Research in International Business and Finance 89 (2026) 103486

Contents lists available at ScienceDirect

Research in International Business and Finance

journal homepage: www.elsevier.com/locate/ribaf

Using on-chain data to predict Bitcoin cycles

Klaus Grobys [a,b], Sebastian Näsman [b,*], Davide Sandretto [c]

[a] Innovation and Entrepreneurship (InnoLab), University of Vaasa, Wolffintie 34, Vaasa 65200, Finland
[b] School of Accounting and Finance, University of Vaasa, Wolffintie 34, Vaasa 65200, Finland
[c] Department of Management, University of Turin, Corso Unione Sovietica 218 bis, Torino 10134, Italy

ARTICLE INFO

JEL Classification:
G11
G12
G13
G40

Keywords:
On-chain data
Market cycles
Behavioral finance
Bitcoin

ABSTRACT

There is limited literature studying the predictive power of on-chain data for Bitcoin price cycles. This paper contributes to this literature by assessing whether three on-chain, trading-based measures help predict the Bitcoin price time series across three major market cycles. We find that these indicators outperform both a buy-and-hold benchmark and random-entry strategies simulated through Monte Carlo analysis. For example, the Sharpe ratio increases from 0.45 for the buy-and-hold benchmark to 1.28 when using the Market Value to Realized Value Z-score measure. This study contributes to the literature by showing that blockchain-based behavioral data provides predictive value in decentralized markets that lack intrinsic valuation anchors. The findings also have practical implications for investors, traders, and regulators, and they challenge traditional notions of market efficiency by providing evidence of recurring behavioral patterns embedded in publicly observable blockchain activity.

1. Introduction

Traditional financial models have difficulty explaining the price movements of crypto assets (Dunbar and Owusu-Amoako, 2022; Liu et al., 2022; Shen et al., 2020). Unlike equities, most cryptocurrencies lack intrinsic value, leading researchers to develop alternative approaches to describe market dynamics, including investor sentiment measures. In this regard, one of the most promising sources of information is on-chain data, which is recorded directly on the blockchain and reflects the actions and sentiment of network participants in near real time (Jagannath et al., 2021). In traditional markets, sentiment is typically inferred from indirect proxies such as surveys or media-based indicators (Bollen et al., 2011; Brown, 1999). In contrast, blockchains provide a transparent and tamper-resistant ledger of transactions, offering a verifiable record of investor behavior that can be used to derive sentiment from revealed behavior rather than self-reported attitudes.

As cryptocurrencies are largely sentiment-driven (Aslanidis et al., 2024; Jia et al., 2022), the on-chain data's ability to directly measure sentiment becomes even more critical. Changes in earnings estimates or cash flows do not drive the bull and bear markets in cryptocurrencies. Instead, they are driven by emotions like greed and fear. This feature motivates the use of sentiment analysis and supports on-chain data as a direct behavioral lens on market conditions. Moreover, the documented inefficiencies of cryptocurrency markets make predictive analysis especially pertinent. Prior evidence suggests indeed that crypto markets are far from informationally efficient (Vidal-Tomás et al., 2019) and that even relatively simple technical trading strategies can deliver abnormal performance (Grobys et al., 2020).

* Corresponding author.
E-mail addresses: klaus.grobys@uwasa.fi (K. Grobys), nasmansebastian@gmail.com (S. Näsman), davide.sandretto@unito.it (D. Sandretto).

https://doi.org/10.1016/j.ribaf.2026.103486
Received 12 February 2026; Received in revised form 11 May 2026; Accepted 12 May 2026
Available online 15 May 2026

In this regard, the present study investigates the ability of three simple on-chain, rule-based trading measures to capture Bitcoin price cyclicality over the period from December 7, 2013 to April 12, 2025, which spans three complete market cycles (2015, 2018, and 2022). The three indicators are the Net Unrealized Profit/Loss (NUPL) ratio, the Market Value to Realized Value Z-score (MVRV Z-score), and Cumulative Value Days Destroyed (CVDD). The first two metrics relate prices to holders' aggregate cost basis (realized value) and can be interpreted through behavioral finance mechanisms. NUPL quantifies the extent to which the investors are sitting on unrealized gains or losses by comparing market capitalization to realized capitalization, where realized capitalization reflects the value of coins at their last on-chain movement. High NUPL values indicate that the representative holder is largely in profit, a condition typically associated with stronger investor sentiment, lower perceived risk, and a greater role for greed (Demeester et al., 2019). In such phases, gains can encourage risk-taking and trend-following behavior (fear of missing out), while the accumulation of profits may also create potential selling pressure, increasing the market's sensitivity to negative shocks. Low or negative NUPL values indicate that investors are holding unrealized losses, which is consistent with weak sentiment, heightened risk aversion, and a greater role for fear. MVRV Z-score, in turn, standardizes the deviation of market value from realized value and can be interpreted as mispricing relative to an aggregate cost-basis anchor. High MVRV Z-scores reflect market valuations far above the cost basis, consistent with overvaluation and very optimistic sentiment, whereas low MVRV Z-scores indicate valuations close to or below cost basis, consistent with undervaluation and weak sentiment across cycles (Mahmudov and Puell, 2018). CVDD differs conceptually, as it reflects long-term holder behavior. By weighting coin movements by their "age," it captures the spending of long-held coins and therefore provides information about capitulation by long-term holders during periods of extreme pessimism. For this reason, CVDD is commonly viewed as a long-horizon bottom-related signal, consistent with the idea that long-term holder activity helps establish a floor during market drawdowns.

This study makes several key contributions to the growing literature on cryptocurrency research. First, we contribute to the current debate on the cryptocurrency market efficiency. Several studies argue that cryptocurrencies are informationally inefficient (Aggarwal, 2019; Al-Yahyaee et al., 2020; Khuntia and Pattanayak, 2018; Urquhart, 2016) while others support the efficient market hypothesis (Sensoy, 2019; Tiwari et al., 2018). In this context, a possible finding that on-chain trading strategies generate meaningful abnormal performance would provide empirical evidence that further supports the inefficiency view.

Second, we extend the literature on technical trading rules in cryptocurrency markets. For example, Grobys et al. (2020) showed that simple moving-average strategies can generate positive excess returns after controlling for the average market return. Hudson and Urquhart (2021) evaluated almost 15,000 technical trading rules using data on Bitcoin and three other popular cryptocurrencies and documented significant predictability and profitability across rule classes and assets. Using high-frequency data, Corbet et al. (2019) also found support for moving-average-based strategies. We build on this evidence by assessing the performance of on-chain trading strategies, rather than conventional approaches, hence exploiting signals that are specific to cryptocurrency markets thanks to the blockchain technology.

Finally, we contribute to the still limited literature that examines the price predictability of cryptocurrencies using on-chain data. Only a small number of studies incorporate on-chain characteristics to explain crypto-asset price movements. Casella and Paletto (2023) show how on-chain data can be used to detect cryptocurrency market regimes, such as bull and bear phases, using machine learning techniques. Jagannath et al. (2021) employ machine learning models to assess whether on-chain variables provide useful additional inputs for predicting Ethereum prices. Similarly, Kim et al. (2022) analyze the relationship between Bitcoin prices and on-chain data using machine learning approaches. More recent studies extend this line of research by combining large sets of on-chain variables with feature-selection methods and machine-learning or deep-learning models to predict short-horizon Bitcoin price movements, often reporting strong predictive performance (Khosravi and Ghazani, 2023; Omole and Enke, 2024, 2025; Dubey and Enke, 2025). Interestingly, Sakkas and Urquhart (2024) study crypto risk compensation using on-chain characteristics. In contrast to this strand of the literature, which is primarily concerned with short-horizon predictive accuracy in high-dimensional modeling settings, we examine whether a small set of economically interpretable on-chain indicators can be used to construct profitable trading strategies for identifying Bitcoin market cycles.

Our findings show that all strategies consistently outperform the buy-and-hold benchmark, and that the CVDD strategy performs better than most of the random-entry strategies simulated through Monte Carlo analysis. The MVRV Z-score emerges as the best-performing approach, delivering higher returns and better risk-adjusted performance than both the benchmark and the NUPL-based strategies across all versions tested. Moreover, the individual trades generated by each strategy also outperform the buy-and-hold benchmark. Overall, these results suggest that on-chain, trading-based strategies capture investor sentiment and can help to predict Bitcoin price cycles. These findings hold important theoretical contributions, further corroborating the inefficiency of cryptocurrencies, and have practical implications for investors and regulators.

The remainder of the paper is organized as follows. The next section reviews the literature on market efficiency and investor sentiment. Section 3 describes the data and methodology, Section 4 presents the empirical results, Section 5 discusses the findings and limitations and the final section concludes.

2. Literature review

2.1. Efficiency of crypto market

The Efficient Market Hypothesis (EMH), first formalized by Fama (1970), suggests that financial markets are informationally efficient, meaning that all available information is already reflected in asset prices. Fama proposed a three-tier classification system—weak-form efficiency, semi-strong efficiency, and strong efficiency—depending on the degree to which information is incorporated into asset prices. This leads to the implication that achieving abnormal returns consistently through fundamental or technical analysis is impossible. The EMH also assumes that investors act rationally and maximize their expected utility, as formalized by Expected Utility Theory (EUT) (Von Neumann and Morgenstern, 1944). However, both assumptions have long been challenged. Empirical findings show that investors do not behave as EUT predicts, and that cognitive biases and emotional responses systematically distort decision-making. Prospect Theory, developed by Kahneman and Tversky (1979), offers an alternative framework, documenting how investors' irrational investment decisions generate predictable market inefficiencies (Barberis et al., 1998; De Bondt and Thaler, 1985).

While it is still an ongoing debate whether stock markets satisfy weak-form efficiency (Verheyden et al., 2015), in the cryptocurrency context it is clear that markets are far from being efficient. Their unregulated nature already makes information dissemination very different from traditional markets and leads to a lower degree of efficiency (Nimalendran et al., 2025). For instance, information sharing does not follow strict reporting standards and instead relies on social media, news websites, and influencers.

Many studies provide empirical evidence that crypto markets are inefficient. For example, Urquhart (2016) studied the market efficiency of Bitcoin using different tests and found that it is weakly efficient. Caporale et al. (2018) identified high return persistence patterns for four digital coins, highlighting market inefficiencies. Moreover, Vidal-Tomás et al. (2019) analyzing the overall cryptocurrency market, documented that it is weak-form inefficient due to the behavior of altcoins and that the creation of new cryptocurrencies over time has not significantly changed market efficiency. Consistently, Balcilar et al. (2017) found short-term autocorrelation in cryptocurrencies, allowing traders to profit from momentum-based strategies. Lahmiri and Bekiros (2019) also found that past Bitcoin prices influence future prices, strongly contradicting weak-form efficiency. Naeem et al. (2021) further documented that cryptocurrency prices are strongly influenced by social media sentiment, which causes investors to act irrationally rather than following objective financial data and metrics. This and other evidence indicate that cryptocurrency markets are far from strong-form efficiency. While traditional markets can be argued to exhibit varying degrees of efficiency, cryptocurrency markets can be seen as largely inefficient in all three forms. Indeed, historical price data has predictive power (violating weak-form efficiency), public information is not fully reflected in prices (violating semi-strong efficiency), and insider trading leads to consistent abnormal returns (violating strong-form efficiency). These inefficiencies create opportunities for traders, which will be investigated through on-chain analysis in this study.

2.2. Investor sentiment

The aggregate of pessimism and optimism, driven by emotions as well as behavioral and psychological biases, is referred to as investor sentiment and it influences individuals' trading and investment decisions (Huynh et al., 2025). Investor sentiment plays a crucial role in financial markets because it can push prices far away from their fundamental value (Baker and Wurgler, 2007). For instance, over-optimism may trigger excessive risk-taking and price run-ups that can evolve into bubbles, which then burst when investors panic and prices fall well below intrinsic value (Pan, 2020). While fundamentals may drive prices in the long run, sentiment tends to generate short-term fluctuations around those fundamentals (Ung et al., 2024).

Unlike corporate fundamentals such as earnings or revenue, which are relatively easy to quantify, investor sentiment is inherently psychological and behavioral, making it difficult to measure directly (Bollen et al., 2011; Da et al., 2011). Measurement is further complicated because investors may not report their feelings truthfully due to social desirability, shame, or their own behavioral biases (Baker and Wurgler, 2007; Brown and Cliff, 2005). Despite these challenges, many different measurement methods have emerged over the years.

A traditional method relies on survey-based indicators, which directly ask investors or market professionals about their expectations. Among the best-known surveys there are the AAII Sentiment Survey, Investor Intelligence, and the UBS/Gallup polls (Brown, 1999; Bu, 2023). For example, Brown and Cliff (2005) used the Investor Intelligence survey to capture sentiment and studied its relationship with asset valuation. More recent approaches exploit digital traces of attention and beliefs, such as search engine data and social media activity. Da et al. (2011) showed that abnormal increases in search volume predict short-term outperformance, while Chen et al. (2014) analyzed investor opinions expressed on social media and their ability to forecast future stock returns and earnings surprises. Similarly, Bollen et al. (2011) used Twitter content to extract mood indicators and documented a significant relationship between these mood measures and the performance of the Dow Jones Industrial Average.

In the cryptocurrency context, sentiment has also been proxied using a range of methods, with search engine activity and social media being among the most prominent sources (Burggraf et al., 2020; Kristoufek, 2013; Long et al., 2025). Kristoufek (2013) examined Bitcoin attention using Google and Wikipedia search data and found a positive association with prices, highlighting a bidirectional relationship between attention and returns. Burggraf et al. (2020) developed a FEARS index based on negative Google search terms (in the spirit of Da et al., 2011) to proxy investor anxiety about micro and macroeconomic concerns, and they found a significant negative relationship between this index and Bitcoin prices, especially for microeconomic concerns. Long et al. (2025) constructed a sentiment index using posts from Twitter and Reddit and investigated its link with short- and long-term cryptocurrency volatility. Overall, this evidence supports the idea that sentiment can help predict cryptocurrency price dynamics in a market where traditional valuation models are ineffective.

However, compared with the abovementioned sentiment proxies, blockchain technology introduces a distinctive source of information for sentiment analysis: on-chain data. Every transaction recorded in the ledger reflects an investor's decision, and each of these decisions is influenced by emotions. This makes on-chain metrics sentiment-driven, enabling them to capture market participants' collective moods and expectations without biases (Jagannath et al., 2021). On-chain data comprises all transactions and interactions permanently recorded on a blockchain. This data is transparent and publicly accessible, and because of the distributed ledger technology, the information cannot be altered.

This paper takes a different perspective from the existing literature and leverages on-chain information to assess the performance of trading strategies based on sentiment signals. Several commonly used metrics help interpret network activity and investor behavior (Casella and Paletto, 2023). Measures of network usage and participation include active addresses and transaction counts, which indicate how many unique participants are transacting and how frequently the network is used. Indicators linked to capital flows and market pressure include exchange inflows and outflows, often interpreted as potential selling pressure when coins move to exchanges, and accumulation signals when coins leave exchanges to cold storage. Metrics capturing economic value and profitability include realized value (realized capitalization), which values coins based on the last time they moved. For miner-related dynamics (in proof-of-work networks), miner revenue can matter because declining revenues may incentivize miners to sell holdings, consequently increasing selling pressure. Finally, for network conditions and security, gas fees on Ethereum reflect demand for block space and smart contract execution (often rising when usage is high), while hash rate in proof-of-work systems is a key proxy for the network's computational security.

A few studies that use on-chain data to study crypto movements are Casella and Paletto (2023), Jagannath et al. (2021), and King et al. (2024). Jagannath et al. (2021), using metrics such as transaction volume and active addresses, showed how they produce more accurate short-horizon price predictions. They argued on-chain metrics can complement traditional indicators by capturing underlying system-level demand and behavior, which provide information about the network's demand. Similarly, Casella and Paletto (2023) examined how on-chain data can be used to identify cryptocurrency market regimes and how forecasting these indicators can support optimal long-term asset allocation decisions.

3. Data and methodology

All trading simulations are conducted using daily closing price data for Bitcoin, obtained from Investing.com, covering the period from December 7, 2013, to April 12, 2025. This sample window is selected because it spans three complete Bitcoin market cycles (2015, 2018, and 2022). In addition, the 7th of December 2013 corresponds to the first date on which Bitcoin had declined by at least 50% from the 2013 cycle peak. We therefore adopt this date as a logical and relatively optimistic entry point for the buy-and-hold benchmark. Table A1 presents the timestamp data for each strategy's entry and exit points.

Data for the on-chain indicators is sourced from bitcoinmagazinepro.com. We evaluate the profitability of three trading strategies based on the NUPL ratio, the MVRV Z-score and CVDD. All strategies rely only on historical on-chain data, making them implementable in practice and thus ruling out any potential look-ahead bias. Transaction costs are not considered in the analysis since the strategies involve only a small number of trades (3 entries and 3 exits) over their full investment horizon, and no short positions are taken, so associated costs are negligible.[1] Moreover, when the NUPL ratio and MVRV Z-score strategies are out of the market, we assume a zero return on cash, which is a conservative assumption. This allows us to compare the performance of the trading strategies coherently with that of the benchmark portfolio.

3.1. NUPL ratio

The NUPL ratio approximates the share of coins that are currently held at an unrealized profit or loss and is commonly used to identify potential market tops and bottoms. High values indicate substantial unrealized gains—often associated with euphoric market conditions—whereas low values are typically linked to fear and capitulation.

Following Demeester et al. (2019), the NUPL ratio is computed using market capitalization and realized capitalization as:

$$NUPL = \frac{Market\ Cap \quad - Realized\ Value}{Market\ Cap} \tag{1}$$

Market capitalization is calculated as the Bitcoin price multiplied by the number of coins in circulation. Realized capitalization (realized value) is computed by valuing each coin at the price at which it was last transferred on-chain and summing across all coins in circulation. Subtracting realized capitalization from market capitalization yields the aggregate unrealized profit/loss, which is then scaled by market capitalization to obtain the NUPL ratio.

From a behavioral finance perspective, the NUPL ratio can be interpreted as a proxy for investor sentiment and for recurrent cognitive patterns such as loss aversion, anchoring, and herding. When the ratio is elevated (e.g., above 0.75), market participants may remain anchored to unrealized gains and delay profit-taking, consistent with heightened optimism, greed, and fear of missing out (FOMO). Conversely, negative figures are indicative of widespread unrealized losses and may coincide with capitulation dynamics, in which investors sell under stress after sustained drawdowns. The tendency for these behavioral regimes to recur across cycles motivates the use of NUPL ratio as a signal for systematic trading rules.

We design three threshold-based trading strategies using the NUPL ratio to assess the practical applicability of NUPL-based signals for market timing. Each NUPL strategy takes a long position when the NUPL ratio falls below 0 and exits the position at a strategy-specific upper threshold. The exit thresholds are 0.67 for the NUPL 1 strategy (conservative), 0.70 for NUPL 2 (moderate), and

[1] An unreported analysis accounting for a conservative transaction cost of 2% per trade confirms that performance metrics remain qualitatively unchanged, with returns and Sharpe ratios only marginally lower than those reported in the main tables. Results are available from the authors upon request.

0.73 for NUPL 3 (aggressive). The entry threshold of 0 and the central exit threshold of 0.70 correspond to the levels suggested by the indicator creators. The entry level is grounded in the economic interpretation of NUPL as a signal of widespread capitulation, when aggregate losses across market participants have likely exhausted downward pressure. The exit level instead identifies the zone in which accumulated unrealized profits become sufficiently large to trigger a behavioral shift toward profit protection, generating the selling pressure that historically precedes cycle reversals. The alternative thresholds of 0.67 and 0.73 allow us to test the sensitivity and robustness of performance to alternative profit-taking levels.

Across the sample, each strategy generates one entry and one exit per cycle, resulting in a total of three trades per strategy. However, during the final cycle (2022–2025), none of the strategies reached their respective exit thresholds. Consequently, all three strategies remain in the same ongoing position. This reduces cross-strategy differentiation, as the third trade is identical across the three specifications.

3.2. MVRV Z-score

The MVRV Z-score is a widely used on-chain metric designed to assess whether a coin is undervalued or overvalued relative to its “fair value.” It combines three key components. Two coincide with those used in the NUPL ratio, namely market value (MV) and realized value (RV). The third component is the Z-score, which standardizes the deviation between MV and RV by the standard deviation of market value. The measure builds on the MVRV framework first introduced by Mahmudov and Puell (2018).

The MVRV Z-score is computed as:

$$Z = \frac{MV - RV}{\sigma(MV)} \tag{2}$$

where MV denotes market value, defined as Bitcoin price multiplied by circulating supply, RV denotes realized value, defined as the aggregate value of coins priced at their last on-chain movement and $\sigma(MV)$ denotes the standard deviation of market value.

During bull market phases, when Bitcoin’s market value rises substantially above realized value, the indicator suggests that market participants are holding large unrealized gains. This condition is typically associated with elevated speculative activity, FOMO, and heightened media attention. Conversely, during bear markets, when market value falls below realized value, the indicator points to capitulation dynamics and extreme investor pessimism.

To further examine the predictive content of the MVRV Z-score, we construct and backtest three rule-based trading strategies over three Bitcoin cycles spanning 2013–2025. These tests serve as a robust approach to prove whether the strategy’s performance is systematic and repeatable across cycles, rather than being driven by cycle-specific conditions or chance. All strategies take a long position when the MVRV Z-score falls below -0.2, consistent with the mid-entry region reported on bitcoinmagazinepro.com and indicative of heightened fear and uncertainty. The entry level identifies the condition in which market value has fallen significantly below realized value, signaling that the average market participant is holding unrealized losses and that capitulation dynamics are sufficiently advanced to represent a favorable entry point. The exit thresholds of 5, 6, and 7 identify the zone in which market value has deviated so far above realized value that the average participant holds large unrealized gains, generating the behavioral pressure toward profit-taking that historically coincides with cycle tops. Exit thresholds vary across strategies to reflect different risk preferences and to provide robustness checks. The conservative specification exits when the MVRV Z-score reaches at least 5, the moderate specification exits when it reaches at least 6, and the aggressive specification exits when it reaches at least 7.

Across the sample, the MVRV-based strategies also generate three trades per specification. As with the NUPL ratio case, the final trade remains open at the end of the sample period, implying that the third trade is identical across the three MVRV Z-score strategies.

3.3. Cumulative value days destroyed (CVDD)

The third on-chain metric selected for more detailed analysis is the CVDD, introduced by Woo (see Woo and Puell, 2019). CVDD is regarded as a reliable indicator for identifying long-term market bottoms, as it combines both time-based and value-based information from Bitcoin’s transaction history. Importantly, CVDD does not predict tops or short-term movements, but it detects zones of extreme undervaluation for long-term position building. This means it is not for active traders but rather for investors wanting to find buying opportunities.

The series is publicly available through platforms such as bitcoinmagazinepro.com. CVDD is built on Coin Days Destroyed (CDD), a metric that weights transactions by both the amount of coins moved and how long they were held. CDD is computed as the number of coins transferred multiplied by the number of days since those coins were last moved. For example, transferring 10 BTC after 100 days destroys 1000 coin days. Because longer-held coins receive greater weight, the metric is often interpreted as capturing long-term holder behavior and, by extension, market sentiment during stress periods.

CVDD extends CDD by aggregating this activity over time in value terms. Operationally, CVDD is obtained by taking the cumulative sum of the USD value of coin days destroyed, normalizing by the age of the Bitcoin network (in days), and applying a scaling factor (commonly around six million) to align the series with Bitcoin’s price level for comparison.

Mathematically, the metric can be approximated as:

$$CVDD = \frac{\sum(Value\ of\ BTC\ Transferred(USD)x\ Coin\ Days\ Destroyed)}{Martket\ Age\quad (in\ days)}\ x\ Scaling\ Constant \tag{3}$$

The CVDD curve is often interpreted as a long-term valuation floor that price approaches during periods of extreme negative sentiment, when even long-term holders capitulate. Historically, Bitcoin's price has tended to converge toward CVDD near major cycle lows, motivating its use as a bottom-detection benchmark.

To test CVDD's bottom-prediction ability, we employ a methodology different from that used for NUPL ratio and MVRV Z-score. The entry signal is triggered when Bitcoin's price approaches the CVDD line, reflecting the economic interpretation of CVDD as the lower bound of Bitcoin's fair value. At this level, the cost basis of long-term holders who have historically sold at the worst possible times represents a floor below which sustained price declines become increasingly unlikely, as further selling pressure is exhausted. A long entry is triggered in day t_c^{CVDD} when Bitcoin's price is within 1% of (or closer to) the CVDD line for each cycle c. The strategy return is then the buy-and-hold log return from t_c^{CVDD} to a common exit date $t_c^{\max}$ (the cycle peak, or the last available observation for the ongoing cycle):

$$R_c^{CVDD} = \ln\left(\frac{P_{t_c^{\max}}}{P_{t_c^{CVDD}}}\right) \tag{4}$$

We then benchmark the CVDD-timed entry against uninformed timing using Monte Carlo simulations, run separately for each of the three cycles in our sample. For each cycle c, we generate $N = 100$ random entry dates by drawing uniformly from a window around t_c^{CVDD}. To capture different degrees of investor timing uncertainty and provide robust tests, we consider three window widths:

- $t_c^{CVDD} \pm 50$ trading days
- $t_c^{CVDD} \pm 75$ trading days
- $t_c^{CVDD} \pm 100$ trading days

For each random draw i, we compute the same buy-and-hold return from the random entry date to the same cycle exit date:

$$R_{c,w}^{(i)} = \ln\left(\frac{P_{t_c^{\max}}}{P_{t_{c,w}^{(i)}}}\right) \tag{5}$$

Within each cycle, all trades share the same exit $t_c^{\max}$, so differences in performance are driven only by entry timing. Indeed, CVDD is designed exclusively as a bottom-detection indicator and does not generate exit signals. We set $t_c^{\max}$to the cycle peak, while for the ongoing cycle (2022–), we use the final closing price in our dataset (April 12, 2025). Importantly, the CVDD-identified bottom is generated using only information available at the time the signal occurs. As a result, the strategy is fully implementable in real time.

3.4. Buy-and-hold benchmark

We compare the payoffs of the on-chain strategies with a buy-and-hold benchmark (B&H). The benchmark is constructed by taking a long position on Bitcoin on the 7th of December 2013, which is the first date on which Bitcoin's price had declined by at least 50% from the 2013 cycle peak. We considered alternative entry definitions for the buy-and-hold investor, including using the same entry dates generated by the strategy signals. However, such an approach would mechanically provide buy-and-hold with unusually favorable timing, since a typical investor is unlikely to enter consistently near cycle lows. Moreover, retail participation often intensifies closer to cycle peaks, which could make even the 7th of December 2013 a relatively optimistic entry point for the average investor.

We nonetheless adopt the 7th of December 2013 because it provides a simple and transparent entry date that is independent of the trading signals and does not weaken the benchmark. This choice helps ensure that any performance differences are not driven by an artificially advantageous benchmark construction and supports the reliability of the results.

To compare the performance of the on-chain strategies with the benchmark, we report several performance measures, including

Table 1
NUPL ratio strategies performance.

Strategy	NUPL 1	NUPL 2	NUPL 3	Buy & Hold
Cum. Log Return	4.77	5.19	7.12	4.79
Cum. Simple Return	117.04	178.04	1232.93	118.94
Ann. Log Return	0.42***	0.46***	0.63***	0.42
t-stat	(2.75)	(2.98)	(3.77)	(1.51)
Ann. Volatility	0.52	0.52	0.56	0.94
Min	-0.50	-0.50	-0.50	-0.85
Max	0.24	0.24	0.24	1.47
Kurtosis	41.87	41.28	32.45	242.33
Sharpe ratio	0.82	0.88	1.12	0.45
ΔSharpe ratio *t*-stat	15.20***	17.78***	26.17***	-
Holding Days	4143	4143	4143	4143

This table reports the descriptive statistics and performance measures for the three different NUPL ratio strategies compared to the buy-and-hold benchmark. The ΔSharpe ratio *t*-test is computed consistent with Opdyke (2007).

cumulative and annualized log returns, annualized volatility, minimum and maximum daily returns, and the Sharpe ratio.

4. Results

In this section, we present diagnostic results for our strategies. Importantly, all statistical inference is conducted using the full daily strategy return series. Table 1 reports the results for the NUPL-based strategies. The findings indicate superior risk-adjusted performance for all three NUPL strategies relative to the B&H benchmark, highlighting the robustness of the outcomes obtained. The most aggressive strategy, NUPL 3, achieves an annualized log return of 0.63, significantly higher than the benchmark's 0.42. It also outperforms the benchmark in risk-adjusted returns, with a Sharpe ratio of 1.12, and in cumulative simple returns. Similar outperformance is observed for NUPL 1 and NUPL 2. Although their absolute performance is slightly lower than NUPL 3, both strategies still exhibit higher Sharpe ratios and annualized log returns than the benchmark, with the exception of NUPL 1, which has the same annualized log return. Moreover, all trading strategies display smaller drawdowns (−0.50) than the B&H (−0.85) and avoid extreme returns, as indicated by lower kurtosis values ranging from 32.45 to 41.87.

To test statistically whether the difference between each strategy's Sharpe ratio and the B&H Sharpe ratio is greater than zero, we follow the approach described in Opdyke (2007). All NUPL Sharpe ratio differences are highly statistically significant, with t-statistics of 15.20, 17.78, and 26.17 for NUPL 1, NUPL 2, and NUPL 3, respectively. We therefore reject the null hypothesis that the Sharpe ratio difference is less than or equal to zero.

In Table 2, we compare each NUPL trade with the benchmark. Consistent with the previous results, all NUPL trades outperform the B&H strategy, delivering higher annualized log returns and lower annualized volatility. Overall, these findings support once again the reliability of the NUPL indicator across cycles under different risk profiles and highlight the ratio's ability to consistently identify favorable entry and exit signals.

Moving to the performance of the MVRV Z-score strategies, Table 3 reports the corresponding results. Similar to the NUPL strategies, all three MVRV Z-score strategies exhibit superior and robust risk-adjusted performance relative to the B&H benchmark. The most aggressive strategy, MVRV Z-score 3, achieves an annualized log return of 0.71, significantly higher than the benchmark's 0.42. It also outperforms the benchmark in risk-adjusted returns, with a Sharpe ratio of 1.28, and in cumulative returns. Overall, all MVRV strategies deliver higher annualized log returns, Sharpe ratios, and cumulative returns than the benchmark. They also outperform the NUPL-based strategies across all metrics, although MVRV Z-score 2 displays higher volatility than its corresponding NUPL counterpart.

Again, we test whether the difference between each MVRV strategy's Sharpe ratio and the B&H Sharpe ratio is greater than zero. All the mentioned differences are highly statistically significant, with t-statistics of 22.58, 28.64, and 31.23 for MVRV 1, MVRV 2, and MVRV 3, respectively. In Table 4, each MVRV Z-score trade outperforms the B&H strategy, mirroring the evidence for NUPL ratio. This suggests that MVRV Z-score strategies are as reliable as the NUPL strategies, while delivering higher risk-adjusted and cumulative returns.

Finally, Table 5 reports the Monte Carlo simulation results for the CVDD strategies. The evidence highlights CVDD's ability to identify cycle bottoms. Across all trades and window ranges, CVDD outperforms most randomly timed entries. In terms of cumulative log returns, CVDD attains an average p-value of 99%, indicating an optimal ability to capture cycle bottoms. Annualized log returns are lower, with an average p-value of 82%. This suggests that although CVDD typically enters very close to the bottom, its holding periods are sometimes longer than those of certain random-entry simulations, which reduces annualized performance. Nevertheless, in Cycle 1 (2013–2017), CVDD achieves a 100% p-value for annualized log returns in the ±75-day range and nearly 100% in the other two ranges.

Overall, these results suggest that different dimensions of sentiment and market structure are captured by each metric, depending on its design. The NUPL ratio reflects short-term sentiment in real time, the MVRV Z-score captures major value dislocations, and CVDD focuses on longer-run structural time-value exhaustion. Despite these differences, it is possible to conclude that the sentiment leaves measurable, traceable footprints in on-chain data that can be used to anticipate turning points in Bitcoin's cyclical market.

Table 2
NUPL ratio trades performance.

Trade	NUPL Ann. Log Return	NUPL Ann. Volatility	Buy & Hold Ann. Log Return	Buy & Hold Ann. Volatility	Outperformed
NUPL1 1	0.67	0.63	0.42	0.94	Yes
NUPL1 2	0.80	0.76	0.42	0.94	Yes
NUPL1 3	0.46	0.51	0.42	0.94	Yes
NUPL2 1	0.75	0.64	0.42	0.94	Yes
NUPL2 2	0.89	0.76	0.42	0.94	Yes
NUPL2 3	0.46	0.51	0.42	0.94	Yes
NUPL3 1	1.17	0.71	0.42	0.94	Yes
NUPL3 2	0.98	0.76	0.42	0.94	Yes
NUPL3 3	0.46	0.51	0.42	0.94	Yes

This table reports the performance measures for the 9 different NUPL ratio trades compared to the buy-and-hold benchmark.

Table 3
MVRV Z-score strategies performance.

Strategy	MVRV Z-score 1	MVRV Z-score 2	MVRV Z-score 3	Buy & Hold
Cum. Log Return	5.81	7.40	8.09	4.79
Cum. Simple Return	333.02	1628.65	3260.62	118.94
Ann. Log Return	0.51***	0.65***	0.71***	0.42
t-stat	(3.42)	(4.02)	(4.29)	(1.51)
Ann. Volatility	0.50	0.55	0.55	0.94
Min	-0.50	-0.50	-0.50	-0.85
Max	0.24	0.24	0.24	1.47
Kurtosis	44.94	35.30	33.03	242.33
Sharpe ratio	1.01	1.19	1.28	0.45
ΔSharpe ratio *t*-stat	22.58***	28.64***	31.23***	-
Holding Days	4143	4143	4143	4143

This table reports the descriptive statistics and performance measures for the three different MVRV Z-score strategies compared to the buy-and-hold benchmark. The ΔSharpe ratio *t*-test is computed consistent with Opdyke (2007).

Table 4
MVRV Z-score trades performance.

Trade	MVRV Z-score Ann. Log Return	MVRV Z-score Ann. Volatility	Buy & Hold Ann. Log Return	Buy & Hold Ann. Volatility	Outperformed
MVRV Z-score 1 1	0.93	0.63	0.42	0.94	Yes
MVRV Z-score 1 2	1.00	0.75	0.42	0.94	Yes
MVRV Z-score 1 3	0.53	0.51	0.42	0.94	Yes
MVRV Z-score 2 1	1.24	0.70	0.42	0.94	Yes
MVRV Z-score 2 2	1.09	0.76	0.42	0.94	Yes
MVRV Z-score 2 3	0.53	0.51	0.42	0.94	Yes
MVRV Z-score 3 1	1.35	0.71	0.42	0.94	Yes
MVRV Z-score 3 2	1.19	0.77	0.42	0.94	Yes
MVRV Z-score 3 3	0.53	0.51	0.42	0.94	Yes

This table reports the performance measures for the 9 different MVRV Z-score trades compared to the buy-and-hold benchmark.

Table 5
CVDD strategies performance.

Strategy	Trade 1	Trade 2	Trade 3
Cum. Log Return	4.75	3.03	1.66
Cum. Simple Return	115.62	20.79	5.26
Ann. Log Return	1.62***	1.04**	0.69**
t-stat	(4.00)	(2.30)	(2.19)
Ann. Volatility	0.69	0.77	0.49
Min	-0.18	-0.50	-0.09
Max	0.24	0.18	0.11
Kurtosis	9.74	26.01	5.32
Sharpe ratio	2.34	1.35	1.41
Holding Days	1069	1061	885
Monte Carlo Simulation (CVDD vs. Random)			
± 50 days Annualized Log Return *p*-value	99%	84%	66%
± 50 days Log Return *p*-value	99%	100%	99%
± 75 days Annualized Log Return *p*-value	100%	71%	76%
± 75 days Log Return *p*-value	100%	100%	97%
± 100 days Annualized Log Return *p*-value	98%	63%	80%
± 100 days Log Return *p*-value	100%	98%	100%

This table reports descriptive statistics and performance measures for the three CVDD trades and the Monte Carlo simulation output. The simulation draws 100 randomly timed entry dates from windows around the cycle bottom, using three alternative ranges: ±50 days, ±75 days, and ±100 days. The simulation p-value indicates the proportion of simulations for which the CVDD strategy's reported metric is higher than the simulated outcomes.

5. Discussion

5.1. How does our study line up with earlier studies?

The present study examines whether three on-chain indicators—NUPL, MVRV Z-score, and CVDD—can be translated into simple rule-based trading strategies capable of identifying Bitcoin market cycles over a long sample spanning three major boom-bust episodes. In contrast to high-dimensional predictive frameworks, our approach is deliberately parsimonious and interpretable. Rather than using large sets of raw blockchain variables within machine-learning architectures, we assess whether a small number of economically grounded indicators can outperform passive and random-timing benchmarks. The results indicate that all three measures contain predictive value, with MVRV Z-score producing the strongest overall risk-adjusted performance and CVDD appearing especially informative for identifying market bottoms.

This contribution is closely related to the recent literature that combines on-chain data with machine-learning and deep-learning methods, but it differs materially in methodological orientation. Casella and Paletto (2023) use stochastic processes and deep-learning models to forecast several on-chain time series in order to identify cryptocurrency market regimes and support long-horizon allocation decisions. Kim et al. (2022) develop a self-attention-based multi-LSTM framework with change-point detection to improve Bitcoin price forecasting. More recent studies by Omole and Enke (2024), Dubey and Enke (2025), and Omole and Enke (2025) adopt even more data-intensive approaches, combining large sets of on-chain features with formal feature-selection procedures and machine-learning or deep-learning models to predict short-horizon Bitcoin price direction or magnitude. Their findings generally show strong predictive performance and substantial trading profitability, especially when feature selection is carefully integrated into the modeling process.

A key commonality across these studies and ours is that they all support the view that on-chain data contain economically meaningful information about Bitcoin market behavior. Whether the objective is market-phase detection or cycle timing, blockchain-based variables repeatedly improve predictive performance relative to more limited benchmark models. In addition, several of the recent studies highlight the importance of variables linked to realized value, unrealized gains and losses, and investor positioning. This is especially relevant for the present paper because NUPL and MVRV Z-score are directly built around those same economic dimensions. Hence, although the modeling frameworks differ sharply, the underlying economic content identified in the machine-learning literature appears broadly consistent with the mechanisms captured by our indicator-based approach.

On the other hand, the recent machine-learning literature is primarily concerned with maximizing short-horizon predictive accuracy through large feature spaces, feature-selection approaches, and complex nonlinear models. By contrast, our study asks whether a small set of interpretable on-chain indicators can identify major cyclical turning points in Bitcoin markets. The comparison should therefore not be framed as one between directly competing models solving the same problem. Rather, the recent literature is mainly focused on predictive optimization, whereas our study is focused on cycle identification and economic interpretability. This distinction is important because it clarifies our contribution: we do not claim to outperform state-of-the-art predictive models in terms of classification accuracy or short-horizon trading profitability; instead, we show that relatively simple and theory-consistent on-chain indicators can still generate economically meaningful signals across full market cycles. Indeed, while the Sharpe ratios of our strategies do not match those of machine learning models (see, for example, Dubey and Enke 2025), they remain consistently higher than the buy-and-hold benchmark. Given the simplicity of the approach, we argue that less complex models are better suited for straightforward long-only implementation.

Furthermore, from a methodological point of view, recent research also shows that forecasting performance depends not only on model choice, but also on preprocessing, feature selection, labeling, and validation design. In that respect, the broader literature underscores that methodological credibility is as important as algorithmic sophistication. Relative to these studies, the main strength of the present paper lies in its transparency, interpretability, and economic coherence. Overall, our findings align well with the broader literature in showing that blockchain-based information is valuable for understanding Bitcoin price dynamics, while also demonstrating that a simpler indicator-based framework can recover part of the same predictive content identified by more complex machine-learning models.

5.2. Limitations

This study has limitations that also open avenues for future research. First, consistent with the existing literature (Burggraf et al., 2020; Kristoufek, 2013), we focus exclusively on Bitcoin. Although Bitcoin is the most liquid and prominent cryptocurrency, it may not represent the entire market, since other cryptocurrencies may behave differently due to distinct characteristics. Future work could examine the link between on-chain data and prices for other assets, such as Ethereum, Solana, and XRP. Second, this study focuses on three on-chain indicators selected on the basis of their prominence in the existing literature (Sakkas and Urquhart, 2024; Casella and Paletto, 2023) and their established interpretation as sentiment-driven signals. However, we acknowledge that the universe of on-chain metrics is considerably broader, and that other indicators could equally qualify as sentiment proxies and potentially generate competitive trading signals. Our choice was also guided by a preference for parsimony in model implementation, avoiding the data-mining risk that would arise from testing a large number of indicators and selecting only those that perform well ex post. Moreover, new on-chain indicators are frequently introduced, so the metrics studied here could become less informative as markets evolve and mature. Future research could therefore systematically evaluate a wider set of on-chain metrics, assessing their sentiment content and their ability to identify market cycle turning points. Third, while NUPL ratio, MVRV Z-score, and CVDD have shown an ability to identify market cycles, we acknowledge that the exit thresholds are not fully objective. In some cases, the thresholds proposed by metric creators may reflect hindsight bias, potentially reducing the indicators' credibility. However, there are no universally accepted rules for defining exit thresholds, so some discretion is unavoidable. To mitigate this concern, we performed sensitivity checks using three alternative exit thresholds, and the results remained robust, suggesting that performance does not rely on hindsight.

A further limitation concerns the role of Artificial Intelligence and Large Language Models (LLMs) in shaping investor sentiment, which this study does not account for. LLMs are increasingly used by retail and institutional investors to interpret market conditions and process information, with the potential to amplify behavioral biases. Lopez-Lira and Tang (2023) document that ChatGPT can predict stock market reactions from news headlines, while Kirtac and Germano (2024) show that GPT-based models significantly outperform traditional sentiment analysis methods in predicting stock returns. In the cryptocurrency context, Nguyen et al. (2025) demonstrate that ChatGPT-constructed sentiment indicators significantly explain Bitcoin returns, even after controlling for established sentiment proxies. To the extent that LLMs are influencing how investors form expectations and process public information, they may also modify the dynamics of the on-chain sentiment signals examined here. Future research could therefore investigate whether and how LLMs may affect predictability of on-chain indicators across Bitcoin market cycles.

6. Conclusion

This study is one of the first to investigate the performance of trading indicators based on on-chain data as measures of investor sentiment. Three indicators are considered: NUPL ratio and the MVRV Z-score, which provide signals for potential entries and exits from trading positions, and CVDD, which aims to identify market bottoms as long-entry opportunities. We use Bitcoin as the test asset over the period from December 7, 2013 to April 12, 2025, which spans three market cycles.

The results show that NUPL ratio and the MVRV Z-score outperform a buy-and-hold benchmark across all the metrics considered. This finding is robust to alternative exit strategies. CVDD also performs well relative to random entry rules, as assessed through Monte Carlo simulations. Overall, the evidence suggests that strategies based on on-chain data can capture investor sentiment and provide timing signals that help predict Bitcoin price cycles. These results are consistent with rejecting the efficient market hypothesis in the crypto context (Caporale et al., 2018; Urquhart, 2016), since such strategies should not systematically work otherwise, and they provide an additional theoretical contribution to this debate.

The findings are relevant to a broad set of market participants. Traders and retail investors may enhance their market analysis and potentially improve performance by identifying different market phases. In particular, these metrics appear to detect sentiment extremes that coincide with price peaks and bottoms. Traders and long-term investors can use these indicators to time major reversals and improve risk management, as traditional risk models often perform poorly in highly volatile markets. On-chain metrics, therefore, offer a complementary measure to monitor signals of overheating and widespread panic.

For regulators, the transparency of blockchains and the availability of on-chain data provide a granular, near real-time view of market sentiment and systematic behavioral shifts. On-chain metrics may serve as early-warning indicators of overheating markets, allowing regulators to act more proactively during periods of elevated systemic risk.

CRediT authorship contribution statement

Klaus Grobys: Writing – original draft, Supervision, Project administration, Investigation. **Sebastian Näsman:** Writing – original draft, Visualization, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Davide Sandretto:** Writing – review & editing, Writing – original draft, Validation, Project administration, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix

Table A1
Entry/exit timestamps for trading strategies

Strategy	Entry	Exit	Holding Days (invested)	Holding Days (0% cash)	Total Days
Buy-and-Hold Benchmark	07/12/2013	12/04/2025	4143	0	4143
NUPL					
NUPL1 Trade 1	05/10/2014	19/05/2017	958		
NUPL1 Trade 2	20/11/2018	26/12/2020	768		
NUPL1 Trade 3	14/06/2022	12/04/2025	1033		
Total NUPL 1			**2759**	**1384**	**4143**

(continued on next page)

Table A1 (*continued*)

Strategy	Entry	Exit	Holding Days (invested)	Holding Days (0% cash)	Total Days
NUPL2 Trade 1	05/10/2014	24/05/2017	964		
NUPL2 Trade 2	20/11/2018	02/01/2021	775		
NUPL2 Trade 3	14/06/2022	12/04/2025	1033		
Total NUPL 2			**2772**	**1371**	**4143**
NUPL3 Trade 1	05/10/2014	06/12/2017	1160		
NUPL3 Trade 2	20/11/2018	07/01/2021	780		
NUPL3 Trade 3	14/06/2022	12/04/2025	1033		
Total NUPL 3			**2973**	**1170**	**4143**
MVRV Z-Score					
MVRV Z-score 1 Trade 1	05/01/2015	24/05/2017	871		
MVRV Z-score 1 Trade 2	25/11/2018	02/01/2021	770		
MVRV Z-score 1 Trade 3	19/06/2022	12/04/2025	1028		
Total MVRV Z-Score 1			**2669**	**1474**	**4143**
MVRV Z-score 2 Trade 1	05/01/2015	27/11/2017	1058		
MVRV Z-score 2 Trade 2	25/11/2018	07/01/2021	775		
MVRV Z-score 2 Trade 3	19/06/2022	12/04/2025	1028		
Total MVRV Z-Score 2			**2861**	**1282**	**4143**
MVRV Z-score 3 Trade 1	05/01/2015	06/12/2017	1067		
MVRV Z-score 3 Trade 2	25/11/2018	19/02/2021	818		
MVRV Z-score 3 Trade 3	19/06/2022	12/04/2025	1028		
Total MVRV Z-Score 3			**2913**	**1230**	**4143**
CVDD					
CVDD 1	14/01/2015	17/12/2017	1069	-	1069
CVDD 2	14/12/2018	08/11/2021	1061	-	1061
CVDD 3	09/11/2022	12/04/2025	885	-	885

Data Availability

Data will be made available on request.

References

Aggarwal, D., 2019. Do bitcoins follow a random walk model? Res. Econ. 73, 15–22. https://doi.org/10.1016/j.rie.2019.01.002.

Al-Yahyaee, K.H., Mensi, W., Ko, H.-U., Yoon, S.-M., Kang, S.H., 2020. Why cryptocurrency markets are inefficient: The impact of liquidity and volatility. North Am. J. Econ. Financ. 52, 101168. https://doi.org/10.1016/j.najef.2020.101168.

Aslanidis, N., Bariviera, A.F., Savva, C.S., 2024. Do online attention and sentiment affect cryptocurrencies' correlations? Res. Int. Bus. Financ. 71, 102488. https://doi.org/10.1016/j.ribaf.2024.102488.

Baker, M., Wurgler, J., 2007. Investor sentiment in the stock market. J. Econ. Perspect. 21, 129–151. https://doi.org/10.1257/jep.21.2.129.

Balcilar, M., Bouri, E., Gupta, R., Roubaud, D., 2017. Can volume predict Bitcoin returns and volatility? a quantiles-based approach. Econ. Model. 64, 74–81. https://doi.org/10.1016/j.econmod.2017.03.019.

Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. J. Financ. Econ. 49, 307–343. https://doi.org/10.1016/S0304-405X(98)00027-0.

Bollen, J., Mao, H., Zeng, X., 2011. Twitter mood predicts the stock market. J. Comput. Sci. 2, 1–8. https://doi.org/10.1016/j.jocs.2010.12.007.

Brown, G.W., 1999. Volatility, sentiment, and noise traders. Financ. Anal. J. 55, 82–90.

Brown, G.W., Cliff, M.T., 2005. Investor Sentiment and Asset Valuation. J. Bus. 78, 405–440. https://doi.org/10.1086/427633.

Bu, Q., 2023. Are all the sentiment measures the same? J. Behav. Financ. 24, 161–170. https://doi.org/10.1080/15427560.2021.1949718.

Burggraf, T., Huynh, T.L.D., Rudolf, M., Wang, M., 2020. Do FEARS drive bitcoin? Rev. Behav. Financ. 13, 229–258. https://doi.org/10.1108/RBF-11-2019-0161.

Caporale, G.M., Gil-Alana, L., Plastun, A., 2018. Persistence in the cryptocurrency market. Res. Int. Bus. Financ. 46, 141–148. https://doi.org/10.1016/j.ribaf.2018.01.002.

Casella, B., Paletto, L., 2023. Predicting Cryptocurrencies Market Phases through On-Chain Data Long-Term Forecasting, in: 2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). Presented at the 2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), pp. 1–4. https://doi.org/10.1109/ICBC56567.2023.10174989.

Chen, H., De, P., Hu, Y. (Jeffrey), Hwang, B.-H., 2014. Wisdom of crowds: the value of stock opinions transmitted through social media. Rev. Financ Stud. 27, 1367–1403. https://doi.org/10.1093/rfs/hhu001.

Corbet, S., Eraslan, V., Lucey, B., Sensoy, A., 2019. The effectiveness of technical trading rules in cryptocurrency markets. Financ. Res. Lett. 31, 32–37. https://doi.org/10.1016/j.frl.2019.04.027.

Da, Z., Engelberg, J., Gao, P., 2011. Search of Attention. J. Financ. 66, 1461–1499.

Demeester, T., Blummer, T., Lescrauwaet, M.,2019. A Primer on Bitcoin Investor Sentiment and Changes in Saving Behavior. Medium. ⟨https://medium.com/⟩@adamant_capital/a-primer-on-bitcoin-investor-sentiment-and-changes-in-saving-behavior-a5fb70109d32.

De Bondt, W.F.M., Thaler, R., 1985. Does the stock market overreact? J. Financ. 40, 793–805. https://doi.org/10.2307/2327804.

Dubey, R., Enke, D., 2025. Bitcoin price direction prediction using on-chain data and feature selection. Mach. Learn. Appl. 20, 100674.

Dunbar, K., Owusu-Amoako, J., 2022. Cryptocurrency returns under empirical asset pricing. Int. Rev. Financ. Anal. 82, 102216. https://doi.org/10.1016/j.irfa.2022.102216.

Fama, E.F., 1970. Efficient capital markets: a review of theory and empirical work. J. Financ. 25, 383–417. https://doi.org/10.2307/2325486.

Grobys, K., Ahmed, S., Sapkota, N., 2020. Technical trading rules in the cryptocurrency market. Financ. Res. Lett. 32, 101396. https://doi.org/10.1016/j.frl.2019.101396.

Hudson, R., Urquhart, A., 2021. Technical trading and cryptocurrencies. Ann. Oper. Res 297, 191–220. https://doi.org/10.1007/s10479-019-03357-1.

Huynh, N., De Mello, L., Li, K., 2025. Evolution of investor sentiment: a systematic literature review and bibliometric analysis. Int. Rev. Econ. & Financ. 100, 104115. https://doi.org/10.1016/j.iref.2025.104115.

Jagannath, N., Barbulescu, T., Sallam, K.M., Elgendi, I., Mcgrath, B., Jamalipour, A., Abdel-Basset, M., Munasinghe, K., 2021. An On-Chain Analysis-Based Approach to Predict Ethereum Prices. IEEE Access 9, 167972–167989. https://doi.org/10.1109/ACCESS.2021.3135620.

Jia, B., Shen, D., Zhang, W., 2022. Extreme sentiment and herding: Evidence from the cryptocurrency market. Res. Int. Bus. Financ. 63, 101770. https://doi.org/10.1016/j.ribaf.2022.101770.

Kahneman, D., Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. Econometrica 47, 263–291. https://doi.org/10.2307/1914185.

Khosravi, M., Ghazani, M.M., 2023. Novel insights into the modeling financial time-series through machine learning methods: Evidence from the cryptocurrency market. Expert Syst. Appl. 234, 121012.

Khuntia, S., Pattanayak, J.K., 2018. Adaptive market hypothesis and evolving predictability of bitcoin. Econ. Lett. 167, 26–28. https://doi.org/10.1016/j.econlet.2018.03.005.

Kim, G., Shin, D.-H., Choi, J.G., Lim, S., 2022. A Deep Learning-Based Cryptocurrency Price Prediction Model That Uses On-Chain Data. IEEE Access 10, 56232–56248. https://doi.org/10.1109/ACCESS.2022.3177888.

King, J.C., Dale, R., Amigó, J.M., 2024. Blockchain metrics and indicators in cryptocurrency trading. Chaos Solitons & Fractals 178, 114305. https://doi.org/10.1016/j.chaos.2023.114305.

Kirtac, K., Germano, G., 2024. Sentiment trading with large language models. Financ. Res. Lett. 62, 105227. https://doi.org/10.1016/j.frl.2024.105227.

Kristoufek, L., 2013. BitCoin meets Google Trends and Wikipedia: Quantifying the relationship between phenomena of the Internet era. Sci. Rep. 3, 3415. https://doi.org/10.1038/srep03415.

Lahmiri, S., Bekiros, S., 2019. Cryptocurrency forecasting with deep learning chaotic neural networks. Chaos Solitons & Fractals 118, 35–40. https://doi.org/10.1016/j.chaos.2018.11.014.

Liu, Y., Tsyvinski, A., Wu, X., 2022. Common Risk Factors in Cryptocurrency. J. Financ. 77, 1133–1177. https://doi.org/10.1111/jofi.13119.

Long, S., Xie, Y., Zhou, Z., Lucey, B., Urquhart, A., 2025. From whales to waves: Social media sentiment, volatility, and whales in cryptocurrency markets. Br. Account. Rev., 101682 https://doi.org/10.1016/j.bar.2025.101682.

Lopez-Lira, A., Tang, Y., 2023. Can ChatGPT forecast stock price movements? Return predictability and large language models. J. Financ. Econ. Forthcom.

Mahmudov, M., Puell, D., 2018. Bitcoin Market-Value-to-Realized-Value (MVRV) Ratio. Medium. ⟨https://medium.com/⟩⟨@kenoshaking/bitcoin-market-value-to-realized-value-mvrv-ratio-3ebc914dbaee⟩.

Naeem, M.A., Mbarki, I., Shahzad, S.J.H., 2021. Predictive role of online investor sentiment for cryptocurrency market: Evidence from happiness and fears. Int. Rev. Econ. & Financ.

Nguyen, B.T., Chu, T.T., Ha, S.X., Nguyen, A.T., 2025. Decoding market sentiment: the power of ChatGPT in explaining Bitcoin returns from X data. China Financ. Rev. Int. https://doi.org/10.1108/CFRI-05-2024-0278.

Nimalendran, M., Pathak, P., Petryk, M., Qiu, L., 2025. Informational Efficiency of Cryptocurrency Markets. J. Financ. Quant. Anal. 60, 1427–1456. https://doi.org/10.1017/S0022109024000310.

Omole, O., Enke, D., 2024. Deep learning for Bitcoin price direction prediction: models and trading strategies empirically compared. Financ. Innov. 10 (1), 117.

Omole, O., Enke, D., 2025. Using machine and deep learning models, on-chain data, and technical analysis for predicting bitcoin price direction and magnitude. Eng. Appl. Artif. Intell. 154, 111086.

Opdyke, J.D., 2007. Comparing Sharpe ratios: So where are the p-values? J. Asset Manag 8, 308–336. https://doi.org/10.1057/palgrave.jam.2250084.

Pan, W.-F., 2020. Does investor sentiment drive stock market bubbles? beware of excessive optimism! J. Behav. Financ. 21, 27–41. https://doi.org/10.1080/15427560.2019.1587764.

Sakkas, A., Urquhart, A., 2024. Blockchain factors. J. Int. Financ. Mark. Inst. Money 94, 102012. https://doi.org/10.1016/j.intfin.2024.102012.

Sensoy, A., 2019. The inefficiency of Bitcoin revisited: a high-frequency analysis with alternative currencies. Financ. Res. Lett. 28, 68–73. https://doi.org/10.1016/j.frl.2018.04.002.

Shen, D., Urquhart, A., Wang, P., 2020. A three-factor pricing model for cryptocurrencies. Financ. Res. Lett. 34, 101248. https://doi.org/10.1016/j.frl.2019.07.021.

Tiwari, A.K., Jana, R.K., Das, D., Roubaud, D., 2018. Informational efficiency of Bitcoin—An extension. Econ. Lett. 163, 106–109. https://doi.org/10.1016/j.econlet.2017.12.006.

Ung, S.N., Gebka, B., Anderson, R.D.J., 2024. An enhanced investor sentiment index. Eur. J. Financ. 30, 827–864. https://doi.org/10.1080/1351847X.2023.2247440.

Urquhart, A., 2016. The inefficiency of Bitcoin. Econ. Lett. 148, 80–82. https://doi.org/10.1016/j.econlet.2016.09.019.

Verheyden, T., De Moor, L., Van den Bossche, F., 2015. Towards a new framework on efficient markets. Res. Int. Bus. Financ. 34, 294–308. https://doi.org/10.1016/j.ribaf.2015.02.007.

Vidal-Tomás, D., Ibáñez, A.M., Farinós, J.E., 2019. Weak efficiency of the cryptocurrency market: a market portfolio approach. Appl. Econ. Lett. 26, 1627–1633. https://doi.org/10.1080/13504851.2019.1591583.

Von Neumann, J., Morgenstern, O., 1944. Theory of games and economic behavior. Princeton University Press.

Woo, W., Puell, D., 2019. Experiments on Cumulative Destruction. Medium. ⟨https://medium.com/⟩@kenoshaking/experiments-on-cumulative-destruction-3a126fbd63b3.