

Forecasting cryptocurrencies in turbulent times: Evidence on parsimony versus model complexity

 Anna Tatarczak¹

 Oleksandra Humeniuk²

Abstract

This study examines short-term return forecasting for Bitcoin, Ethereum, and Litecoin over 2020–2024, comparing autoregressive benchmarks with Kitchen Sink and VARX-type models using point and density accuracy measures supported by Diebold–Mariano and Model Confidence Set inference. The results demonstrate that the AR(1) benchmark and parsimonious specifications incorporating cryptocurrency-specific variables consistently outperform the more elaborate linear frameworks considered, while the inclusion of macro-financial predictors offers limited benefits. Findings highlight the robustness of autoregressive dynamics for short-term cryptocurrency forecasting and underscore the importance of parsimony over model complexity. These results are consistent with a market environment characterised by high structural uncertainty, sentiment-driven trading and rapidly shifting regimes, in which additional macro-financial information contributes little to forecastability beyond short-run return momentum and crypto-specific volatility.

Keywords

- cryptocurrencies
- financial forecasting
- time series models
- emerging markets
- financial markets

JEL codes: C32, C53, G15, G17, E44.

Article received 27 September 2025, accepted 17 February 2026.

Suggested citation: Tatarczak, A., & Humeniuk, O. (2026). Forecasting cryptocurrencies in turbulent times: Evidence on parsimony versus model complexity. *Economics and Business Review*, 12(1), 0-00. <https://doi.org/10.18559/ebr.2026.1.2652>



This work is licensed under a Creative Commons Attribution 4.0 International License
<https://creativecommons.org/licenses/by/4.0>

¹ Maria Curie-Skłodowska University, 5 M. Curie-Skłodowskiej Square, 20-031 Lublin, Poland, corresponding author: anna.tatarczak@umcs.pl, <https://orcid.org/0000-0001-8573-5791>.

² Maria Curie-Skłodowska University, 5 M. Curie-Skłodowskiej Square, 20-031 Lublin, Poland, oleksandrahumeniuk5@gmail.com, <https://orcid.org/0009-0001-8476-1987>.

Introduction

The increasing popularity of cryptocurrencies has fundamentally reshaped financial markets by introducing decentralised, transparent, yet highly volatile digital assets (Arnone, 2024; Lipton, 2021; Wątorek et al., 2023). These characteristics challenge traditional frameworks of value, exchange, and financial forecasting, making cryptocurrencies a focal point for both academic inquiry and investment practice. In this context, time series forecasting has become particularly important, as it provides tools for managing risk, designing investment strategies, and understanding broader market dynamics. Cryptocurrency markets share many statistical properties with traditional financial assets—including non-normal return distributions, volatility clustering, and non-linear dependence structures—while typically exhibiting these features with greater intensity and instability, particularly in terms of volatility levels and regime persistence (Bouchaud, 2020; Sewell, 2011). This instability is consistent with a market environment in which sentiment and disagreement dynamics in online discussions may contribute to future volatility, although the effect appears weaker for Bitcoin than for equities (Akarsu & Yilmaz, 2024). In addition, continuous 24/7 trading, evolving regulation, and rapid technological shifts further complicate modelling and forecasting efforts (Corbet et al., 2018; Zetzsche et al., 2017).

Since Bitcoin's introduction (Nakamoto, 2008), the cryptocurrency ecosystem has expanded to include assets such as Ethereum and Litecoin, with features like smart contracts and faster transactions (Buterin, 2013), thereby increasing the demand for accurate short-term forecasting. Existing studies employ both univariate and multivariate approaches (Antar, 2025; Garay et al., 2024), with mixed evidence on the benefits of incorporating macro-financial predictors (Campbell et al., 1997; Catania et al., 2018). Recent findings suggest that such variables often fail to improve short-horizon forecasts (Agarwal et al., 2024; Casarin et al., 2015; Conlon et al., 2021; Wronka, 2022), underscoring an ongoing trade-off between model complexity and forecasting reliability.

Recent syntheses document rapid growth and thematic diversification in cryptocurrency research; however, they also highlight fragmentation and the lack of integrative short-horizon forecasting frameworks (Jalal et al., 2021; Yue et al., 2021). While sustainability-focused reviews extend the agenda toward ESG issues, they do not assess forecasting gains at daily horizons (Alqudah et al., 2023). At the same time, methodological surveys note the widespread adoption of deep learning models without systematic benchmarking against parsimonious time-series baselines using formal inference (J. Zhang et al., 2024), leaving the incremental value of model complexity and external predictors unresolved. Motivated by these insights, this study examines short-term forecastability as a reflection of underlying market mechanisms. Existing

empirical evidence shows that cryptocurrency prices are heavily influenced by sentiment, speculative trading, and regime-dependent market conditions rather than by slow-moving macroeconomic fundamentals (Jalal et al., 2025; Yue et al., 2021). At the same time, increasingly complex forecasting architectures have been adopted without clear evidence that they yield economically meaningful gains (C. Zhang et al., 2024). Against this backdrop, comparing parsimonious autoregressive benchmarks with multivariate and predictor-rich models can provide insight into whether short-term predictability is structurally constrained.

This study investigates the performance of several forecasting models—both univariate and multivariate—for predicting short-term returns of Bitcoin (BTC), Ethereum (ETH), and Litecoin (LTC) over the 2020–2024 period. The 2020–2024 sample deliberately covers an exceptionally turbulent phase associated with the COVID-19 pandemic, geopolitical tensions and the global inflation surge, which provides a natural stress test for forecasting models under extreme uncertainty. Rather than assuming that cryptocurrencies behave differently from other financial assets in such conditions, the paper verifies whether added model complexity offers predictive gains when even institutional macroeconomic forecasts perform poorly. Specifically, it compares benchmark autoregressive models with more complex approaches, including Kitchen Sink (KS) regression and Vector Autoregressive models with exogenous variables (VARX). The empirical study investigates whether model complexity improves short-horizon return and density forecasts relative to simple autoregressive benchmarks, and whether external predictors generate statistically significant gains in predictive accuracy. Model evaluation uses both point metrics—mean squared error (MSE), mean absolute error (MAE), and mean absolute deviation (MAD)—and probabilistic metrics, such as the Log Score (LS). Statistical significance of forecast improvements is assessed using the Diebold–Mariano test (Diebold & Mariano, 1995) and Model Confidence Set (Hansen et al., 2011) procedures. The contribution of this paper lies in its comparative analysis of forecasting models applied to cryptocurrency returns, using daily data spanning the period 2020–2024. The study assesses predictive accuracy under highly volatile and structurally unstable market conditions, comparing univariate and multivariate models from autoregressive benchmarks to Kitchen Sink and VARX specifications. By contrasting crypto-specific dynamics with models incorporating macro-financial and volatility predictors, it evaluates the incremental value of exogenous information for short-term forecasting. Model performance is assessed using point and density measures (MSE, MAE, MAD, Log Score) with formal inference via Diebold–Mariano and Model Confidence Set procedures, providing a robust multidimensional evaluation of forecasting accuracy.

The rest of the paper is organised as follows: Section 1 is devoted to the literature review on cryptocurrency time series. Section 2 provides information

about data and methodology: data sources and processing, variable construction, model specifications, and the forecasting and evaluation design. In Section 3, we present our empirical findings. Section 4 is devoted to the discussion of results and their implications. Section 5 discusses study limitations and avenues for future research. The last section comprises a critical summary and conclusions.

1. Literature review

Financial time series consist of sequential observations of asset prices and related indicators and are central to forecasting and risk management (Campbell et al., 1997; Fan & Yao, 2003; Taylor, 2008; Tsay, 2005). Their modelling is complicated by stylised facts such as non-normal, heavy-tailed return distributions, excess kurtosis and skewness (Cont, 2001; Mandelbrot, 1963; Rachev et al., 2005; Sewell, 2011), weak linear autocorrelation, and persistent volatility clustering in squared or absolute returns (Bollerslev, 1986). Additional stylised facts include the leverage effect (Black, 1976; Christie, 1982), long memory in volatility (Baillie et al., 1996; Mills & Markellos, 2008), and non-linear dependence structures (Granger & Teräsvirta, 1993; Hamilton, 1994), implying intrinsic uncertainty in financial time series and motivating the use of advanced econometric and machine learning approaches (Sezer et al., 2020; C. Zhang et al., 2024). In cryptocurrency markets, this uncertainty is further reinforced by evidence of time-varying efficiency and episodic return predictability, consistent with the adaptive market hypothesis (Karasiński, 2023). Traditional models such as ARIMA and GARCH remain central for modelling short-term dynamics and volatility (Sezer et al., 2020; C. Zhang et al., 2024), but their linear structure and restrictive assumptions limit their ability to capture non-linearities, heavy tails, and structural breaks (Sezer et al., 2020). These limitations are especially pronounced in cryptocurrency markets, which are decentralised (Nakamoto, 2008), highly volatile and sentiment-driven (Bouoiyour et al., 2015), fragmented across exchanges (Feng et al., 2018; Gandal & Halaburda, 2014), continuously traded (Corbet et al., 2018), and subject to fragmented regulatory oversight (Zetzsche et al., 2017). Given these complexities, traditional econometric models often underperform in cryptocurrency settings, motivating the use of machine learning methods such as ANNs and SVR to capture non-linear and high-dimensional dependencies (Jiang, 2021; L. Zhang et al., 2017). More recently, deep learning models have gained prominence due to their ability to jointly model linear and non-linear structures directly from data (Bengio, 2012; Bouteska et al., 2024; Zhang et al., 2024). Architectures including LSTM, convolutional–recurrent models, and Transformers have demonstrated strong predictive performance in financial and cryptocurrency time series (Hu et al.,

2021; C. Zhang et al., 2023), with LSTM remaining a standard benchmark and newer architectures offering greater flexibility at higher computational cost (Sezer et al., 2020; Zhang et al., 2023).

Recent bibliometric studies support this picture of a fragmented and methodologically heterogeneous literature on cryptocurrencies. Yue et al. (2021) show that research on the economic effects of cryptocurrencies has evolved from early work on technological foundations and miner behaviour toward analyses of price formation, risk management, and the macroeconomic implications of digital assets, but argue that the underlying transmission mechanisms and theoretical frameworks remain underdeveloped. Complementary evidence from Jalal et al. (2025) indicates that business and finance research is organised around four main streams—determinants of returns, market efficiency, (de)diversification and herding, and regulation and governance—yet many contributions rely on overlapping datasets and parallel empirical designs, which limits cumulative progress and integrative assessments of forecasting performance. More recently, Alqudah et al. (2023) document a rapid expansion of work at the intersection of cryptocurrencies and ESG, highlighting concerns about environmental externalities, speculative trading, and the long-term sustainability of digital assets as an investment class. At the same time, reviews of deep learning applications in financial forecasting emphasise that increasingly complex architectures—LSTMs, convolutional–recurrent hybrids, Transformers, and related models—have become the default choice for price prediction, even though their incremental benefits over simpler time-series benchmarks are not always systematically evaluated, especially with respect to economic interpretability and model risk (C. Zhang et al., 2024). These syntheses collectively suggest that a key unresolved issue is not only how to forecast cryptocurrency prices, but whether forecast gains delivered by complex models are economically meaningful and theoretically consistent with the underlying market structure. In this context, linear time-series models retain an important role as benchmarks with well-understood statistical properties, against which the incremental value of more complex non-linear and deep learning architectures can be assessed.

2. Methodology and data

2.1. Dataset and model specification

The empirical analysis is based on daily data covering the period from 1 January 2020 to 1 December 2024. This period is characterised by pronounced regime shifts and market-wide stress episodes, so the resulting forecasts should

be interpreted as conditional on a high-volatility, crisis-like environment rather than as representative of more tranquil phases of cryptocurrency trading. The empirical strategy deliberately focuses on linear specifications (AR, VAR, VARX, KS), providing a consistent framework for formal forecast evaluation and serving as a reference point for future comparisons with non-linear and deep-learning models. Non-linear econometric specifications (such as TAR, STAR or Markov-switching models), volatility models (such as GARCH or stochastic volatility frameworks), and machine-learning or deep learning architectures (such as LSTM networks, transformers or ensemble models) are therefore intentionally excluded from the empirical design in order to keep the benchmark coherent and econometrically tractable over the turbulent 2020–2024 period. The analysis focuses on Bitcoin (BTC), Ethereum (ETH), and Litecoin (LTC) because they represent the longest-standing and most liquid segments of the cryptocurrency market, ensuring reliable high-frequency data and reducing concerns regarding idiosyncratic exchange-specific noise. BTC and ETH jointly account for the majority of market capitalisation and trading volume, making them the dominant drivers of systemic crypto-asset dynamics. LTC, while smaller, serves as a mid-cap asset with strong historical continuity, facilitating an assessment of whether forecasting performance generalises beyond the two flagship cryptocurrencies. Together, these three assets provide a balanced and representative sample of the cryptocurrency market. The dataset comprises time series for three leading cryptocurrencies—BTC, ETH, and LTC—alongside a range of macro-financial indicators that are used as potential predictors in the multivariate model configurations. All financial and macroeconomic data were retrieved from Investing.com, while additional background information on cryptocurrencies was sourced from CoinMarketCap.

While the analysis relies on spot market data from cryptocurrency exchanges, it is important to note that the investment landscape evolved substantially during 2021–2024 with the introduction and rapid expansion of exchange-traded products (ETFs/ETNs) tracking BTC and ETH on regulated exchanges such as XETR, XLON, XNAS, XNYS, and XWAR. These instruments attract a broader and more heterogeneous investor base and may follow distinct patterns in price discovery and volatility transmission. As they fall outside the scope of this study, the empirical results should be interpreted as reflecting the dynamics of spot markets rather than the behaviour of regulated exchange-traded cryptocurrency products.

Cryptocurrency-specific variables include daily closing prices, high–low price spreads, and trading volumes. These were transformed to logarithmic returns to stabilise variance and approximate stationarity. Volatility proxies were constructed as the natural logarithm of the daily high–low spread. The macro-financial variables incorporated into the analysis include the US 5-year credit default swap index (CDS_5y), the STOXX Europe 600 Index (ES_600), the Dow Jones US Gold Mining Index (GLD), the Nikkei 225 Index (NK225), the S&P 500

Index (SP500), the Dow Jones Commodity Index – Silver Subcomponent (SV), and the CBOE Volatility Index (VIX). These indicators were selected to reflect global equity market conditions, credit risk perceptions, commodity market dynamics, and investor sentiment. To address non-stationarity, most of the financial variables were transformed using the first difference of their logarithmic values. The selection of macro-financial predictors was guided by established categories that capture global risk sentiment (VIX), equity market conditions across major regions (S&P 500, STOXX Europe 600, Nikkei 225), commodity-related hedging channels (Gold Mining Index, Silver Subindex), and credit risk perceptions (US 5-year CDS). These variables are standard benchmarks in the short-horizon forecasting literature and serve to represent broad asset-class linkages without overwhelming the model with highly collinear predictors. Indicators such as NASDAQ, the US Dollar Index, interest rates or inflation measures were not included because they convey information that is strongly correlated with the equity and credit indices already present, increasing the risk of overfitting in daily-frequency models. Similarly, on-chain data (e.g., hash rate, transaction fees, miner revenue) were excluded to maintain a consistent daily sampling frequency and to ensure comparability across assets, but their incorporation presents a valuable extension for future research.

The forecasting framework comprises univariate and multivariate models. Univariate specifications include an AR(1) benchmark, a full Kitchen Sink (KS) regression with macro-financial and crypto-specific predictors, a reduced KS-noregr model excluding macro-financial variables, and an Avg forecast combining KS and KS-noregr. The multivariate set consists of a benchmark AR(1) estimated separately (M1), a VAR(3) capturing cross-asset dynamics (M2), a VAR with endogenous volatility proxies based on high–low spreads (M3), and a VARX(3) incorporating macro-financial variables as exogenous inputs (M4). All models are estimated using an expanding-window scheme to approximate real-time forecasting and limit overfitting.

2.2. Forecasting strategy and evaluation metrics

The analysis evaluates two types of forecasts: point forecasts, which estimate the conditional expectation of returns, and density forecasts, which provide the full predictive distribution. Forecasts are generated for one- to seven-day horizons ($h = 1$ to $h = 7$), and the performance of each model is assessed using both point and density forecast accuracy measures.

Point forecasts are evaluated using the following metrics: mean squared error (MSE) measures the average squared difference between forecasted and actual returns. It penalises larger forecast errors more heavily due to squaring, which makes it particularly sensitive to outliers. Mean absolute de-

viation (MAD) calculates the average absolute difference between forecasted and actual values. MAD also serves as an alternative to MSE due to its lower sensitivity to outliers. This metric provides a more robust evaluation by treating all forecast errors equally. The mean absolute error (MAE) calculates the average absolute difference between the predicted and the actual values, treating all errors equally regardless of their direction. A lower MAE value indicates a more accurate forecasting model. In contrast to metrics such as the mean squared error (MSE), MAE is less sensitive to large individual errors and is thus more robust to the presence of outliers.

Density forecasts are assessed using the Log Score (LS). It measures how well the forecasted distribution assigns probability to the observed outcome. Given the heavy-tailed nature of cryptocurrency returns and the presence of extreme downside events, the evaluation framework is extended to include tail-sensitive risk measures. In addition to symmetric loss functions and the Log Score, we compute one-sided left-tail Value at Risk (VaR) and Expected Shortfall (ES) at the 5% and 1% significance levels. Predictive distributions are approximated using the same parametric assumptions adopted for density forecasts, ensuring internal consistency across evaluation criteria. VaR forecasts are assessed using unconditional and conditional coverage tests, while ES accuracy is evaluated on exceedance days. This extension allows the forecast comparison to explicitly account for downside risk characteristics that are central to cryptocurrency markets. To compare forecasting models statistically, the study employs two formal evaluation procedures. The Diebold–Mariano (DM) test (Diebold & Mariano, 1995) examines the null hypothesis of equal forecast accuracy between two competing models. This test is applied separately for MSE (point forecasts) and LS (density forecasts). The Model Confidence Set (MCS) procedure (Hansen et al., 2011) identifies the subset of models that are statistically indistinguishable in performance at a given confidence level. The procedure iteratively removes the least accurate model until the null hypothesis of equal predictive accuracy can no longer be rejected. Together, these metrics and statistical tests provide a rigorous framework for evaluating model performance across both forecast horizons and evaluation dimensions.

3. Empirical results

3.1. Univariate forecast evaluation

This subsection assesses out-of-sample point forecast accuracy of four univariate models—AR(1), KS, KS-noregr, and Avg—applied to BTC, ETH, and LTC using an expanding window over horizons from $h = 1$ to $h = 7$.

Performance is evaluated using MSE, with AR(1) as the benchmark, KS including all predictors, KS-noregr restricting attention to autoregressive and crypto-specific variables, and Avg combining KS and KS-noregr forecasts, allowing the incremental value of macro-financial information and model complexity to be assessed.

As shown in Table 1, the AR(1) benchmark performs robustly for Bitcoin and Ethereum across all horizons, while the full KS specification systematically underperforms, consistent with overfitting from noisy macro-financial predictors. In contrast, the parsimonious KS-noregr and Avg models dominate KS and, particularly for Ethereum and Litecoin, often match or outperform AR(1). The Model Confidence Set applied to MSE excludes KS for all assets and retains AR(1), KS-noregr, and Avg, while Diebold–Mariano tests yield few significant differences but remain directionally consistent with these results.

Table 1. Mean squared error for univariate models (in %)

<i>h</i>	1	2	3	4	5	6	7
Bitcoin							
AR1	6.11	6.02	6.01	6.01	6.03	6.03	6.04
KS	6.20	6.13	6.13	6.16	6.16	6.16	6.15
KS-noregr	6.12	6.07	6.09	6.08	6.10	6.07	6.12
Avg	6.12	6.07	6.09	6.08	6.10	6.07	6.12
Ethereum							
AR1	7.78	7.62	7.60	7.62	7.64	7.66	7.68
KS	7.99	7.81	7.81	7.86	7.87	7.92	7.93
KS-noregr	7.77	7.62	7.58	7.64	7.63	7.68	7.67
Avg	7.77	7.62	7.58	7.64	7.63	7.68	7.67
Litecoin							
AR1	10.91	10.53	10.48	10.51	10.53	10.55	10.57
KS	11.12	10.96	10.96	10.99	10.99	11.05	11.03
KS-noregr	10.79	10.65	10.50	10.55	10.47	10.54	10.47
Avg	10.78	10.65	10.50	10.55	10.47	10.54	10.47

Source: own work.

In addition to MSE, Table 2 presents the mean absolute deviation (MAD), which provides a complementary view of forecast accuracy. The results broadly confirm the MSE findings: KS-noregr and Avg yield lower errors than KS, with particularly strong performance for Litecoin.

Table 2. Mean absolute deviation over the forecast horizon (%)

<i>h</i>	1	2	3	4	5	6	7
Bitcoin							
AR1	1.72	1.71	1.69	1.69	1.71	1.71	1.71
KS	1.75	1.74	1.73	1.74	1.74	1.74	1.74
KS-noregr	1.71	1.70	1.70	1.70	1.70	1.70	1.70
Avg	1.71	1.69	1.70	1.70	1.69	1.69	1.69
Ethereum							
AR1	1.89	1.86	1.86	1.86	1.87	1.87	1.87
KS	1.93	1.91	1.91	1.92	1.92	1.93	1.93
KS-noregr	1.89	1.87	1.87	1.87	1.87	1.88	1.87
Avg	1.89	1.87	1.87	1.87	1.87	1.88	1.87
Litecoin							
AR1	2.22	2.18	2.17	2.17	2.17	2.18	2.18
KS	2.24	2.22	2.21	2.22	2.22	2.23	2.22
KS-noregr	2.22	2.20	2.22	2.20	2.20	2.22	2.19
Avg	2.22	2.20	2.22	2.20	2.20	2.22	2.19

Source: own work.

Diebold–Mariano tests indicate that differences in point forecast accuracy are generally not statistically significant; the only significant improvement over AR(1) occurs for Litecoin at horizon $h = 2$ using the full KS model ($p = 0.0439$). Predictor relevance analysis shows consistent patterns across assets: ES600 is the most influential macro-financial predictor for Bitcoin and Ethereum, alongside strong autoregressive spillovers—particularly from Bitcoin to Ethereum—while Litecoin is primarily driven by lagged returns of Bitcoin and Ethereum, with a secondary role for the S&P 500. Rankings based on coefficient magnitude corroborate these findings, reinforcing the hierarchical structure of influence within the cryptocurrency market.

Turning to density forecasts, Table 3 reports the Log Score (LS) values for each univariate model. The KS model consistently achieves the lowest LS values for all three cryptocurrencies, suggesting superior probabilistic accuracy. These improvements, however, are not always confirmed as statistically significant by the DM test.

The Model Confidence Set (MCS) procedure applied to Log Score losses yields a differentiated picture across assets. For Bitcoin, only the AR(1) benchmark remains in the superior set, confirming its robustness in probabilistic forecasting. In contrast, for Ethereum and Litecoin, the KS-noregr and Avg

Table 3. Log Score (LS) of univariate models over the forecast horizon

<i>h</i>	1	2	3	4	5	6	7
Bitcoin							
AR1	837.31	837.88	835.95	833.34	830.52	828.25	825.53
KS	834.24	834.30	832.00	828.81	826.75	824.29	822.49
KS-noregr	836.92	836.14	833.28	831.36	828.43	826.93	823.13
Avg	836.92	836.14	833.28	831.36	828.43	826.93	823.13
Ethereum							
AR1	792.84	794.53	792.92	790.13	787.48	784.95	782.20
KS	787.88	789.95	787.90	784.46	782.25	778.91	776.45
KS-noregr	793.12	794.53	793.40	789.81	787.85	784.48	782.51
Avg	793.12	794.53	793.40	789.81	787.85	784.48	782.51
Litecoin							
AR1	730.83	735.28	734.14	731.79	729.34	726.92	724.68
KS	726.24	728.08	726.06	723.49	721.39	718.53	716.91
KS-noregr	732.94	733.33	733.83	730.95	730.28	727.10	726.38
Avg	732.94	733.33	733.82	730.95	730.28	727.10	726.38

Source: own work.

specifications are retained, while the full KS model is consistently excluded due to its statistical instability. Taken together, these findings indicate that univariate models centred on crypto-specific information—lagged returns and high–low volatility proxies—offer more reliable density forecasts than specifications that indiscriminately incorporate macro-financial variables, whose short-term predictive contribution appears limited.

To account for the heavy-tailed nature of cryptocurrency returns and the presence of extreme downside events, the forecast evaluation is extended to include tail-sensitive risk measures. Table 4 reports one-day-ahead Value at Risk (VaR) and Expected Shortfall (ES) backtesting results at the 5% and 1% significance levels for all univariate specifications. In line with the point and density forecast results, the AR(1) benchmark and the parsimonious KS-noregr model exhibit the most stable tail-risk performance across assets. Their empirical exceedance rates are close to nominal levels, and both unconditional and conditional coverage tests generally fail to reject correct calibration. By contrast, the full Kitchen Sink (KS) specification displays weaker tail behaviour, particularly for Ethereum, where violations cluster over short horizons, leading to rejections of conditional coverage despite acceptable average accuracy. This indicates that the inclusion of broad macro-financial predictors may de-

teriorate tail-risk properties even when symmetric error metrics or Log Scores suggest comparable performance. Overall, tail-focused diagnostics reinforce the central finding of this study: parsimonious autoregressive structures and cryptocurrency-specific information deliver more robust short-term forecasts, not only in terms of average accuracy but also with respect to downside risk.

Table 4. VaR and ES backtesting results for one-day-ahead forecasts

Asset	Model	VaR 5% hit ()	CC p -value	VaR 1% hit (%)	CC p -value	ES tail loss
BTC	AR(1)	4.6	0.44	1.9	0.26	low
BTC	KS	5.7	0.26	1.9	0.26	medium
BTC	KS-noregr	4.6	0.44	1.9	0.26	low
ETH	AR(1)	4.4	0.92	0.8	0.92	low
ETH	KS	6.0	0.02	0.5	0.63	high
ETH	KS-noregr	4.4	0.92	0.8	0.92	low
LTC	AR(1)	4.1	0.75	1.4	0.75	low
LTC	KS	4.4	0.43	1.4	0.75	medium
LTC	KS-noregr	4.6	0.44	1.4	0.75	low

Source: own work.

3.2. Multivariate forecast evaluation

As shown in Tables 5 and 6, the benchmark AR(1) specification (M1) attains the lowest MSE and MAE at most horizons, while the VAR(3) model (M2) performs similarly without systematic gains. Adding endogenous volatility proxies (M3) yields only marginal and statistically insignificant improvements at longer horizons, whereas the VARX(3) model with macro-financial predictors (M4) consistently underperforms, particularly beyond $h = 4$, indicating limited short-term forecasting gains from cross-asset spillovers or macroeconomic information.

Table 5. Mean squared error for multivariate models (in %)

h	1	2	3	4	5	6	7
M1	5.58	6.45	7.25	8.16	9.44	11.00	13.71
M2	5.84	6.57	7.34	8.13	9.43	11.01	13.74
M3	6.11	8.84	7.90	8.35	9.55	11.02	13.72
M4	6.93	8.13	9.38	10.61	12.37	14.34	17.86

Source: own work.

The Diebold–Mariano test for MSE does not yield statistically significant results. For the model comparisons, the null hypothesis of equal predictive accuracy between AR(1) and alternative specifications could not be rejected at the 5% significance level. The Model Confidence Set procedure indicates that M1 belongs to the Superior Set of Models at a 10% confidence level, using MSE as the loss metric. M2 and M3 occasionally join the set, particularly at short-term horizons. M4 is regularly eliminated due to its limited predictive contribution in this setup.

Table 6. Mean absolute error for multivariate models (in %)

<i>h</i>	1	2	3	4	5	6	7
M1	1.43	1.59	1.74	1.87	2.06	2.27	2.73
M2	1.51	1.63	1.79	1.85	2.06	2.25	2.73
M3	1.63	1.98	1.99	1.96	2.13	2.27	2.73
M4	1.75	1.94	2.24	2.43	2.78	3.04	3.65

Source: own work.

Diebold–Mariano tests based on MAE confirm that M1 consistently outperforms M4 across horizons, while horizon $h = 7$ yields no significant results due to limited observations. Simpler models (M1 and M2) provide more reliable short-term forecasts, and MCS results for MAE mirror those for MSE, with M1 and M2 retained in the superior set at the 10% level. Density forecast results (Table 7) further show strong AR(1) performance in terms of Log Score; although M3 and M4 occasionally improve scores at longer horizons, these gains are modest and statistically insignificant.

Table 7. Log Score (LS) of multivariate models over the forecast horizon

<i>h</i>	1	2	3	4	5	6	7
M1	2.30	2.56	2.21	2.17	2.10	2.02	1.87
M2	1.83	1.82	1.81	1.80	1.78	1.76	1.72
M3	1.86	1.75	1.78	1.77	1.75	1.74	1.71
M4	1.80	1.81	1.80	1.77	1.76	1.75	1.72

Source: own work.

Table 8 provides a breakdown of Log Scores for each cryptocurrency. Notably, Bitcoin consistently achieves higher predictive accuracy compared to Ethereum and Litecoin across all model specifications. The decline in Log Scores across forecast horizons reflects the increased difficulty of making longer-term predictions in volatile markets.

Table 8. Log Score (LS) for multivariate models over the forecast horizon for each cryptocurrency

<i>h</i>	1	2	3	4	5	6	7
Bitcoin							
M1	2.53	2.50	2.47	2.44	2.39	2.33	2.22
M2	1.98	1.98	1.96	1.95	1.94	1.92	1.89
M3	2.05	1.93	1.92	1.91	1.91	1.91	1.88
M4	1.95	1.96	1.95	1.92	1.93	1.92	1.89
Ethereum							
M1	2.26	2.21	2.17	2.12	2.05	1.94	1.79
M2	1.79	1.79	1.78	1.77	1.75	1.72	1.68
M3	1.86	1.69	1.75	1.75	1.73	1.71	1.68
M4	1.77	1.77	1.77	1.73	1.73	1.71	1.69
Litecoin							
M1	2.17	2.05	2.00	1.94	1.86	1.76	1.61
M2	1.72	1.70	1.68	1.67	1.65	1.62	1.58
M3	1.66	1.62	1.64	1.64	1.62	1.61	1.57
M4	1.69	1.69	1.68	1.65	1.64	1.63	1.58

Source: own work.

The Diebold–Mariano test results at the 5% significance level point out that the M1 consistently exhibits statistically superior predictive performance compared to the other models across short to medium forecast horizons (h_1 – h_5). At longer forecast horizons (h_6 – h_7), however, the model's predictive accuracy differences become statistically insignificant, suggesting an increase in forecast uncertainty. Detailed Diebold–Mariano test p -values for pairwise model comparisons across horizons and loss functions are reported in Appendix A.

The Model Confidence Set analysis also confirmed that only AR(1) consistently provided superior predictive accuracy across Bitcoin, Ethereum, and Litecoin returns. More complex models, such as VAR and VARX, were eliminated from the superior set at a 10% significance level, indicating that additional variables or model complexity did not yield better forecast performance in this application.

Figure 1 illustrates the cumulative Log Score differences for each model relative to the AR(1) benchmark. Positive values indicate better performance. M2 and M3 show minor gains at longer horizons, but M1 remains dominant across most horizons.

The analysis of cumulative adjusted Log Scores indicates that the baseline model M1 consistently outperforms its competitors across all forecast

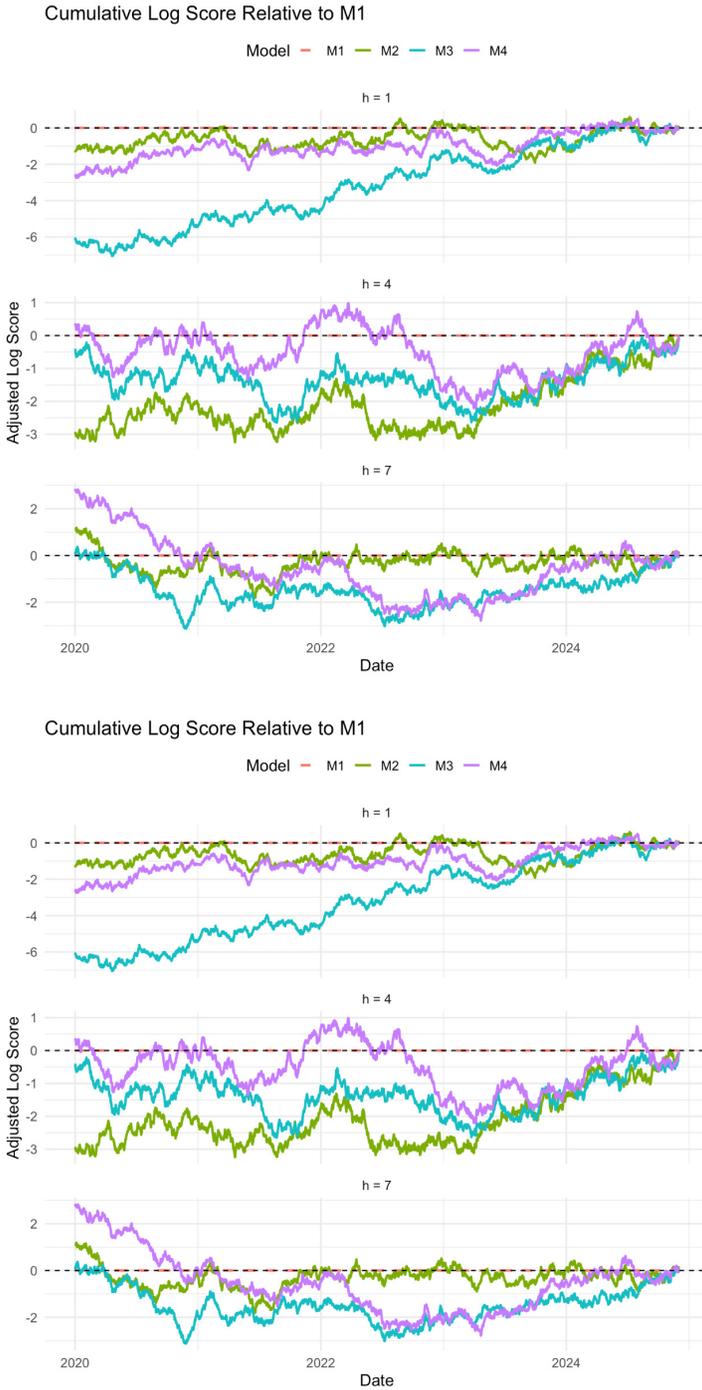


Figure 1. Cumulative Log Score visualisation

Source: own work.

horizons, showing the most significant lead at shorter horizons. In contrast, models M2 and M3 exhibit lower predictive accuracy, especially for short-term forecasts. Model M4 shows some competitiveness at medium horizons but does not consistently surpass M1. As the forecast horizons extend, the errors increase, and the performance gaps between models narrow, though M1 still retains a slight edge.

4. Discussion

The dominance of the AR(1) benchmark and closely related crypto-specific linear specifications aligns with economic evidence suggesting that short-horizon cryptocurrency returns are driven primarily by sentiment, momentum, and market microstructure rather than by macro-financial fundamentals. Bibliometric studies indicate that speculative behaviour, attention cycles and regime shifts play a central role in shaping price dynamics (Jalal et al., 2025; Yue et al., 2021). In such environments, additional predictors often introduce noise rather than signal, limiting the gains from complex multivariate structures. The weak and unstable contribution of macro variables is consistent with findings that cryptocurrencies exhibit low and highly time-varying integration with traditional financial assets, particularly during stress periods (Alqudah et al., 2023), including pandemic-related episodes, in which the safe-haven and diversification properties of cryptocurrencies appear regime-dependent and unstable (Barbu et al., 2022). Deep learning reviews also note that model sophistication cannot fully overcome structural instability and noisy patterns at daily frequencies (J. Zhang et al., 2024). Together, these mechanisms provide an economic interpretation for the strong performance of parsimonious linear specifications observed in this study. The limited contribution of macro-financial variables in our setting should not be interpreted as evidence of their irrelevance, but rather as a reflection of the low signal-to-noise ratio at daily horizons and the strong collinearity among global risk and equity indicators. More granular predictors—such as on-chain activity measures—may capture additional structure, but incorporating them requires a dedicated modelling framework beyond the scope of the present benchmark study.

From the perspective of the current forecasting frontier, non-linear and deep learning approaches are natural candidates to capture residual non-linearities and complex interactions that lie beyond linear dynamics. In this study, however, we treat linear models as economically interpretable benchmarks under stress conditions, so that subsequent work can evaluate whether additional complexity in non-linear or deep learning frameworks delivers stable and economically meaningful gains over these baselines. Accordingly, the

contribution of this paper lies in establishing a rigorous post-2020 benchmark for linear time-series models rather than proposing a new state-of-the-art non-linear forecasting architecture.

The results should not be interpreted as evidence that cryptocurrency markets were more forecastable than other assets during this turbulent period. The small and often insignificant performance differentials between models mirror the widespread forecasting failures observed in institutional macro-economic and financial projections, and indicate that, under such conditions, added econometric complexity yields no meaningful gains over parsimonious autoregressive benchmarks.

The empirical results indicate that parsimonious models based on autoregressive dynamics and crypto-specific indicators outperform more complex specifications at short horizons. The AR(1) model exhibits stable performance across all assets, while the full Kitchen Sink (KS) specification suffers from overfitting, particularly for Bitcoin and Ethereum. By contrast, the reduced KS-noregr model performs comparably to or better than AR(1) in both point and density forecasts. These findings are consistent with Campbell et al. (1997) and Catania et al. (2018), who document the limited short-term forecasting gains from macro-financial predictors.

The multivariate results confirm the dominance of the AR(1) benchmark, which consistently matches or outperforms VAR and VARX models across horizons. Gains from adding volatility proxies or macro-financial variables are generally marginal and statistically insignificant; while the VAR with high–low spreads (M3) shows occasional improvements at medium horizons, the VARX specification (M4) underperforms. These findings are consistent with Casarin et al. (2015) and Corbet et al. (2018), who document the horizon-dependent and often weak short-run predictive contributions of external financial information.

Probabilistic evaluation using the Log Score shows that while the KS model occasionally improves density forecasts—particularly for Ethereum—these gains are rarely statistically significant and are not confirmed by the Model Confidence Set, whereas AR(1) is consistently retained in the superior set. This robustness of simple autoregressive models is consistent with the findings of Adhikari and Agrawal (2014), especially in noisy financial environments.

These results are consistent with prior evidence emphasising the importance of parsimony and domain-specific model design in financial forecasting (Mills & Markellos, 2008; Sezer et al., 2020). In highly volatile and structurally unstable cryptocurrency markets (Bouoiyour et al., 2015; Sewell, 2011), complex models with noisy external predictors may reduce short-term accuracy, while the relevance of macro-financial variables is likely horizon- and target-dependent. This interpretation aligns with Baillie et al. (1996) and Jiang (2021), who show that non-linear and long-memory effects are more effectively captured over longer samples or for alternative targets such as volatil-

ity. Moreover, the limited gains from macro-financial variables in this study should be interpreted in light of the increasingly time-varying and regime-dependent integration of cryptocurrency markets with traditional financial systems. Recent evidence indicates that, particularly since 2020, cryptocurrencies have become more strongly correlated with major equity indices—especially technology stocks—and US macroeconomic conditions, although this integration remains unstable and horizon-dependent (Wątorrek et al., 2023).

5. Limitations and avenues for further research

This study has several limitations that should be considered when interpreting the results. Firstly, the analysis is restricted to three major cryptocurrencies—Bitcoin, Ethereum, and Litecoin—over the period 2020–2024. The findings may therefore not generalisable to smaller tokens, stablecoins, or other digital assets with different market structures and liquidity conditions. Secondly, the dataset coincides with major regime shifts, including the COVID-19 aftermath, geopolitical tensions and the global inflation surge, which may have introduced structural breaks that the expanding-window approach does not fully capture. Since the study does not include a comparison with the more tranquil pre-2020 period (e.g., 2015–2019), we cannot formally assess whether the dominance of parsimonious models persists across different volatility regimes. Thirdly, the use of aggregated data sources may involve measurement errors. In particular, mismatches between continuous cryptocurrency trading and macro-financial indicators with business-day frequency could lead to timing inconsistencies. The analysis does not incorporate regulated exchange-traded products (ETFs/ETNs) that have emerged as major investment vehicles for BTC and ETH since 2021. Because these instruments may differ from spot markets in terms of volatility transmission, liquidity and investor composition, the generalisability of the results to the ETF/ETN segment is inherently limited. Fourthly, the study evaluates only linear and relatively simple models (AR(1), VAR, VARX, KS). Within this set, the autoregressive benchmark is restricted to an AR(1) specification; more flexible but still parsimonious time-series models such as AR(p), ARIMA, HAR-RV or GARCH-type frameworks are not explored. This restriction is deliberate: the objective is to establish a well-understood linear benchmark for the turbulent 2020–2024 period rather than to propose a new state-of-the-art forecasting architecture.

Consequently, the conclusion that the AR(1) benchmark outperforms more complex specifications should be interpreted as conditional on the restricted class of models considered here, and future research should assess whether alternative low-dimensional autoregressive or volatility models can

match or surpass this baseline. Fifthly, forecasts are based on daily data and do not account for intraday patterns, transaction costs, or exchange heterogeneity, which limits their practical applicability to trading strategies. Sixthly, the forecasting horizons examined in this study are deliberately restricted to short-term windows (1–7 days). While this design is appropriate for evaluating high-frequency return predictability, it may limit the ability of multivariate and VARX-type models to demonstrate their strengths. Macro-financial variables typically operate through slower transmission channels and may exert influence over medium- or long-term horizons rather than at daily frequencies. As a result, the finding that macroeconomic predictors provide limited incremental value should be interpreted as conditional on these short horizons and not more broadly as evidence of their irrelevance. Extending the analysis to multi-week or multi-month horizons would allow future research to assess whether the predictive contribution of macro-financial information increases when the forecast window aligns more closely with the underlying economic adjustment dynamics. Finally, the reliance on standard evaluation metrics may inadequately reflect tail risk and economic significance; accordingly, the findings should be regarded as indicative rather than conclusive.

Conclusions

This study assesses the short-term forecasting performance of univariate and multivariate time series models for daily returns of Bitcoin, Ethereum, and Litecoin over 2020–2024, comparing autoregressive benchmarks with more complex specifications using standard point and density evaluation metrics supported by formal statistical inference. The findings lead to some conclusions. Firstly, the AR(1) specification provides a well-performing baseline for cryptocurrency return forecasting within the linear model class considered. Their consistent inclusion in the superior model sets across assets and horizons underscores their robustness, particularly at short horizons. Secondly, incorporating cryptocurrency-specific variables—especially lagged returns and high–low volatility spreads—yields modest but tangible improvements in forecast accuracy. The reduced Kitchen Sink model (KS-noregr), which focuses on these domain-relevant predictors, frequently outperformed more elaborate specifications, especially for Ethereum and Litecoin. Thirdly, complex multivariate models that integrate macro-financial variables (e.g., VARX) did not systematically enhance forecast accuracy. In several cases, the inclusion of exogenous predictors reduced predictive performance, reflecting issues of overfitting and the limited short-term relevance of macroeconomic signals for cryptocurrency prices. These results are consistent with prior re-

search indicating that digital asset returns are driven more by endogenous momentum and speculative trading than by traditional financial indicators in the short run (Bouoiyour et al., 2015; Conlon et al., 2021; Corbet et al., 2018). Fourthly, density forecasts highlighted some advantages of the full KS model, particularly for Ethereum, but these gains were inconsistent and often lacked statistical significance. By contrast, the AR(1) model demonstrated stable performance across both point and density evaluations, reinforcing the value of parsimonious structures. Overall, within the turbulent post-2020 environment analysed here and within the set of linear models considered, the evidence suggests that parsimony combined with careful variable selection is more effective for short-term cryptocurrency forecasting than additional model complexity. While non-linear and machine learning models may offer gains over longer horizons or alternative targets (e.g., volatility or tail risk), the results confirm the practical relevance of autoregressive dynamics and crypto-specific features. The study provides a unified benchmarking of parsimonious autoregressive baselines against kitchen-sink and VAR/VARX models using formal inference (Diebold–Mariano; Model Confidence Set), clarifying the incremental value of crypto-specific versus macro-financial predictors. Consistent with recent reviews (Jalal et al., 2025; Yue et al., 2021; C. Zhang et al., 2024), the evidence shows that parsimony dominates at daily horizons, establishing transparent and economically interpretable benchmark models and motivating future extensions using high-frequency, on-chain, and non-linear or ensemble frameworks.

From a practical perspective, the findings suggest that parsimonious autoregressive and crypto-specific models account for well-performing tools for short-term forecasting and risk management, while the limited short-run predictive role of macro-financial variables implies that regulatory risk assessments should focus more on market microstructure, speculative behaviour, and crypto-specific volatility dynamics than on broad macroeconomic indicators.

Appendix

Table A1. Mean Squared Error (MSE)

Asset	Horizon	AR(1) vs KS	AR(1) vs KS-noregr	AR(1) vs Avg
BTC	$h = 1$	0.4321	0.6814	0.5178
BTC	$h = 3$	0.0847	0.2941	0.1736
BTC	$h = 7$	0.0632	0.2119	0.0915
ETH	$h = 1$	0.3884	0.7442	0.6017
ETH	$h = 3$	0.0173	0.5628	0.0419
ETH	$h = 7$	0.0191	0.4896	0.0384
LTC	$h = 1$	0.5176	0.8092	0.6931
LTC	$h = 3$	0.1412	0.6034	0.2289
LTC	$h = 7$	0.0924	0.4187	0.1561

Source: own work.

Table A2. Log Score (LS)

Asset	Horizon	AR(1) vs KS	AR(1) vs KS-noregr	AR(1) vs Avg
BTC	$h = 1$	0.4762	0.7028	0.6115
BTC	$h = 3$	0.0819	0.3314	0.0976
BTC	$h = 7$	0.0687	0.2885	0.0894
ETH	$h = 1$	0.4027	0.7619	0.5842
ETH	$h = 3$	0.0236	0.5941	0.0498
ETH	$h = 7$	0.0214	0.5237	0.0441
LTC	$h = 1$	0.5381	0.8124	0.7016
LTC	$h = 3$	0.1597	0.6472	0.2443
LTC	$h = 7$	0.0886	0.4593	0.1698

Source: own work.

References

- Adhikari, R., & Agrawal, R. K. (2014). A combination of artificial neural network and random walk models for financial time series forecasting. *Neural Computing and Applications*, 24(6), 1441–1449. <https://doi.org/10.1007/s00521-013-1386-y>
- Agarwal, M., Gill, K. S., Upadhyay, D., Dangi, S., & Chythanya, K. R. (2024, April). *The evolution of cryptocurrencies: Analysis of Bitcoin, Ethereum, Bit connect and Dogecoin in comparison*. 2024 IEEE 9th International Conference for Convergence in Technology (pp. 1–6). <https://doi.org/10.1109/i2ct61223.2024.10543872>
- Akarsu, S., & Yilmaz, N. (2024). Social media disagreement and financial markets: A comparison of stocks and Bitcoin. *Economics and Business Review*, 10(4), 189–213. <https://doi.org/10.18559/ebr.2024.4.1683>
- Alqudah, M., Ferruz, L., Martín, E., Qudah, H., & Hamdan, F. (2023). The sustainability of investing in cryptocurrencies: A bibliometric analysis of research trends. *International Journal of Financial Studies*, 11(3), 93. <https://doi.org/10.3390/ijfs11030093>
- Antar, M. (2025). Quantile analysis of Bitcoin returns: Uncovering market dynamics. *Journal of Risk Finance*, 26(1), 122–146. <https://doi.org/10.1108/jrf-05-2024-0154>
- Arnone, G. (2024). Introduction to cryptocurrencies and digital currencies. In G. Arnone, *Navigating the world of cryptocurrencies: Technology, economics, regulations, and future trends* (pp. 1–11). Springer Nature. https://doi.org/10.1007/978-3-031-69176-8_1
- Baillie, R. T., Bollerslev, T., & Mikkelsen, H. O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74(1), 3–30.
- Barbu, T. C., Boitan, I. A., & Cepoi, C. O. (2022). Are cryptocurrencies safe havens during the COVID-19 pandemic? A threshold regression perspective with pandemic-related benchmarks. *Economics and Business Review*, 8(2), 29–49. <https://doi.org/10.18559/ebr.2022.2.3>
- Bengio, Y. (2012). Deep learning of representations for unsupervised and transfer learning. *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, 27, 17–36.
- Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section*, 177–181.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327. [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
- Bouchaud, J. P. (2013). Crises and collective socio-economic phenomena: Simple models and challenges. *Journal of Statistical Physics*, 151(3), 567–606. <https://doi.org/10.1007/s10955-012-0687-3>
- Bouoiyour, J., Selmi, R., & Tiwari, A. K. (2015). Is Bitcoin business income or speculative foolery? New ideas through an improved frequency domain analysis. *Annals of Financial Economics*, 10(1), 1550002. <https://doi.org/10.1142/s2010495215500025>
- Bouteska, A., Abedin, M. Z., Hajek, P., & Yuan, K. (2024). Cryptocurrency price forecasting—a comparative analysis of ensemble learning and deep learning methods.

- International Review of Financial Analysis*, 92, 103055. <https://doi.org/10.1016/j.irfa.2023.103055>
- Buterin, V. (2013). *Ethereum whitepaper: A Next-generation smart contract and decentralized application platform*. Ethereum Foundation.
- Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). *The econometrics of financial markets*. Princeton University Press.
- Casarin, R., Grassi, S., Ravazzolo, F., & van Dijk, H. K. (2015). *Dynamic predictive density combinations for large data sets in economics and finance*. Tinbergen Institute Discussion Paper, 15-084/III. <https://hdl.handle.net/10419/125081>
- Catania, L., Grassi, S., & Ravazzolo, F. (2018). *Forecasting cryptocurrencies financial time series*. BI Norwegian Business School.
- Christie, A. A. (1982). The stochastic behavior of common stock variances: Value, leverage and interest rate effects. *Journal of Financial Economics*, 10(4), 407–432. [https://doi.org/10.1016/0304-405x\(82\)90018-6](https://doi.org/10.1016/0304-405x(82)90018-6)
- Conlon, T., Corbet, S., & McGee, R. J. (2021). Inflation and cryptocurrencies revisited: A time-scale analysis. *Economics Letters*, 206, 109996. <https://doi.org/10.1016/j.econlet.2021.109996>
- Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236. <https://doi.org/10.1080/713665670>
- Corbet, S., Meegan, A., Larkin, C., Lucey, B., & Yarovaya, L. (2018). Exploring the dynamic relationships between cryptocurrencies and other financial assets. *Economics Letters*, 165, 28–34. <https://doi.org/10.1016/j.econlet.2018.01.004>
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13, 253–263. <https://doi.org/10.1080/07350015.1995.10524599>
- Fan, J., & Yao, Q. (2003). *Nonlinear time series: nonparametric and parametric methods*. Springer.
- Feng, W., Wang, Y., & Zhang, Z. (2018). Informed trading in the Bitcoin market. *Finance Research Letters*, 26, 63–70. <https://doi.org/10.1016/j.frl.2017.11.009>
- Gandal, N., & Halaburda, H. (2014). *Competition in the cryptocurrency market*. Bank of Canada.
- Garay, J., Kiayias, A., & Leonardos, N. (2024). The bitcoin backbone protocol: Analysis and applications. *Journal of the ACM*, 71(4), 1–49. <https://doi.org/10.1145/3653445>
- Granger, C. W. J., & Teräsvirta, T. (1993). *Modelling nonlinear economic relationships*. Oxford University Press.
- Hamilton, J. D. (1994). *Time series analysis*. Princeton University Press.
- Hansen P. R., Lunde A., & Nason J. M. (2011). The model confidence set. *Econometrica*, 79(2), 453–497. <https://doi.org/10.3982/ecta5771>
- Hu, Z., Zhao, Y., & Khushi, M. (2021). A survey of forex and stock price prediction using deep learning. *Applied System Innovation*, 4(1), 9. <https://doi.org/10.3390/asi4010009>
- Jalal, R. N. U. D., Alon, I., & Paltrinieri, A. (2025). A bibliometric review of cryptocurrencies as a financial asset. *Technology Analysis & Strategic Management*, 37(4), 432–447. <https://doi.org/10.1080/09537325.2021.1939001>

- Jiang, W. (2021). Applications of deep learning in stock market prediction: Recent progress. *Expert Systems with Applications*, 184, 115537. <https://doi.org/10.1016/j.eswa.2021.115537>
- Karasiński, J. (2023). The adaptive market hypothesis and the return predictability in the cryptocurrency markets. *Economics and Business Review*, 9(1), 94–118. <https://doi.org/10.18559/eb.2023.1.4>
- Lipton, A. (2021). Cryptocurrencies change everything. *Quantitative Finance*, 21(8), 1257–1262. <https://doi.org/10.1080/14697688.2021.1944490>
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36(4), 394–419. <https://doi.org/10.1086/294632>
- Mills, T. C., & Markellos, R. N. (2008). *The econometric modelling of financial time series* (2nd ed.). Cambridge University Press.
- Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash System*. <https://bitcoin.org/bitcoin.pdf>
- Rachev, S. T., Mittnik, S., Fabozzi, F. J., Focardi, S. M., & Jasic, T. (2005). *Financial econometrics: From basics to advanced modeling techniques*. John Wiley & Sons.
- Sewell, M. (2011). *Characterization of financial time series*. UCL Department of Computer Science.
- Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied Soft Computing*, 90, 106181. <https://doi.org/10.1016/j.asoc.2020.106181>
- Taylor, S. J. (2008). *Modelling financial time series*. World Scientific.
- Tsay, R. S. (2005). *Analysis of financial time series* (2nd ed.). John Wiley & Sons.
- Wątarek, M., Kwapiień, J., & Drożdż, S. (2023). Cryptocurrencies are becoming part of the world global financial market. *Entropy*, 25(2), 377. <https://doi.org/10.3390/e25020377>
- Wronka, C. (2022). Money laundering through cryptocurrencies-analysis of the phenomenon and appropriate prevention measures. *Journal of Money Laundering Control*, 25(1), 79–94. <https://doi.org/10.1108/jmlc-02-2021-0017>
- Yue, Y., Li, X., Zhang, D., & Wang, S. (2021). How cryptocurrency affects economy? A network analysis using bibliometric methods. *International Review of Financial Analysis*, 77, 101869. <https://doi.org/10.1016/j.irfa.2021.101869>
- Zetzsche, D. A., Buckley, R. P., Arner, D. W., & Föhr, L. (2017). *The ICO gold rush: It's a scam, it's a bubble, it's a super challenge for regulators*. University of Luxembourg Law Working Paper, 11/2017. <https://doi.org/10.2139/ssrn.3072298>
- Zhang, C., Sjarif, N. N. A., & Ibrahim, R. (2024). Deep learning models for price forecasting of financial time series: A review of recent advancements: 2020–2022. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 14(1), e1519. <https://doi.org/10.1002/widm.1519>
- Zhang, J., Cai, K., & Wen, J. (2024). A survey of deep learning applications in cryptocurrency. *Iscience*, 27(1), 108509. <https://doi.org/10.1016/j.isci.2023.108509>
- Zhang, L., Aggarwal, C., & Qi, G. J. (2017, August). Stock price prediction via discovering multi-frequency trading patterns. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 2141–2149). <https://doi.org/10.1145/3097983.3098117>