the forecast date. The initial training sample covers January 2018 to December 2020, and 2021 is used for model validation and hyperparameter selection. The first fully out-of-sample forecasts begin in January 2022. The estimation window subsequently expands through time and the models are re-estimated periodically through April 2026. A token listed after 2018 becomes part of the forecasting cross-section only after satisfying the minimum trading-history requirement.

We evaluate level forecasts using the out-of-sample coefficient of determination,

$$R_{OOS}^2 = 1 - \frac{\sum (r_{i,t+1} - \hat{r}_{i,t+1})^2}{\sum r_{i,t+1}^2}, \quad (8)$$

together with the mean absolute error and root mean squared error. We also calculate the weekly cross-sectional correlation between predicted and realized returns. The latter measures whether a model correctly ranks tokens even when the predicted level of returns is imprecise. Pairwise differences in forecast loss are evaluated using the Diebold–Mariano test (Diebold & Mariano, 1995).

We then examine where the forecasts obtain their information. For the nonlinear models, variable importance and SHAP (SHapley Additive exPlanations) values (Lundberg & Lee, 2017) are calculated for each characteristic. The analysis is conducted both at the individual-variable level and after aggregating characteristics into the five economic groups in Table 1. We further compare predictor importance across low- and high-volatility regimes. This allows us to determine whether a model's additional predictive ability is associated with a stable set of characteristics or with relationships that emerge only under particular market conditions. To avoid reducing model quality to a single statistic, the empirical interpretation therefore treats return-level accuracy, cross-sectional ranking, and portfolio performance as distinct but complementary dimensions of forecast usefulness.

3.5 Portfolio Tests

Finally, we examine whether the forecasts identify economically meaningful differences in subsequent token returns. Each week, we rank all eligible tokens by their predicted one-week-ahead return and assign them to five portfolios. The highest-prediction portfolio contains the top quintile and the lowest-prediction portfolio contains the bottom quintile. We form both equal-weighted and value-weighted portfolios and rebalance them weekly.

For each model, the long-short return is

$$R_{H-L,t} = R_{Q5,t} - R_{Q1,t}. \quad (9)$$

We report the mean returns of the five portfolios and examine whether realized performance changes monotonically with predicted returns. We then evaluate the long-short strategies using their mean return, volatility, Sharpe ratio, maximum drawdown, and factor-adjusted alpha. The latter is obtained from

$$R_{p,t}^e = \alpha_p + \beta_{MKT} CMKT_t + \beta_{SIZE} CSIZE_t + \beta_{MOM} CMOM_t + \epsilon_{p,t}. \quad (10)$$

Statistical inference for mean portfolio returns and factor-adjusted alphas is based on Newey–West (1987) heteroskedasticity- and autocorrelation-consistent standard errors.

The robustness design additionally specifies portfolio turnover, returns net of transaction costs, higher market-capitalization thresholds, and minimum portfolio-size requirements. These checks are important for determining whether apparent predictability is concentrated in small or difficult-to-trade tokens. Because the current draft does not contain the underlying weekly forecast panel or the full alternative-specification output, the final submission should report these sensitivity results explicitly rather than treating them as implied by the baseline estimates.

Our final tests split the out-of-sample period by market regime. For each prediction model, we separately calculate forecast accuracy, cross-sectional ranking ability, and portfolio performance in low- and high-volatility states. We then compare the performance gain of nonlinear models relative to OLS across the two regimes. The principal quantity of interest is

$$\Delta Perf_s = Perf_{Nonlinear,s} - Perf_{OLS,s}, \quad (11)$$

where s denotes the market state. A larger $\Delta Perf_s$ in the high-volatility regime would indicate that nonlinear structure becomes more informative when cryptocurrency-market conditions are less stable; similar performance across states would instead suggest that greater model flexibility adds little beyond a parsimonious return specification.

3.6 Identification, Validation, and Robustness Logic

The empirical strategy is designed to limit four threats that are especially relevant in machine-learning asset-pricing studies: look-ahead bias, survivorship bias, overfitting, and economic over-interpretation. Temporal ordering ensures that hyperparameter selection and model estimation use only information available before the forecast date. The point-in-time asset universe reduces survivorship bias by allowing tokens to enter and leave the sample as eligibility changes.

Cross-sectional standardization is performed within each formation week, preventing future information from entering predictor scaling.

Validation is multi-dimensional rather than metric-specific. Forecast-level performance is assessed with out-of-sample R^2 and forecast loss; relative asset ordering is evaluated using Pearson and Spearman information coefficients; economic implementation is evaluated through forecast-sorted portfolios with HAC inference. These tests are intentionally non-equivalent. Consistency across them would strengthen evidence of robust predictability, whereas divergence identifies the stage at which statistical information fails to translate into economic value.

The design also incorporates several robustness dimensions already embedded in the empirical protocol: alternative market-capitalization thresholds, weekly re-estimation through an expanding window, transaction-cost and turnover evaluation, factor-adjusted portfolio inference, and separate analysis across low- and high-volatility regimes. Because the current manuscript reports the primary specification most fully, robustness evidence should be interpreted as a validation framework rather than as a substitute for the central out-of-sample results.

Table 2. Model comparison and inferential role in the empirical design.

Model	Primary role	Functional flexibility	Main inferential risk	Observed pattern
OLS	Linear economic benchmark	Low	Missed nonlinearities/interactions	Moderate IC; negative OOS R^2
Elastic Net	Regularized linear benchmark	Low–moderate	Shrinkage can suppress weak signals	Better OOS R^2 than OLS; incomplete IC reporting
Random Forest	Nonlinear threshold/interactions	Moderate–high	Variance and local instability	Best OOS R^2 among models; strongest Spearman IC
XGBoost	Sequential nonlinear approximation	High	Residual overfitting in unstable samples	Weak IC and portfolio spread
Neural Network	Flexible function approximation	Very high	Bias–variance instability	Substantially weakest OOS R^2
Combination	Model uncertainty diversification	Mixed	Dilution of strong individual signals	Most significant state-dependent loss improvement

Threat	Why it matters	Mitigation in this study	Remaining limitation
Look-ahead bias	Inflates forecast performance	Temporally ordered train/validation/test split and one-day information gap	Data revisions and vendor histories may still matter
Survivorship bias	Favors persistent large assets	Dynamic point-in-time universe with entry and exit	Historical coverage is less complete for short-lived tokens
Overfitting	Complex models learn noise	Strict out-of-sample evaluation, simple benchmarks, expanding-window re-estimation	Hyperparameter uncertainty remains
Metric dependence	One metric can overstate usefulness	OOS R^2 , ICs, and portfolio tests evaluated separately	No single summary metric captures all objectives
Trading frictions	Paper profits may be untradeable	Turnover, transaction-cost logic, liquidity screens, HAC inference	Exchange-specific costs and short-sale constraints vary
Regime confounding	Volatility may proxy for other changes	Two-state volatility model interpreted as a boundary condition	Regime does not isolate a causal volatility channel

Table 3. Validity threats and design-based mitigation.

3.7 Pre-specified Robustness and Sensitivity Tests

A publication-level robustness package should test whether the headline conclusions survive reasonable changes in the sample, model, inference procedure, and implementation assumptions. The objective is not to search for a specification that produces significance, but to establish whether the qualitative ordering of models and the concentration of predictive information are stable under perturbations that are economically and statistically defensible.

First, sample sensitivity should be evaluated by re-estimating the full forecasting exercise under alternative minimum market-capitalization thresholds of USD 2 million, USD 5 million, and USD 10 million, alongside the USD 1 million baseline. The same exercise should impose minimum cross-sectional portfolio sizes. The relevant comparison is

whether model rankings, information coefficients, and forecast-sorted spreads remain directionally similar as the investable universe becomes more liquid and concentrated.

Second, predictor-group ablations should remove one economic block at a time—size/price, past returns, liquidity, risk, and return-distribution variables—and re-estimate the principal models. This provides a stronger test of incremental information than SHAP importance alone because it measures the deterioration in true out-of-sample performance when an information block is withheld.

Third, pairwise predictive-accuracy differences should be evaluated from the weekly forecast-loss series using the Diebold–Mariano test. Because Table 2 reports aggregate out-of-sample statistics rather than the underlying loss sequence, valid pairwise test statistics cannot be reconstructed from the manuscript alone. The final replication package should therefore report a complete model-by-model comparison matrix using the same out-of-sample weeks.

Fourth, economic robustness should be assessed using turnover-adjusted and transaction-cost-adjusted portfolio returns, including a break-even cost calculation. Reporting the transaction cost that reduces the high-minus-low strategy to zero is more informative than selecting one arbitrary cost assumption, particularly in cryptocurrency markets where trading frictions vary materially across assets and periods.

Fifth, explainability should be tested for stability rather than reported only as a single full-sample ranking. Regime-specific SHAP rankings, rank correlations between low- and high-volatility states, and dependence plots for the leading predictors would show whether the economic interpretation of the forecasting function is persistent or state contingent. These analyses require observation-level SHAP output and therefore cannot be recovered from the category shares currently reported in Table 4.

Taken together, these checks create a robustness ladder from design integrity to economic implementation. Figure 13 summarizes the sequence. Importantly, the items that require raw forecast errors, alternative-universe re-estimation, transaction-cost series, or observation-level SHAP values are specified here as required validation analyses; they are not presented as completed empirical findings in the absence of those underlying outputs.




Figure 2. Robustness Ladder for Final SSCI-Level Validation

Table 4. Robustness package required for final submission.

Test	Specification	Primary output	Data requirement	Current status
Sample thresholds	Re-estimate at \$2m, \$5m, \$10m market-cap cutoffs	Model ranks, OOS R^2 , ICs, H–L returns	Alternative-universe model outputs	Required before final submission
Predictor ablation	Drop one predictor block at a time	Δ OOS R^2 , Δ Spearman IC, Δ forecast loss	Re-estimated model outputs	Required before final submission

Test	Specification	Primary output	Data requirement	Current status
Pairwise accuracy	Diebold–Mariano on weekly loss differentials	DM statistic and p-value matrix	Weekly model forecast-loss series	Cannot be reconstructed from aggregate tables
Transaction costs	Net returns across a cost grid	Net H–L return and break-even cost	Weekly turnover and cost assumptions	Required before economic claim
SHAP stability	Compare regimes and dependence functions	Rank correlation, dependence shape	Observation-level SHAP values	Cannot be reconstructed from category shares
Temporal stability	Rolling / subperiod evaluation	Metric dispersion and rank persistence	Weekly OOS forecasts	Recommended

4. Empirical Results

4.1 Evolution of the Cryptocurrency Sample

We begin by examining how the investable cryptocurrency universe changes over time. The point-in-time sample is reconstructed separately at each historical snapshot rather than fixing the asset set using end-of-sample market capitalization. Figure 1 reports the evolution of the eligible universe, the distribution of asset histories, and market-capitalization concentration.




Figure 3. Evolution of the Point-in-Time Cryptocurrency Investment Universe, 2020-2025

The number of eligible assets remains broadly stable across annual snapshots but the composition of the sample changes materially. The universe contains 46 eligible assets in 2020, declines to 42 by 2023, and rises again to 46 in 2025. The relatively stable number of assets therefore masks substantial entry and exit within the cross-section.




Figure 4. Distribution of Cryptocurrency AssetHistorical Presence across Point-in-TimeSnapshots

The distribution of historical presence confirms this turnover. A large proportion of assets appear in only one annual snapshot, while a much smaller group remains continuously represented throughout the sample. Eighteen assets appear in all six snapshots. The remaining observations are distributed across assets with substantially shorter histories. This pattern produces an inherently unbalanced panel and illustrates the importance of preserving the contemporaneous investment universe rather than retrospectively selecting surviving assets.

Market concentration remains high throughout the period.




Figure 5. Market Capitalization Concentration acrossPoint-in-Time Cryptocurrency Snapshots

The largest asset accounts for approximately 78% of the capitalization of the eligible top-market-cap segment in 2020 and 2021. Its share falls sharply to approximately 53% in 2022 before rising to about 68% in 2025. The three largest assets account for between approximately 85% and 94% of market capitalization, while the corresponding share for the five largest assets remains above 90% throughout the sample.

The decline in concentration in 2022 coincides with a period in which several alternative cryptoassets attained substantial market capitalization. Concentration subsequently increases again as the market becomes more heavily weighted toward the largest assets. This variation implies that both value-weighted market returns and size-related characteristics change meaningfully through time.

The same structural change is visible in individual market rankings.




Figure 6. Point-in-Time Market Capitalization Rank Dynamics of Selected Cryptocurrencies.2020-2025

Bitcoin and Ethereum remain consistently at the top of the market, but the relative position of other assets changes substantially. Dogecoin rises from approximately rank 31 in 2020 to rank 7 by 2025. Solana enters the high-capitalization segment later but reaches approximately rank 5 in 2022, falls during 2023, and returns to rank 5 in 2024 and 2025. Litecoin follows the opposite pattern, falling from rank 6 in 2020 to approximately rank 21 by 2025. XRP, BNB, Cardano, Chainlink, and Polkadot also exhibit substantial changes in relative position.

These movements are accompanied by large changes in market capitalization.




Figure 7. Point-in-Time Market Capitalization Evolution of Selected Cryptocurrencies (USD Billion, Log Scale)

Bitcoin's point-in-time market capitalization increases from approximately USD 134 billion in 2020 to almost USD 1.95 trillion in 2025, while Ethereum rises from about USD 15 billion to approximately USD 438 billion. The expansion is not uniform across assets. Several tokens experience rapid increases followed by substantial contractions, particularly around the 2021–2023 cycle. The resulting cross-section therefore contains pronounced variation in scale, age, and market position over time.

Taken together, Figures 1–3 show that the cryptocurrency market is neither a fixed cross-section nor a stable ranking of assets. The number of eligible assets changes relatively little, but their identities, market shares, and relative capitalization change considerably. We next examine whether token characteristics constructed from this evolving market contain out-of-sample information about subsequent returns.

4.2 Out-of-Sample Return Prediction

Table 2 compares the out-of-sample performance of the competing forecasting models. The models differ in their degree of regularization and functional flexibility, ranging from OLS and Elastic Net to Random Forest, XGBoost, a shallow neural network, and a forecast combination.

Table 5. Out-of-Sample Forecasting Performance

Model	OOS R ²	Pearson IC	Spearman IC
OLS	-1.28%	0.092	0.088
Elastic Net	-0.63%	—	—
Random Forest	-0.29%	0.075	0.116
XGBoost	-0.84%	0.005	0.020
Neural Network	-18.76%	0.051	0.058
Combination	-1.04%	0.073	0.085

The first result is that one-week-ahead cryptocurrency returns remain difficult to predict in levels. The out-of-sample R² is negative for every specification. Random Forest performs best at −0.29%, followed by Elastic Net at −0.63%. XGBoost records −0.84%, while OLS produces −1.28%. The neural network performs substantially worse, with an out-of-sample R² of −18.76%.




Figure 8. Out-of-Sample Return Forecast Accuracy across Competing Models

The differences among the linear, regularized, and tree-based models are modest relative to the deterioration observed for the neural network. Increasing functional flexibility therefore does not produce a monotonic improvement in forecast performance. Random Forest performs better than the linear benchmark, but greater complexity beyond this point does not systematically improve accuracy.

The ranking results provide a different ordering.




Figure 9. Cross-Sectional Ranking Performance: Pearson and Spearman Information Coefficients

OLS generates the highest Pearson information coefficient at 0.092, whereas Random Forest produces the highest Spearman rank correlation at 0.116. The corresponding correlations for XGBoost are substantially smaller, at 0.005 and 0.020. The forecast combination performs between the linear and tree-based alternatives.

The divergence between the two evaluation criteria is important. A model can perform poorly when predicting the level of realized returns yet retain information about the relative ordering of assets. Random Forest illustrates this distinction: its out-of-sample R^2 remains negative, but its rank correlation is the strongest among the models. By contrast, XGBoost does not translate its greater functional flexibility into either stronger level forecasts or stronger cross-sectional rankings.

The results therefore indicate that the relevant forecasting problem is not simply whether a model can reproduce realized return magnitudes. Cross-sectional differentiation provides an additional dimension of predictive performance.

4.3 Economic Value of Forecast-Sorted Portfolios

We next examine whether differences in predicted returns translate into differences in realized portfolio performance. Assets are ranked according to their one-week-ahead return forecasts, and the high-minus-low spread compares assets with the strongest predicted returns with those receiving the weakest forecasts.

Table 6. Forecast-Sorted Portfolio Performance

Model	Weekly H-L return	HAC t-statistic
OLS	0.03%	0.05
Elastic Net	-0.29%	-0.59
Random Forest	0.70%	1.18
XGBoost	-0.02%	-0.03
Neural Network	-0.11%	-0.20
Combination	0.35%	0.59

Random Forest generates the largest spread, approximately 0.70% per week. The forecast combination produces a smaller positive spread of 0.35%, while the OLS portfolio is close to zero. Elastic Net, XGBoost, and the neural network generate negative or economically negligible spreads.




Figure 10. Economic Value of Forecast-Sorted Portfolios across Models

The Random Forest spread is positive but imprecisely estimated. Its HAC t-statistic is 1.18 and therefore does not reach conventional levels of statistical significance. The portfolio results consequently provide weaker evidence than the ranking statistics.

The contrast between Random Forest and XGBoost is particularly informative. Both specifications allow nonlinear relationships, yet their portfolio results differ sharply. Random Forest produces the strongest rank correlation and the largest portfolio spread, whereas XGBoost generates almost no high-minus-low return. The economic value of model flexibility therefore depends on the form of the nonlinear approximation rather than on complexity alone.

More generally, the results separate three dimensions of forecasting performance. Out-of-sample R^2 measures the accuracy of predicted return levels, information coefficients evaluate cross-sectional ordering, and forecast-sorted portfolios measure whether this ordering translates into economic return differences. The model rankings are not identical across the three measures.

4.4 Sources of Predictive Information

We next examine which characteristics contribute most strongly to the predictions of the tree-based models. Figure 7 reports mean absolute SHAP values aggregated into economic categories.

Table 7. Predictor Importance by Economic Category

Predictor category	Random Forest	XGBoost
Risk	37.8%	32.3%
Return distribution	25.3%	26.0%
Past returns	24.7%	26.7%
Liquidity	8.7%	12.3%
Price	3.5%	2.8%

Risk characteristics account for the largest share of predictive importance in both models. Their contribution reaches 37.8% in Random Forest and 32.3% in XGBoost. Past-return and return-distribution variables jointly account for approximately half of total importance, whereas liquidity contributes between 8.7% and 12.3%. Nominal price contributes less than 4%.




Figure 11. Economic Predictor Importance from SHAPDecomposition

The concentration of variable importance is notable. The nonlinear models do not distribute predictive weight evenly across the full characteristic set. Instead, most of the forecasting signal is concentrated in a relatively small group of variables associated with recent risk, the distribution of returns, and past performance.

Within Random Forest, the most important individual characteristics include longer-horizon realized volatility, market beta, skewness, short-horizon momentum, and kurtosis. These variables capture different dimensions of recent risk rather than a single directional price signal.

This pattern also explains why a characteristic may appear weak in a univariate return sort but remain important in a nonlinear model. Its predictive contribution can depend on the values of other characteristics. Volatility, for example, need not have a stable monotonic relation with subsequent returns if its effect differs across levels of beta, skewness, or recent momentum. Tree-based models can capture such conditional relationships without imposing a common coefficient throughout the predictor space.

The relative importance of liquidity is smaller than that of risk and past-return information. This suggests that, within the current matched sample, the predictive structure is driven more heavily by differences in recent risk conditions than by variation in trading activity alone.

4.5 Market Volatility and Relative Model Performance

We next examine whether the relative performance of nonlinear forecasts changes with aggregate market conditions. A two-state Markov-switching specification applied to realized cryptocurrency-market volatility identifies a low- and a high-volatility state.

The two states are economically distinct. Mean realized market volatility is approximately 3.7% in the low-volatility state and 8.1% in the high-volatility state. The high-volatility regime is persistent, with an estimated transition probability

$$P(S_t = High | S_{t-1} = High) = 0.912, \quad (12)$$

while the corresponding probability for the low-volatility state is approximately


$$P(S_t = Low | S_{t-1} = Low) = 0.699. \quad (13)$$

We then relate the posterior probability of the high-volatility state to each model's forecast loss relative to OLS.

Table 8. High-Volatility Probability and Relative Forecast Loss

Model	Coefficient	t-statistic	p-value
Random Forest	-0.000260	-1.94	0.053
XGBoost	-0.000512	-1.87	0.062
Combination	-0.000331	-2.12	0.034
Neural Network	-0.000880	-0.77	0.443

The coefficients for Random Forest, XGBoost, and the forecast combination are negative. Because the dependent variable measures model forecast loss relative to OLS, the negative coefficients indicate that these models perform relatively better as the probability of the high-volatility state rises.

**Figure 12.** State-Dependent Forecast Performance under Cryptocurrency Market Volatility Regimes

The effect is statistically significant for the combination forecast at the 5% level. Random Forest and XGBoost display similar negative relations, with p-values of 0.053 and 0.062, respectively. The neural-network coefficient is statistically insignificant.

The result changes the interpretation of the unconditional model comparison. Average forecasting differences are small, but relative model performance is not constant through time. The advantage of the tree-based and combination forecasts becomes more favorable as market volatility increases.

The same pattern is not equally strong when performance is measured using cross-sectional information coefficients. The state-dependent relation is therefore concentrated primarily in forecast loss rather than in asset ranking. Higher market volatility changes the relative ability of the models to estimate return magnitudes, but it does not produce an equivalent improvement in their ordering of subsequent asset returns.

This distinction is consistent with the earlier separation between level accuracy and cross-sectional performance. Market conditions affect the two dimensions of forecasting differently.

4.6 Integrated Multi-Criterion Assessment

The preceding tests evaluate different dimensions of predictive performance separately. A stronger interpretation emerges when the metrics are considered jointly rather than treated as interchangeable measures of model quality. Figure 9 summarizes the evidence from Tables 2 and 3 in a common matrix. Cell labels report the original statistics, while the background scale standardizes each metric within its own column solely to facilitate visual comparison across models. The standardized shading is descriptive and is not used for statistical inference.




Figure 13. Integrated Evaluation Matrix of Forecast Accuracy, Ranking Ability, and Economic Value.

The integrated comparison reinforces the absence of a universally dominant model. Random Forest is the most balanced specification across the reported criteria: it has the least-negative out-of-sample R^2 , the highest Spearman information coefficient, and the largest positive high-minus-low return. OLS remains competitive on Pearson correlation, while the neural network is clearly weak on return-level accuracy. The combination forecast occupies an intermediate position and becomes more informative once market-state dependence is considered.




Figure 14. Forest Plot of the Relationship between High-Volatility-State Probability and Relative Forecast Loss.

Figure 10 makes the uncertainty around the regime estimates explicit. The combination forecast is the only specification whose reported p-value falls below 0.05, whereas the Random Forest and XGBoost estimates are directionally similar but borderline at conventional thresholds. The neural-network estimate is considerably less precise. The figure therefore supports a cautious state-dependent interpretation: evidence that model flexibility becomes more useful in volatile markets is strongest for forecast combination and suggestive, rather than definitive, for the individual tree models.




Figure 15. Multi-Objective Comparison between Return-Level Accuracy and Cross-Sectional Ranking Performance

The multi-objective view clarifies why model choice should depend on the intended use of the forecast. Random Forest combines the strongest rank signal with the least-negative return-level fit and the largest positive portfolio spread, but its economic result remains statistically imprecise. OLS ranks reasonably well despite a weaker out-of-sample R^2 , and the combination forecast offers moderate performance across dimensions. These patterns support evaluating forecasting systems along a performance frontier rather than selecting models from a single headline metric.

4.7 Robustness Evidence and the Boundary of the Current Results

The current results already provide several forms of triangulation. Model conclusions are compared across return-level fit, cross-sectional rank correlation, forecast-sorted portfolio spreads, predictor importance, and volatility-state dependence. The fact that these measures do not produce an identical model ranking is informative rather than problematic: it shows that predictive accuracy, ranking ability, and economic implementation are distinct empirical objects.

At the same time, the evidence should not be described as a complete robustness exercise. The manuscript reports the baseline forecasting table, portfolio spreads, category-level SHAP importance, and state-dependent loss regressions, but it does not report the full outputs from alternative capitalization screens, cost grids, predictor ablations, weekly pairwise forecast-loss tests, or regime-specific SHAP dependence analyses. Those tests require underlying data objects that are not contained in the current document.

Accordingly, the strongest defensible conclusion at this stage is that the main findings are internally consistent across several non-equivalent evaluation criteria, while external sensitivity to investable-universe definitions, trading costs, predictor blocks, and explainability stability remains to be documented in the final replication package. This distinction strengthens the credibility of the paper by separating reported evidence from analyses that are specified but not yet observable in the manuscript.

4.8 Summary of Empirical Findings

The results reveal substantial structural change in the cryptocurrency cross-section. The number of eligible assets remains relatively stable across annual snapshots, but the identity, market capitalization, and ranking of individual assets change materially. Market concentration also varies through time, with the dominance of the largest cryptoasset falling sharply during 2022 before increasing again.

Against this changing market structure, future cryptocurrency returns remain difficult to predict in levels. All models generate negative out-of-sample R^2 . Random Forest produces the smallest forecasting error, whereas the neural network performs substantially worse. Greater model complexity therefore does not produce a uniform improvement in predictive accuracy.

Cross-sectional ranking provides additional information. Random Forest generates the highest Spearman information coefficient and the largest forecast-sorted return spread. However, the corresponding portfolio return is estimated imprecisely and does not reach conventional levels of statistical significance.

The predictive information of the nonlinear models is concentrated primarily in risk, return-distribution, and past-return characteristics. Liquidity and price contribute smaller shares. The results therefore point toward a relatively concentrated predictive structure rather than a diffuse collection of independent signals.

Finally, the relative performance of the nonlinear and combination forecasts varies with aggregate market volatility. Random Forest, XGBoost, and the combination forecast improve relative to OLS as the probability of the high-volatility state increases. This improvement is clearer in forecast loss than in cross-sectional ranking.

Overall, the results, including the integrated comparison in Figures 9-11, distinguish three separate empirical dimensions:

- return-level predictability
- cross-sectional ranking
- and economic portfolio performance

The three dimensions are related, but they do not move together across models or market states.

5. Discussion

5.1 A Changing Cross-Section and the Stability of Return Predictability

The evolution of the point-in-time universe changes how the return-prediction results should be interpreted. Although the number of eligible cryptoassets varies relatively little between 2020 and 2025, the identities and market weights of those assets change considerably. Bitcoin and Ethereum remain dominant, but the rise of Solana and Dogecoin, the relative decline of Litecoin, and the substantial movements in XRP, BNB, Cardano, and other assets show that the cross-section is repeatedly reconstituted. Market concentration changes at the same time: the share represented by the largest asset falls sharply in 2022 and subsequently increases again. The forecasting problem is therefore not based on a stable set of securities whose characteristics merely fluctuate through time. The securities themselves, and their relative positions within the market, change.

This feature matters because a characteristic can retain the same statistical definition while its economic content changes. A token classified as relatively small in a mature and concentrated market is not necessarily comparable with a similarly ranked token during a period of rapid market entry. Likewise, high trading volume can indicate established liquidity for one asset but speculative turnover for another. Cross-sectional standardisation controls differences in measurement scale, but it cannot make the economic meaning of those observations invariant through time. Consequently, a stable coefficient on size, momentum, or liquidity is a stronger assumption in cryptocurrencies than it may initially appear.

The three-factor evidence of [Liu et al. \(2022\)](#) shows that market, size, and momentum capture a substantial proportion of the cross-section of cryptocurrency expected returns. Importantly, their result is one of parsimony: several apparently different strategies can be accounted for by a small set of common dimensions. The present evidence adds a temporal qualification to that interpretation. A parsimonious factor structure can coexist with unstable factor loadings. Indeed, Liu et al. themselves find that momentum loadings appear more transient than size loadings. Thus, the relevant question is not only whether a characteristic belongs to the return-generating structure, but whether its contribution remains stable as the cryptocurrency universe changes.

This distinction can be expressed as

$$E_t(r_{i,t+1}) = g_t(z_{i,t}), \quad (14)$$

rather than assuming a time-invariant mapping $g(\cdot)$. Both the distribution of characteristics $z_{i,t}$ and the relationship between those characteristics and subsequent returns can evolve. In this setting, model flexibility has potential value not because cryptocurrencies necessarily require a very high-dimensional pricing function, but because the same small set of predictors may operate differently across market environments.

5.2 Model Complexity and the Effective Dimension of Cryptocurrency Returns

The out-of-sample results provide little support for a simple relation between computational complexity and predictive performance. Random Forest produces the smallest forecast error, but the difference from OLS is modest. XGBoost does not improve further, while the neural network performs substantially worse. The ordering of the models therefore differs from the conventional argument that increasingly flexible functions should progressively uncover more return predictability.

One explanation is that the effective dimension of cryptocurrency expected returns is considerably lower than the observed number of characteristics. The dataset contains information on past returns, risk, liquidity, price, and distributional properties, but these variables need not represent independent sources of expected returns. Volatility, beta, skewness, extreme returns, and momentum may partly describe different manifestations of a smaller number of underlying market conditions. In that case, increasing the number of nonlinear interactions expands the estimation problem faster than it expands economically relevant information.

This interpretation is consistent with the evidence of [Cakici et al. \(2024\)](#). Despite using 40 cryptocurrency characteristics and a wide range of machine-learning methods, they find that model complexity provides limited incremental value and that predictive information is concentrated in a handful of comparatively simple variables. They offer several possible explanations, including a lower-dimensional crypto return structure, rapid changes in the cryptocurrency universe, and insufficient historical information to estimate complex nonlinear patterns reliably. The present model ordering is consistent with all three possibilities, but it also provides a useful distinction among different forms of complexity.

Random Forest and XGBoost are both tree-based methods, yet their results are materially different. Random Forest relies on averaging across many partially independent trees. This aggregation can reduce the influence of individual noisy relationships. Boosting, by contrast, sequentially attempts to correct preceding forecast errors. In a market containing extreme observations and unstable relationships, such residual fitting can also place greater weight on patterns that do not persist out of sample. The weak XGBoost result therefore does not indicate an absence of nonlinear structure; rather, it suggests that the amount of nonlinear structure that can be estimated reliably is limited.

The neural-network result strengthens this interpretation. Its substantially poorer out-of-sample R^2 is difficult to reconcile with the proposition that the principal obstacle is insufficient functional flexibility. Instead, it points toward a bias–variance problem. Cryptocurrency returns provide a noisy target, the cross-section changes rapidly, and the available historical sample is short relative to conventional equity datasets. Increasing flexibility under these conditions can improve in-sample approximation while worsening estimation of the persistent component of expected returns.

Cakici et al. reach a closely related conclusion when contrasting cryptocurrency and equity markets: nonlinearities and interactions appear less beneficial in a rapidly changing cryptocurrency universe than they do in stocks. The relevant implication is not that nonlinear modelling is unnecessary. It is that model complexity must be justified by incremental out-of-sample information, rather than by the capacity of the algorithm alone.

5.3 From Individual Characteristics to Conditional Predictability

The variable-importance results reveal a second distinction. Risk characteristics account for the largest share of SHAP importance, followed by past returns and return-distribution variables. Yet several of these characteristics do not necessarily produce strong unconditional return spreads when examined individually. The apparent discrepancy between weak univariate effects and substantial multivariate importance is informative.

A univariate portfolio sort implicitly asks whether

$$E(r_{i,t+1}|z_{i,t}) \tag{15}$$

changes monotonically with one characteristic. A nonlinear prediction model instead permits

$$E(r_{i,t+1}|z_{i,t}, x_{i,t}) \tag{16}$$

to change with combinations of characteristics. Thus, volatility can contain little information on its own while becoming informative conditional on beta, skewness, momentum, or other aspects of the recent return distribution.

For example, high realized volatility need not have one constant implication for expected returns. A high-volatility asset with positive momentum and limited downside asymmetry can represent a different state from a high-volatility asset dominated by extreme negative observations. A linear coefficient averages these observations. A tree model can separate them. The predictive importance assigned to volatility therefore does not require volatility to generate a strong unconditional premium.

This distinction also prevents an overly strong asset-pricing interpretation of SHAP values. High importance does not imply that volatility, beta, or skewness constitutes a priced risk factor. SHAP measures the contribution of a variable to the model's forecasts; it does not establish that investors receive compensation for bearing the associated risk. The evidence supports conditional predictive relevance, not a causal or equilibrium pricing relation.

The composition of predictor importance also differs from the broad-universe evidence of Cakici et al., where price, past alpha, momentum, and illiquidity play particularly important roles and past returns and liquidity dominate at the category level. In the present results, risk and return-distribution characteristics are more prominent. This difference is economically meaningful. Cakici et al. show that their strongest machine-learning alphas are concentrated among difficult-to-trade, illiquid, and volatile cryptocurrencies. A more concentrated high-market-cap prediction sample naturally contains less extreme variation in some liquidity measures. Once part of this variation is compressed, differences in risk and return distributions can become relatively more informative.

The evidence therefore suggests that predictor importance depends not only on the forecasting algorithm, but also on which segment of the cryptocurrency market is being priced. There is no reason to expect a single importance ranking to apply equally to mature high-capitalisation assets and newly entering, thinly traded tokens.

5.4 What the Volatility Regime Result Actually Identifies

The regime results provide the clearest evidence that the return-prediction function is time varying. Random Forest, XGBoost, and the forecast combination improve relative to OLS as the posterior probability of the high-volatility state increases. The effect is strongest for the combination forecast and is close to conventional significance levels for the two tree-based models.

The immediate interpretation might be that nonlinear prediction becomes more valuable during turbulent markets. The comparison between forecast loss and information coefficients, however, allows a more precise conclusion. The improvement appears primarily in return-level forecast loss. A comparable improvement is not found in cross-sectional ranking.

This distinction matters because the high-volatility state is a common market condition. All assets experience the same aggregate regime at a given date. If greater model flexibility improves return-level predictions but does not materially improve the ordering of assets, the models may be adapting more successfully to changes in the common component of the return distribution rather than uncovering stronger differences between individual tokens.

Formally, returns can be separated into a market component and a relative component,

$$r_{i,t+1} = \beta_{i,t} r_{m,t+1} + u_{i,t+1}. \quad (17)$$

An aggregate volatility regime can materially alter the first term even if the ranking information contained in $u_{i,t+1}$ changes much less. The observed pattern—improved forecast loss without a comparable increase in IC—is consistent with this possibility. Nonlinear models may better accommodate how beta, recent volatility, and tail characteristics interact with a changing aggregate risk environment, while gaining less information about which individual asset will outperform its peers.

This interpretation also clarifies why the unconditional differences among models are modest. Low- and high-volatility periods are pooled in the average OOS statistics. If nonlinear flexibility matters primarily in one state, the full-sample comparison dilutes its value. Conversely, selecting the most complex model solely because it performs relatively well during extreme states would sacrifice the advantage of simpler specifications during more stable periods.

The performance of the forecast combination is informative in this respect. Its relative forecast-loss advantage strengthens significantly as high-volatility probability rises, even though no individual model dominates across every

performance measure. Combining forecasts reduces dependence on a single functional form precisely when the mapping between characteristics and returns is least stable. The result is consistent with a forecasting environment in which model uncertainty increases together with market uncertainty.

This does not establish that volatility causes nonlinear predictability. The regime is an estimated description of the market state, and several economic changes—liquidity conditions, investor attention, leverage, risk tolerance, or market composition—may occur simultaneously. The evidence instead identifies a boundary condition: the relative usefulness of model flexibility is systematically related to aggregate cryptocurrency-market volatility.

5.5 Statistical Predictability Does Not Imply Economic Exploitability

The portfolio results place an important restriction on the economic interpretation. Random Forest produces the strongest cross-sectional ranking and a high-minus-low return of approximately 0.70% per week, but the corresponding HAC t-statistic is only 1.18. The magnitude is economically noticeable, but the estimate is too imprecise to support a claim of reliably exploitable abnormal performance.

This result differs from the broad-universe evidence of Cakici et al., where weak predictive R^2 coexists with strong portfolio profitability across models. Their portfolios successfully separate winners from losers even when level forecasts appear statistically disappointing. They further show that much of this profitability is concentrated in assets characterised by high illiquidity, idiosyncratic risk, bid-ask spreads, and other limits to arbitrage.

The weaker portfolio evidence here therefore contains information rather than simply representing a failed prediction exercise. A market-capitalisation-based sample places greater weight on assets for which trading is easier and arbitrage capacity is generally greater. If some of the strongest cryptocurrency anomalies survive because they occur in assets that are costly or difficult to trade, removing much of that segment should reduce the observed portfolio premium. The present pattern is consistent with such an interpretation.

However, a pure small-asset explanation would also be too simple. Liu et al. (2022), for example, find that their momentum strategy is substantially stronger among larger cryptocurrencies than among the below-median-size group. Different signals therefore need not share the same relation with size or liquidity. Momentum can be stronger among larger coins while broader machine-learning alphas remain concentrated among hard-to-arbitrage assets. This reinforces the view that cryptocurrency predictability is composed of several conditional relationships rather than a single small-coin premium.

The distinction between the three performance measures is consequently substantive:

- R_{00S}^2 asks whether the magnitude of future returns is forecast accurately;
- IC asks whether assets are ordered correctly;
- and R_{H-L} asks whether that ordering produces an economically meaningful payoff.

Random Forest performs relatively well on the latter two dimensions but not sufficiently strongly to establish statistically reliable portfolio profits. The regime results add another layer: nonlinear methods improve relative forecast loss as volatility increases, but the same state dependence is not clearly transmitted to ranking. Thus,

$$\text{forecast improvement} \neq \text{ranking improvement} \neq \text{arbitrage profit.} \quad (18)$$

This separation is particularly important in cryptocurrency research because extreme returns can make economic gains appear large even when they are generated by a small number of observations. Cakici et al. explicitly show that profitability in broad cryptocurrency samples is critically dependent on extreme returns and on small, illiquid, volatile assets. Economic significance therefore requires evidence not merely of positive portfolio returns, but of persistence after diversification, trading constraints, and the removal of extreme observations.

5.6 Implications for Cryptocurrency Asset Pricing and Forecasting

The empirical evidence points toward a return structure that is parsimonious in its information content but unstable in its conditional mapping. These two features are not contradictory.

The concentration of SHAP importance suggests that only a limited set of characteristics carries substantial predictive information. This accords with the broader evidence that apparently large cryptocurrency factor sets can be compressed into a smaller number of return dimensions. Liu et al. show this directly through their market-size-momentum model.

At the same time, the changing point-in-time universe and the regime results indicate that the way these dimensions translate into future returns is not constant.

The resulting forecasting problem is therefore different from both a fixed linear factor model and an unrestricted black-box prediction problem. A fixed linear model can be too rigid when the role of risk, momentum, or liquidity changes with market conditions. An excessively flexible model, however, can estimate transitory relationships that disappear as the asset universe and volatility state change. The superior relative performance of Random Forest compared with the neural network is consistent with a middle ground: enough flexibility to capture thresholds and interactions, but sufficient regularisation to avoid treating every historical pattern as persistent.

This interpretation also changes the role of machine learning in the analysis. Its value is not primarily that it generates a larger number of predictors. Instead, it provides a way to investigate whether known characteristics have state-dependent and interaction-dependent predictive content. The asset-pricing question consequently shifts from “Which variable predicts returns?” to “Under what market conditions does that variable contain predictive information?”

That distinction is central for a market whose constituent assets, concentration, volatility, and trading environment change rapidly.

The results also suggest that average model comparisons are insufficient for evaluating cryptocurrency forecasts. A model that performs only marginally better over the full sample can provide more information in particular market states, while a model with better level forecasts need not improve portfolio selection. Forecasting performance should therefore be evaluated jointly across time, cross-section, and economic outcomes, rather than reduced to a single R^2 or Sharpe ratio.

Finally, the evidence places a boundary on claims about market efficiency. The results show some predictability, but they do not demonstrate an easily exploitable market inefficiency. The strongest portfolio spread remains statistically imprecise, and the improvement of nonlinear methods in high-volatility states is concentrated in forecast error rather than portfolio ranking. A more defensible interpretation is that cryptocurrency expected returns contain conditional structure that simple unconditional models do not fully capture. Whether that structure represents compensation for risk, temporary mispricing, limits to arbitrage, or changing market participation cannot be identified from predictive performance alone.

The combination of a dynamic asset universe, concentrated predictor importance, limited average gains from complexity, and state-dependent forecast performance therefore points toward a market in which the main challenge is not discovering ever more signals. It is identifying which relationships persist, when they change, and whether the information they contain survives the transition from statistical prediction to economic implementation.

5.7 Theoretical, Methodological, and Practical Implications

Theoretically, the results support a view of cryptocurrency expected returns that is low-dimensional in information content but unstable in conditional mapping. The concentration of SHAP importance suggests that a relatively small set of risk, past-return, and distributional characteristics carries most of the predictive signal, while regime dependence indicates that the mapping from those characteristics to future returns changes with market conditions. This combination reconciles parsimonious factor evidence with selective benefits from nonlinear models.

Methodologically, the findings caution against equating algorithmic sophistication with empirical contribution. The neural network's weak out-of-sample performance and the limited advantage of XGBoost show that additional functional capacity can increase estimation error when the target is noisy and the cross-section evolves rapidly. In this setting, the relevant benchmark is not in-sample fit but incremental out-of-sample information relative to transparent economic models.

For empirical asset-pricing research, the multi-criterion evaluation framework is itself consequential. A forecasting model can add information at the level of return magnitudes, rankings, or portfolio outcomes, but evidence at one level does not logically imply evidence at the others. Reporting all three prevents statistical improvements from being overstated as economically exploitable anomalies.

For practitioners, the results favor adaptive model governance rather than permanent commitment to the most complex specification. Simpler models remain competitive in average conditions, while tree-based or combined

forecasts appear relatively more useful as market volatility rises. A practical implication is that model selection may be conditioned on market state and monitored using out-of-sample diagnostics, rather than fixed ex ante on the basis of computational sophistication alone.

6. Conclusion

This study examines whether cryptocurrency return predictability can be improved by combining a dynamic point-in-time asset universe with linear and nonlinear forecasting models. The central finding is that the main forecasting challenge is not a lack of model complexity, but the instability of the relationship between token characteristics and future returns.

The evidence points to a relatively concentrated information structure. Risk, recent returns, and return-distribution characteristics contain most of the predictive information, while adding greater functional flexibility does not consistently improve forecasts. Random Forest performs more favorably than the other nonlinear specifications, whereas deeper model complexity provides little additional benefit. The relevant advantage of nonlinear methods is therefore selective rather than general.

More importantly, predictive relationships vary with market conditions. Nonlinear and combined forecasts become relatively more useful as the probability of a high-volatility market state increases. This suggests that cryptocurrency return predictability is better represented by a time-varying conditional structure than by a single stable forecasting function. The same predictors may remain relevant across time, while their relative importance and interactions change with the market environment.

The results also place a clear boundary on the economic interpretation of forecast improvements. Better return-level forecasts do not necessarily generate better cross-sectional rankings, and stronger rankings do not automatically translate into statistically reliable portfolio profits. Predictive accuracy, asset selection, and economic value should therefore be treated as separate empirical questions. The integrated multi-criterion comparison further shows that no specification dominates all objectives simultaneously, making the choice of model contingent on whether the research objective is return-level forecasting, asset ranking, or portfolio implementation.

The study contributes to cryptocurrency asset-pricing research by shifting attention from whether machine learning can outperform conventional models on average to when additional model flexibility becomes informative. The evidence suggests that a useful forecasting framework should combine parsimonious economic predictors with sufficient flexibility to accommodate changing market states, rather than relying either on fixed linear relationships or unrestricted model complexity.

Several limitations remain. Cryptocurrency markets continue to evolve rapidly, and the economic interpretation of characteristics can change as new assets enter and older assets lose market relevance. Future research should therefore examine whether the state-dependent patterns remain stable across broader token universes, alternative liquidity screens, transaction-cost assumptions, and longer out-of-sample periods. Further work is also required to distinguish predictive associations from underlying economic mechanisms.

The broader implication is straightforward: cryptocurrency return prediction is primarily a problem of identifying stable information within an unstable market structure. Model sophistication matters only when it captures that instability rather than merely adding complexity.

7. Limitations

This study is subject to several limitations.

First, the reconstruction of the cryptocurrency universe depends on the availability of historical point-in-time market-capitalization and price data. Although the dynamic sample reduces survivorship and look-ahead biases, historical coverage is not equally complete for all assets. Newly listed, delisted, renamed, or migrated tokens are more likely to have discontinuous observations. The matched sample may therefore place relatively greater weight on larger and more continuously traded cryptoassets.

Second, the cryptocurrency market changes rapidly. Exchange structure, regulation, investor participation, liquidity, and the composition of traded assets all evolved materially during the sample period. Relationships estimated from

earlier observations may therefore become less informative in later periods. This structural instability is particularly relevant for nonlinear models, which can estimate complex historical interactions that may not persist out of sample.

Third, the analysis identifies predictive relationships rather than causal mechanisms. Variable-importance and SHAP measures indicate which characteristics contribute to forecasts, while the regime analysis shows how model performance varies with aggregate volatility. Neither establishes that volatility, momentum, liquidity, or other characteristics causally determine expected returns. Further identification would require complementary designs such as event studies, natural experiments, or exogenous shocks.

Fourth, economic performance remains sensitive to implementation assumptions. Weekly portfolio formation reduces turnover relative to daily strategies, but actual profitability can still depend on transaction costs, bid-ask spreads, exchange availability, short-selling constraints, and market depth. These frictions are especially important for smaller and less liquid tokens. Consequently, statistically predictable returns should not be interpreted automatically as fully exploitable trading opportunities.

Finally, the two-state volatility regime provides a parsimonious description of changing market conditions but cannot capture every form of structural change. Cryptocurrency markets may exhibit more gradual transitions or multiple states associated with liquidity, investor attention, leverage, or regulatory shocks. Future research could therefore extend the analysis using richer state variables, alternative regime specifications, and longer point-in-time datasets.

Comparing many model variants creates an additional model-selection risk. White's (2000) reality check and Hansen's (2005) superior predictive ability test provide approaches for evaluating selected forecasts while accounting for data snooping. These are useful extensions for a future robustness analysis; no results from either procedure are claimed here.

These limitations do not alter the distinction established by the empirical analysis between forecast accuracy, cross-sectional ranking, and economic performance. They instead define the conditions under which the results should be interpreted and provide directions for testing whether the identified predictive relationships persist in a broader and more mature cryptocurrency market.

Data Availability Statement

The analysis uses historical cryptocurrency market data and publicly observable market-state variables described in Section 3. The current manuscript does not identify a citable archival package. Before publication, I report the exact data vendors or APIs, access dates, token-identity mapping rules, and licensing restrictions, and should archive all redistributable derived data needed to reconstruct the estimation sample.

Code Availability Statement

For full reproducibility, the publication package archive code for data screening, characteristic construction, model estimation, hyperparameter selection, out-of-sample forecasting, SHAP analysis, regime estimation, portfolio formation, robustness tests, and figure production. The archive should also include environment specifications, random seeds, and saved out-of-sample predictions so that the reported tables and pairwise forecast comparisons can be regenerated exactly.

Acknowledgments

This work has not received any funding support.

Conflict of Interest

I disclose any financial or non-financial interests that could be perceived as influencing the research. If none exist, the final accepted manuscript should state that the authors declare no competing interests.

Ethics Statement

The study analyzes market-level and asset-level financial data and does not involve human participants, personal data, or animal subjects. Institutional ethics approval is therefore generally not applicable, subject to the journal's editorial policy.

References

- Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1), 31–56. [https://doi.org/10.1016/S1386-4181\(01\)00024-6](https://doi.org/10.1016/S1386-4181(01)00024-6)
- Baur, D. G., Hong, K., & Lee, A. D. (2018). Bitcoin: Medium of exchange or speculative assets? *Journal of International Financial Markets, Institutions and Money*, 54, 177–189. <https://doi.org/10.1016/j.intfin.2017.12.004>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Cakici, N., Shahzad, S. J. H., Będowska-Sójka, B., & Zaremba, A. (2024). Machine learning and the cross-section of cryptocurrency returns. *International Review of Financial Analysis*, 94, 103244. <https://doi.org/10.1016/j.irfa.2024.103244>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794. Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939785>
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263. <https://doi.org/10.1080/07350015.1995.10524599>
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223–2273. <https://doi.org/10.1093/rfs/hhaa009>
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384. <https://doi.org/10.2307/1912559>
- Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business & Economic Statistics*, 23(4), 365–380. <https://doi.org/10.1198/073500105000000063>
- Liu, Y., & Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *The Review of Financial Studies*, 34(6), 2689–2727. <https://doi.org/10.1093/rfs/hhaa113>
- Liu, Y., Tsyvinski, A., & Wu, X. (2022). Common risk factors in cryptocurrency. *The Journal of Finance*, 77(2), 1133–1177. <https://doi.org/10.1111/jofi.13119>
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In I. Guyon, U. von Luxburg, S. Bengio, H. Wallach, R. Fergus, V. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* 30 (pp. 4765–4774). Curran Associates, Inc. <https://arxiv.org/abs/1705.07874>
- Makarov, I., & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2), 293–319. <https://doi.org/10.1016/j.jfineco.2019.07.001>
- Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708. <https://doi.org/10.2307/1913610>
- Welch, I., & Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4), 1455–1508. <https://doi.org/10.1093/rfs/hhm014>
- White, H. (2000). A reality check for data snooping. *Econometrica*, 68(5), 1097–1126. <https://doi.org/10.1111/1468-0262.00152>
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>