Sebastian Näsman

Using On-Chain Data to Predict Cryptocurrency Cycles

School of Accounting and Finance
Master's thesis in Finance
Master's Degree Program in Finance

Vaasa 2025

UNIVERSITY OF VAASA	
School of Accounting and Finance	
Author:	Sebastian Näsman
Title of the Thesis:	Using On-Chain Data to Predict Cryptocurrency Cycles
Degree:	Master of Science in Economics and Business Administration
Programme:	Master's Degree Program in Finance
Supervisor:	Klaus Grobys
Year:	2025 **Pages:** 76

ABSTRACT:

This thesis investigates whether on-chain data—transparent, publicly available data recorded directly on blockchains—can be used to measure investor sentiment and forecast cryptocurrency market cycles. The research focuses on three core on-chain indicators: Net Unrealized Profit/Loss (NUPL), MVRV Z-score, and Cumulative Value Days Destroyed (CVDD). Through manual analysis of Bitcoin's historical price data and a Monte Carlo simulation, the study evaluates the performance of each indicator as a contrarian signal across three major market cycles. The findings support two hypotheses: (1) that on-chain metrics capture investor sentiment, and (2) exhibit a statistically significant relationship with price movements. The results indicate that NUPL ratio and MVRV Z-score can generate long signals that significantly outperform a buy-and-hold benchmark, with cumulative returns exceeding 326,000% in the strongest cases and annualized returns up to 175,6 %. CVDD entry strategies, tested using simulations with random control entries within ±50 to ±100-day ranges around the cycle bottom, consistently showed superior returns relative to chance. The study contributes to the literature by demonstrating that blockchain-based behavioral data offers predictive value in decentralized markets lacking intrinsic valuation anchors. It also complements recent advances in deep learning models trained on on-chain data, highlighting rule-based metrics' blend of interpretability and statistical strength. The results carry practical implications for investors, traders, and regulators, and challenge traditional notions of market efficiency by providing evidence of repeated behavioral patterns embedded in public blockchain activity.

KEYWORDS: On-chain data, Cryptocurrency market, Investor sentiment, Bitcoin, Market cycles, Behavioral finance, Machine learning, Sentiment analysis, Predictive analytics, Market efficiency

Contents

Figures

Tables

1 Introduction

The financial world rarely witnesses innovations as disruptive and polarizing as cryptocurrencies. The first cryptocurrency, Bitcoin, was born from the ashes of the 2008 financial crisis as an attempt to replace fiat money and the old broken system. Over the years, the industry has gone through many boom-and-bust cycles and eventually grown to over a trillion-dollar valuation. As this ecosystem grows and evolves, so does the need for more accurate and robust tools for analyzing this unique market and its participants. Although the cryptocurrency market is diverse, this thesis mainly focuses on Bitcoin due to its dominant market position and data availability. Still, the insights gained from the leading asset can be relevant across the broader market due to the high correlation between Bitcoin and other cryptocurrencies.

Traditional financial models have difficulty explaining the price movements of crypto assets. Unlike stocks, most cryptocurrencies lack intrinsic value, leading researchers to find new ways to reflect the market dynamics. One of the most promising new data sources is on-chain data, which is directly recorded on the blockchain and reflects the actions and sentiments of network participants in real time. Unlike in traditional markets, where sentiment is gathered through different sentiment proxies like surveys or media sentiment, blockchain offers a direct, transparent, and tamper-proof measure of sentiment free of biases. Every action is logged on the blockchain permanently, offering researchers a verifiable trail of investor behavior, enabling them not to measure what people say but what they do.

As cryptocurrencies are largely sentiment-driven, 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 dynamic opens the door for sentiment analysis and the use of on-chain data as the purest measure of that. Also, the sentiment-driven nature and the inefficiencies in the cryptocurrency market make predictive analysis more beneficial than in traditional markets. As explained in this thesis, crypto markets violate the

Efficient Market Hypothesis (EMH) assumptions, and on-chain metrics, especially when fed to machine learning models, can use those inefficiencies to predict market reversals successfully.

These ideas guide this thesis forward in combination with financial theory, behavioral insights, and empirical evidence. Through historical case studies and recent academic research, this thesis tries to measure whether those behavioral traces left on the blockchain can be used to measure sentiment and predict market cycles. The goal is not only to see whether specific metrics prove to be valuable in price forecasting but also to improve our general understanding of sentiment and market dynamics. Insights gained from digital assets may ultimately reshape our understanding of market psychology itself, transforming not only sentiment analysis but also the foundational theories that are the core of modern financial markets.

1.1 Purpose of the Study

This thesis uses a thorough literature review and empirical analysis of historical and present data to explore the relationship between on-chain data and cryptocurrency price trends, with an emphasis on the leading digital asset, Bitcoin. The main objective is to find whether on-chain metrics can reliably predict cryptocurrency price movements. By focusing on historical and real-time data recorded on the blockchain, this study aims to provide insights for investors and traders.

The study identifies and analyzes critical on-chain metrics connected to the price of several cryptocurrencies. These include Net Unrealized Profit/Loss (NUPL), MVRV Z-score, and Cumulative Value Days Destroyed (CVDD), which are examined through historical case studies. In addition, metrics such as active addresses, miner revenue, exchange inflows/outflows, and Spent Output Profit Ratio (SOPR) are explored in machine learning research.

This thesis combines blockchain's technical data with financial market analysis. The study's findings can help investors make more informed investment decisions. Additionally, the study contributes to academic research on on-chain data and cryptocurrency price prediction.

The research will test two main hypotheses:

H1: On-chain metrics can be used to measure the investor sentiment of the cryptocurrency market.

The first hypothesis is that we can gauge the overall market sentiment through the activity of blockchain participants. On-chain data consists of people's actions on the blockchain and, thus, also their emotions. Every transaction on the market is also linked to a particular emotion, whether it is greed, fear, or something else. For this reason, it is possible to measure investor sentiment using on-chain data.

H2: A statistically significant relationship exists between on-chain metrics and crypto-currency prices.

The second hypothesis is linked to the first: if you can measure investor sentiment through on-chain metrics, you should also find a correlation with price movements. Sentiment works as a contrarian indicator. Market tops occur when people become excessively greedy, while market bottoms form when people become overly fearful.

These two hypotheses support each other. Previous literature has found a link between investor sentiment and returns (Lee et al. 1991; Wang et al. 2021 and Jiangshan et al. 2021). If on-chain metrics can be used to measure investor sentiment, they should have a relationship with cryptocurrency prices. This suggests that if we can measure investor sentiment with on-chain metrics, they should also help predict cryptocurrency prices more accurately.

1.2 Structure of the Study

The thesis is structured into five main sections. The first section introduces the research topic, outlines the purpose of the study, and presents the structure of the thesis. The second section explains cryptocurrencies, traditional finance theories, investor sentiment, and on-chain data. The third section forms the core of the thesis, presenting three case studies of on-chain metrics—NUPL ratio, MVRV Z-score, and CVDD—and reviewing recent machine learning applications that use on-chain data to predict cryptocurrency prices. The fourth section discusses the practical and academic implications of this thesis, what limitations there might be, and recommendations for future research. Lastly, the conclusion presents the final thoughts and binds everything together.

2 Preliminaries

2.1 Cryptocurrencies

2.1.1 Historical Background and Development

Cryptocurrencies are digital assets enabling secure peer-to-peer transactions using decentralized ledger technology. Unlike traditional currencies controlled by centralized entities, cryptocurrencies work in distributed networks that use consensus mechanisms to validate transactions (Bonneau et al., 2015; Tripathi et al., 2023). Bitcoin, the most well-known cryptocurrency, was first introduced by an anonymous entity called Satoshi Nakamoto in 2008. However, the market has expanded significantly since then, and the number of alternative cryptocurrencies (altcoins) has risen to over 11 million (coinmarketcap.com).

The first concept of digital money dates far before Bitcoin and has roots in cryptographic research dating back to the 1980s. Bonneau et al. (2015) explain in their study that one of the first proposals was made by David Chaum, who introduced the e-cash system with the idea of cryptographically secured, untraceable digital payments. In the following decades, many other proposals, such as Wei Dai's B-money and Nick Szabo's Bit Gold, were also made, which built the foundations for Bitcoin's design.

Bitcoin's innovative value lies in solving the double-spending problem without a centralized authority (Bonneau et al., 2015; Tripathi et al., 2023). Nakamoto achieved this through the combination of Proof of Work (PoW) with a publicly accessible ledger, which allowed network participants (miners) to compete in the validation of transactions and thus provide security and immutability to the system. This breakthrough enabled the fast growth of Bitcoin and the subsequent creation of many other digital currencies with new consensus algorithms, transaction processing speeds, additional privacy measures, and scalability (Bonneau et al., 2015).

2.1.2 Key Characteristics of Cryptocurrencies

Several critical characteristics of cryptocurrencies differentiate them from traditional financial systems. Cryptocurrencies are fundamentally decentralized. Traditional banking systems have central control of transactions. In contrast, cryptocurrencies rely entirely on distributed networks to verify and store transaction records (Bonneau et al., 2015). This decentralization also presents governance and scalability issues, besides being more secure and less prone to censorship.

Different consensus procedures exist, which cryptocurrency systems use to confirm transactions and defend against attacks on their networks. The most common mechanisms include Proof-of-Work (PoW): Bitcoin uses this system, where miners solve complex cryptographic challenges to append new blocks to the blockchain. Tripathi et al. (2023) note that some advantages of PoW are its secure nature and ability to reduce chances for Denial-of-service attacks. However, PoW requires significant computational power, leading to high energy use, according to Bonneau et al. (2015).

Proof-of-Stake (PoS): PoS presents a new blockchain consensus method that chooses validators from those who hold more cryptocurrency than non-validators. Proof-of-Stake (PoS) has two main benefits: increasing energy efficiency and being highly scalable (Tripathi et al., 2023). The cryptocurrency sector also analyzes new consensus systems, including delegated proof-of-stake (DPoS), proof-of-authority (PoA), proof-of-burn (PoB), and other systems that deliver distinct levels of security, efficiency, and decentralization.

Transactions in a blockchain are generally considered transparent and immutable. This means that confirmed transactions on a blockchain cannot be modified. Also, the same transaction history is accessible to all network members, minimizing fraud and manipulation risks, as Bonneau et al. (2015) noted. However, this has not always been true. For example, following the 2016 DAO hack and the loss of over 50 million worth of Ether, the

Ethereum community decided to implement a hard fork reversing transactions associated with the hack (Cryptopedia, 2025). This shows that blockchain immutability can be altered in exceptional situations. Still, these features remain essential, especially when discussing the potential of Central Bank Digital Currencies (CBDCs), where transparency and security are key components (Alvarez et al., 2023).

Bitcoin transactions operate under pseudonyms but maintain no complete privacy. Blockchain analysis tools enable tracking transactions between addresses and linking them to real-world identities. To enhance transaction anonymity, several cryptocurrencies, including Monero and Zcash, employ ring signatures and zero-knowledge proofs as advanced cryptographic techniques (Bonneau et al., 2015).

2.1.3 Challenges and Future Directions

The adoption and stability of cryptocurrencies remain at risk due to multiple challenges that affect them. As a result of its original design, Bitcoin processes roughly seven transactions per second. Network usage expansion produces congestion alongside elevated transaction fees, which become significant problems. The Lightning Network, together with sharding, exists to enhance scalability while preserving the decentralized nature of the network, according to Bonneau et al. (2015).

Blockchain technology is generally secure; however, Bonneau et al. (2015) and Tripathi et al. (2023) point out that several threats still exist. These include 51% attacks, selfish mining, and Sybil attacks. 51% attacks mean that a single entity gains the majority control of the mining power, giving them the ability to manipulate transactions. Selfish mining is a mining strategy where the miner withholds newly discovered blocks to gain a competitive advantage. Lastly, Sybil attacks are when attackers create fake identities (nodes) to influence and manipulate the network. The continuous development of security measures needs to be maintained in order to protect cryptocurrency network integrity (Bonneau et al., 2015).

Cryptocurrencies function within a rapidly changing regulatory environment. Governments and financial institutions strive to determine the correct classification and implementation guidelines for digital assets. Taxation issues, anti-money laundering (AML) compliance, and consumer protection concerns continue to be vital priorities. For example, Alvarez et al. (2023) find that in El Salvador, where Bitcoin has been adopted as legal tender, these regulatory uncertainties still hinder adoption and business acceptance.

In addition to the already mentioned challenges, cryptocurrencies face scrutiny for their impact on climate change. Proof-of-work-based cryptocurrencies need significant computational power, which has sparked environmental usage concerns. Current efforts support the movement toward energy-efficient consensus mechanisms, including PoS and hybrid systems (Bonneau et al., 2015). Even with the efforts to move toward more environmentally friendly consensus mechanisms, environmental concerns still significantly affect the adoption of digital currencies.

2.2 Efficient Market Hypothesis (EMH) vs. Market Inefficiencies

The Efficient Market Hypothesis (EMH) is one of the most popular finance theories, first formalized by Eugene Fama (1970). The theory suggests that the financial markets are "informationally efficient". This means that all available information is already reflected in the asset prices. This leads to the realization that achieving abnormal returns consistently with fundamental or technical analysis is impossible. However, this theory has long been criticized and challenged by numerous empirical findings demonstrating that investors do not act in practice like the EMH would suggest, leading to market inefficiencies.

2.2.1 Levels of Market Efficiency

To explain how information is reflected in prices, Fama created a three-tier classification system: weak-form efficiency, semi-strong efficiency, and strong efficiency. These different tiers represent how information is incorporated into asset prices.

The most basic level of efficiency is the weak-form. In this tier, current prices reflect all past price and volume information, meaning that trading strategies based on analyzing past price patterns or volume data cannot produce abnormal returns in the long term. The rationale is that if there were any exploitable pattern or anomaly, investors would have already exploited it, erasing any potential profit opportunity.

Different empirical studies have tested this hypothesis. One notable paper is Kendall and Hill (1953), which later influenced Eugene Fama's development of EMH. Kendall and Hill (1953), who first introduced the famous Random Walk Theory, argue that stock movements are random and follow the so-called random walk. Some later studies have since found deviations from that. Lo and McKinlay (1988) find a short-term autocorrelation in stock prices, indicating that, in practice, stocks might not follow a perfect random walk. However, these deviations are insufficient to dispute the idea of weak-form efficiency.

It is still an ongoing debate whether stocks fall under the weak-form efficiency; however, with cryptocurrencies, there is no honest debate about that. Many studies have demonstrated that past price trends in cryptocurrencies may persist long, indicating a significant momentum effect. For example, Balcilar et al. (2017) find a short-term autocorrelation in cryptocurrencies, allowing traders to profit from momentum-based strategies. Also, Lahmiri and Bekiros (2019) find that past Bitcoin prices influence future prices, strongly contradicting the weak-form efficiency.

The semi-strong efficiency goes beyond past prices, indicating that all publicly available information is also reflected in asset prices. This means that neither technical nor fundamental analysis can consistently create abnormal returns. In fundamental analysis, investors analyze company financials, industry trends, and economic indicators to find undervalued companies. As all of these are publicly available information, under semi-strong efficiency, they will not provide any edge to the investor as they are already fully incorporated into asset prices.

In traditional markets, there is evidence for and against the semi-strong efficiency. For example, Fama et al. (1969) find that stock prices immediately react to earnings announcements, indicating at least semi-strong efficiency. Other studies, however, also find conflicting evidence. The study by Bernard and Thomas (1989) illustrated a phenomenon called post-earnings announcement drift (PEAD), where prices drift slowly over weeks in the same direction following an earnings announcement. This can be seen as evidence of some inefficiency.

Regarding cryptocurrencies, it is much more straightforward to say that they do not follow semi-strong efficiency. Their unregulated nature already makes sharing information very different from traditional markets. Unlike in regulated markets, information sharing does not follow strict reporting standards and instead relies on social media, news websites, and influencers. This makes the information less trustworthy and thus can affect investors' reaction time. One notable example of this is the Bitconnect scandal in 2018. Bitconnect was a Ponzi scheme that rose to over 2 billion market capitalization before falling to zero (Fortune, 2018). Many financial experts tried to warn investors well before the collapse, and it was a well-known fact that the coin was suspected to be a Ponzi scheme. Still, many decided to buy into the coin despite this information, indicating that not all investors find or trust the information available to them. Naem et al. (2021) also find that cryptocurrency prices are strongly influenced by social media sentiment, which causes investors to act irrationally instead of following objective financial data and metrics.

Strong-form efficiency is the highest level of market efficiency. A market is considered to have a strong-form efficiency when all public and private information is reflected in asset prices. This means insider trading is impossible, as even private information is instantaneously reflected in the price.

Unlike with the weak-form and semi-strong form efficiency, where researchers disagree on whether our market falls into those categories, with strong-form efficiency, there is really no debate. For example, Jaffe (1974) and Seyhun (1986) find that corporate insiders can consistently earn abnormal returns. This means that private information gives insiders a significant trading advantage and shows that the private information is not reflected in the asset prices. Many regulatory agencies, like the U.S. Securities and Exchange Commission (SEC), try to stop insider trading with different insider trading laws, which shows that the markets cannot be fully efficient; otherwise, these laws would not be needed.

The situation is even worse in the cryptocurrency markets, as there is no clear regulation or oversight of this issue. Many crypto projects are controlled by a few wallets ("whales") that can manipulate the price and share insider information, taking full advantage of the information asymmetry in the cryptocurrency markets. In 2024 and 2025, several celebrity meme coin projects were clear examples of insider trading.

One example is the $HAWK meme coin launched in December 2024 by influencer Haliey Welch, also known as the "Hawk Tuah" girl. According to Forbes (2024), this meme coin quickly rose to nearly 500 million in market capitalization before crashing by 95%, all in under a day. Investigations revealed that insiders controlled 80% to 90% of the coin's supply and then used that to facilitate a classic pump-and-dump scheme. Welch has since been accused of insider trading and market manipulation.

Another example is the meme coin $LIBRA, launched in February of 2025 and promoted by the Argentine president, Javier Milei. The token briefly exceeded a four-billion-dollar valuation before crashing by 90% within hours (Axios, 2025). The insiders also heavily controlled this project, owning 70% of the coin's supply. Many have accused the project of a "rug pull" scam, and Argentina's opposition has even threatened to start an impeachment trial against the president, Javier Milei (Reuters, 2025).

These cases highlight the information asymmetry in cryptocurrency markets and the lack of regulation and oversight. Strong-form efficiency does not exist in the cryptocurrency market, as insiders constantly gain abnormal returns by abusing private information. Many studies support this view, such as the study by King et al. (2024), which shows how large wallet holders can buy or sell assets in anticipation of price movements. This and all the other evidence indicate that cryptocurrency markets fall 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. Empirical research points out that historical price data has predictive power (violating weak-form efficiency), public information is not fully reflected in the 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 researched through on-chain analysis in this study.

2.2.2 Expected Utility Theory (EUT) and Prospect Theory: Why Rational Decision-Making Fails, Leading to Inefficiencies

The Efficient Market Hypothesis (EMH) assumes that all information is reflected in asset prices. This assumption requires that investors act rationally and always make the most optimal decisions based on the available data. This assumption aligns with Expected Utility Theory (EUT), which believes investors try to maximize their expected utility when making decisions.

Von Neumann and Morgenstern first formalized EUT in their groundbreaking book "Theory of Games and Economic Behavior" (1944). It is a normative theory, meaning it describes how people should act if they were perfectly rational. EUT assumes that people assign probabilities to different outcomes and then choose the option that maximizes their expected utility. This means that investors do not try to maximize monetary gains but instead maximize their personal utility. This is because investors have different risk appetites. For example, a risk-averse person might opt for a lower return with greater certainty, unlike a risk-seeking investor willing to invest in high-risk, high-reward investments. This theory still has many assumptions, such as that individuals have stable risk preferences, process information efficiently, and are not influenced by cognitive biases.

However, empirical evidence has shown that investors do not always act rationally, which can create predictable market inefficiencies. These findings show the limitations of EUT and explain why it cannot be a descriptive theory that tries to explain how investors act in practice. Prospective Theory, developed by Kahneman and Tversky in 1979, provides an alternative model to capture better the behavioral biases and irrationality that drive investment decisions.

As said, EUT assumes that investors can always assess probabilities correctly and constantly make rational decisions. However, it has been demonstrated that financial markets are driven by emotion, biases, and heuristics rather than pure rationality. For example, De Bondt and Thaler (1985) show that investors tend to overreact to news, which leads to momentum effects and sharp price reversals, contradicting EUT's predictions. Barberis and Huang (2001) find that investors put too much weight on recent experiences, leading to wrong market predictions, further indicating the lack of rationality in investor decision-making.

One of the most significant contradictions to EUT is the disposition effect. The disposition effect is when investors hold onto losing positions much longer than they do with winning positions. If the assumptions in EUT are accurate, investors should evaluate each trade independently based on the expected risk and return. However, Shefrin and Statman (1985) show that investors are more sensitive to losses than gains of the same size, indicating emotion-based decision-making. This behavior contradicts EUT entirely and points to Prospect Theory as a potentially more accurate model of actual investor behavior.

Prospect Theory introduces a new concept called loss aversion, which directly challenges the assumptions of EUT. Loss aversion refers to investors' inclination to feel the pain of losses more strongly than the joy of equivalent gains. Because of loss aversion, investors take excessive risks to avoid realizing a loss. Unlike in EUT, where investors evaluate outcomes only based on the final wealth, in Prospect Theory, investors evaluate their situation based on a reference point. These reference points are usually derived from the investor's purchase price, prior expectations, or psychological level, such as the all-time high or round number.

One of Prospect Theory's advantages is its ability to explain speculative bubbles and the momentum effect. Extreme optimism leads to excessive risk-taking, and people run after gains because they fear missing out (FOMO). This phenomenon has been visible in many instances during the history of financial markets. The first example could be the Dutch Tulip Mania (1634-1637) and, more recently, the Dot-com bubble of the late 90s and the 2017 and 2021 bull markets in cryptocurrencies. In all these examples, excessive optimism and risk-taking led to unsustainable price increases that theories like EMH or EUT cannot explain.

The volatile nature of cryptocurrencies is a clear example of Prospect Theory in action. Studies such as Naeem et al. (2021) show how cryptocurrency prices are strongly affected by sentiment and emotion. Without concrete fundamentals in cryptocurrencies, it is even easier to notice how different emotions move prices up and down. For example, the 2022 collapse of Terra Luna (LUNA) and the stablecoin UST concretely showed this dynamic, as the initial losses caused a massive panic that eventually led to a significant drop in the whole market and 60 billion getting wiped out of the market, according to Forbes (2022).

Another concept introduced by Prospect Theory is the framing effect, which means that individuals are influenced by how the choice is presented, not just its intrinsic value. For example, they might feel differently about a company with a 90% survival rate than a company with a 10% failure rate, even though they are the same. This effect is becoming even more pronounced daily as our world is becoming increasingly more of an attention-driven, fast-paced environment where both traditional and social media compete for your time and attention. This information, presented to you in news titles and TikTok videos, impacts your thinking and judgment and, thus, your market behavior.

2.2.3 Noise Trader Risk and Market Inefficiencies: Why Irrational Traders Drive Price Movements

The assumptions of market efficiency are challenged by the presence of irrational investors, also called noise traders. Fisher Black (1986) first introduced noise trading, which refers to investors who make decisions based on sentiment, speculation, misinformation, and emotions rather than fundamental analysis. Noise traders push prices away from their fundamental values and do it unpredictably. De Long et al. (1990) argue that noise traders create inefficiencies that are hard to arbitrage away by rational traders. These inefficiencies then lead to market bubbles and crashes. As noise traders are more likely to use emotion over logic, their impact is more substantial on markets where asset valuation is not straightforward.

One key assumption of EMH is that rational traders will arbitrage away any mispricing. However, this is not as easy as it sounds. Usually, arbitrage opportunities are limited, leading to extended periods of mispricing. Several market frictions allow noise trader risk to persist.

First, fundamental risk makes it hard for rational investors to correct mispricing. Even when the asset is clearly mispriced, arbitrageurs must face the risk of losing money if they wish to profit from it. Noise traders can push prices further, which could liquidate the arbitrager before the market corrects toward a fair price. An example of this is the Dot-com bubble, where it was evident for a long time that the prices were highly elevated and that there was excessive risk-taking, but still, it did not cause an imminent correction or a crash. Many “rational” traders tried to short the market but eventually got liquidated as the market kept climbing higher. They were right about the mispricing but not about the timing. Eventually, the market did come crashing down, but for many of those “rational” traders, the crash came way too late.

The second market friction is timing difficulty due to the unpredictable nature of noise traders. De Long et al. (1990) explain that noise traders' irrationality means they also move the markets in irrational ways. This leads to the difficulty of timing, as mispricing lasts much longer than many would assume, sometimes even for years.

A third limitation is short-selling constraints, which limit arbitrageurs' ability to profit from mispricing. In short selling, you borrow an asset, then sell it in the market, repurchase it later, hopefully at a lower price, and give the share back to the original owner. The profit you make is the difference between your buy and sell price minus the borrowing costs. In theory, rational investors should be able to use short selling to correct the mispricing. Still, regulatory restrictions, borrowing costs, and institutional constraints limit the investor’s ability to execute short trades effectively.

Regulators might sometimes impose short-selling bans, especially during market downturns. These bans are meant to protect investors, but Brunnermeier and Pedersen (2005) argue that they might allow mispricing to persist for longer. These bans create inefficiencies in the markets as rational investors cannot correct mispricing even if they want to. These short-selling constraints can be even more pronounced in retail-driven markets like cryptocurrencies. Short selling is often only available through leverage trading, which leads to high fees, liquidation risks, and limited liquidity.

2.3 Investor Sentiment

Investor sentiment refers to investors' attitudes toward a particular asset or market. It is not so much about how something is but how it feels. Investor sentiment comprises emotions like optimism, pessimism, fear, and greed, influencing individuals' trading and investing decisions. Unlike a company’s earnings or revenue, which are easily measurable, investor sentiment is a psychological and behavioral factor, making it much harder to measure (Bollen et al., 2011; Da et al., 2015).

Investor sentiment impacts financial markets significantly as it can push prices far beyond their fundamental value. For example, over-optimism can lead to excessive risk-taking and price surges, sometimes even leading to economic bubbles that eventually crash significantly below their actual value as people start panicking. The fundamental factors might drive financial prices in the longer term, but sentiment creates those fluctuations above and below the mean. Alvarez et al. (2023) highlight that although Bitcoin gained legal tender status in El Salvador, adoption rates remained low as the investor sentiment was extremely negative. This is just one example of how sentiment can affect the real-world adoption of financial innovations.

2.3.1 Measuring Investor Sentiment

Measurement of investor sentiment is challenging because of its complex nature. As mentioned, something that is not tangible is harder to measure. Many external factors can influence people’s decisions, which can be hard to isolate. Also, investors might not disclose their feelings truthfully because of shame and behavioral biases (Baker and Wurgler, 2006; Brown and Cliff, 2005). Despite these challenges, many different measurement methods have emerged over the years.

One of the most recognized ways to measure investor sentiment is through closed-end funds. Lee et al. introduce this idea in their 1991 paper “Investor Sentiment and the Closed-end Fund Puzzle”. You value a closed-end fund by calculating its net asset value (NAV), which is done by subtracting the fund’s liabilities from its assets and dividing it by the total outstanding shares. This means that the closed-end fund’s value should equal the combined value of all its stocks. The trading price should be the same in efficient markets as the fund's NAV. This is rarely the case, and instead, closed-end funds typically trade at a discount or a premium. Lee et al. (1991) explain that the difference between the trading price and the NAV is the measure of sentiment. Significant discounts indicate a bearish sentiment; in the case of premiums, investors feel overly optimistic.

Other more recent ways to measure sentiment are, for example, search engine data, sentiment surveys, and social media analysis. These have grown more popular over the years as technology has evolved. Da et al. (2015) use search engine data to measure sentiment and find that abnormal increases in search volumes predict short-term outperformance. Alvarez et al. (2023) also find that increases in search activity coincide with significant price movements in Bitcoin. This is not surprising, as increased media coverage of Bitcoin, such as reports on local news channels, often leads individuals to search for information like "how to buy Bitcoin." This behavioral response illustrates why search engine query data can serve as a proxy for investor sentiment.

One of the more traditional ways to measure sentiment is through surveys. The most well-known sentiment surveys are the AAII Sentiment Survey, Investor Intelligence, and UBS/Gallup polls. Brown and Cliff (2005) use the Investor Intelligence survey to measure sentiment and find that over-optimism leads to lower future returns. This illustrates sentiment's contrarian nature as, typically, extremes in sentiment lead to major turning points in the market. Baker and Wurgler (2006) note that surveys have limitations, such as response biases and delayed data collection. However, they remain valuable tools because of their long historical continuity and ability to capture broad market moods.

The rise of the internet has changed many things, and measuring sentiment is also one of them. Social media has become a key source of sentiment data and is used heavily to measure sentiment nowadays. For example, Bollen et al. (2011) use Twitter/X feeds to analyze sentiment, and they find a significant correlation between different moods and the Dow Jones Industrial Average (DJIA) performance. Their model can predict DJIA trends using social media sentiment with 87,6 % accuracy.

As technology has evolved, so has sentiment measurement. From traditional methods like closed-end fund discounts to social media sentiment, sentiment measurement has grown to include many different methods with unique advantages. All these methods still prove that it is possible to measure sentiment and use that to predict prices, which should not be possible from the assumptions of the EMH. The next section examines how sentiment is typically measured in cryptocurrency markets and how it can be a powerful tool for price prediction, just like in traditional markets.

2.3.2 Cryptocurrency Market Sentiment

Measuring investor sentiment in the cryptocurrency market can be more beneficial for investors than in traditional markets. Unlike stocks, cryptocurrencies do not give investors the right to a company's cash flows. For this same reason, conventional valuation methods cannot compute the fair price of any given cryptocurrency. Sentiment, however, solves this problem.

Any price above zero is a sign of market inefficiency and a construct of sentiment. For this reason, sentiment is the optimal measurement when evaluating a fair price for a cryptocurrency. For the same reason, we have no joint hypothesis problem when we measure investor sentiment in the cryptocurrency market.

Investor sentiment in cryptocurrencies has been measured using different methods and data sources. The few most prominent have been search engine data and social media activity, highlighted by studies from Kristoufek (2013), Burggraf et al. (2020), and Muhammad et al. (2021). Kristoufek analyzes Bitcoin search volume from Google and Wikipedia and finds that increased investor attention correlates positively with price, finding a bidirectional relationship.

Burggraf et al. (2020) develop a FEARS index from negative Google search terms, following the footsteps of Da et al. (2014). The FEARS index reflects investor anxiety towards micro- and macroeconomic concerns. The paper finds a significant negative relationship between the FEARS index and the price of Bitcoin. This is true, especially with regard to microeconomic concerns. In parallel, Muhammad et al. (2021) also use the FEARS index but expand the sentiment analysis to social media by tracking happiness from Twitter/X data. They find that the Twitter/X Happiness sentiment index strongly predicts cryptocurrency returns but only in extreme sentiment states.

These studies indicate the sentiment's ability to predict cryptocurrency price trends and prove it is an essential tool in a market where traditional valuation methods are useless. These researchers utilized measurement methods already widely used for stock market sentiment. However, blockchain technology offers a new type of data that is valuable in sentiment analysis: on-chain data.

Unlike in traditional markets, where sentiment needs to be measured indirectly through surveys, search engine data, or social media, on-chain data provides a direct and real-time way of measuring sentiment. Every transaction recorded in the ledger expresses an investor's decision, and every one of those decisions is influenced by emotions. This makes on-chain metrics sentiment-driven, enabling them to capture market participants' collective moods and expectations without biases. On-chain metrics do not just correlate with sentiment, but they are the sentiment itself. Before going into price prediction through on-chain analysis, it is wise to explain what on-chain data is. In the next section, this thesis explains why on-chain data is crucial for sentiment analysis and price prediction.

2.4 On-chain Data

On-chain data comprises all transactions and interactions permanently recorded on a blockchain. This data is transparent and publicly accessible, making blockchains inherently auditable and secure. Due to its unchangeable nature, once information is added to the blockchain, it cannot be altered, establishing a reliable and tamper-proof environment (Jagannath et al., 2021).

One of the most significant technological breakthroughs introduced by Nakamoto (2009) was the possibility of achieving consensus in a decentralized network without relying on trust. The proof-of-work protocol and a chain of timestamped and cryptographically secured blocks were used to accomplish this. Through this design, it was not only possible to prevent double-spending but also to create permanent, verifiable records that work as the foundation for on-chain analysis. Transactions are stored in a distributed public network, which creates a time-series dataset of investor behavior that is impossible to manipulate after it has been recorded.

On-chain data can be divided into transaction data, block details, smart contract data, and participation metrics. Transaction data includes all the transactions made between wallet addresses, timestamped, and stored in individual blocks on the chain. Block details consist of information about the blocks themselves. This information can include block size, block height, and the mining difficulty associated with each block. Smart contract data records any interactions with smart contracts, including when contracts are executed, altered, or terminated.

In addition to categorizing on-chain data into transaction data, block details, and so on, we can divide them into single metrics. These single metrics usually have specific roles in describing on-chain activity, such as how active users are, how secure the network is, and how much value is transferred. Since these single metrics are used later in this thesis, especially in the machine learning section, it is beneficial to explain them here first. It is essential to understand what different metrics measure and why they are important for a better foundation for later analysis. The following list explains some of the most used and critical metrics in the context of this thesis:

- **Active Addresses:** It measures the number of unique addresses making at least one transaction during a period. It tells how active the network is and how high the user participation is.

- **Transaction Count:** Transaction count measures the total number of confirmed transactions during a specific period. A higher number usually means more participation and speculation in the network.
- **Realized Value (Realized Capitalization):** Realized Value represents the total value of all coins based on their value when last moved.
- **Miner Revenue:** It is the sum of income miners receive from block rewards and transaction fees. If revenues are dropping, it might cause selling pressure as miners try to combat the loss of revenue by selling their coins.
- **Exchange Inflows/Outflows:** Tracks the amount of cryptocurrencies sent to and from exchanges. Inflows are usually seen as potential selling pressure, as many hold their coins in cold storage and move them to exchanges when they want to sell. Conversely, outflows usually signal accumulation, as investors have just bought coins and are confident enough to move them to cold storage, meaning they are not planning to sell them immediately.
- **Gas Fees (Ethereum):** Gas fees represent the cost of transacting and executing smart contracts. Gas fees rise when there is a lot of traffic, which could indicate a strong demand and a lot of speculation.
- **Hash Rate:** Hash rate measures the computational power needed for mining in proof-of-work blockchains. A higher hash rate means a more secure network.
- **Block Height:** This calculates the number of blocks added to the chain since the genesis block (the first block). It is used to segment historical data.
- **SOPR (Spent Output Profit Ratio):** It is used to measure whether coins in the network are sold at a profit (>1) or at a loss (<1). It tries to capture market sentiment without relying on price trends.

Casella and Paletto (2023), Kim et al. (2022), and King et al. (2024) all use some of these metrics in their on-chain analysis, showing their relevance in behavioral proxies. These metrics are sentiment and risk appetite proxies, which can be hard to capture from traditional data sets. This is why they are widely used in different cryptocurrency price prediction models, as sentiment and price are especially intertwined in cryptocurrency markets.

The wide variety of on-chain data has many different use cases, but market analysis is one of the more widespread. Jagannath et al. (2021) explain that analysts use on-chain data to predict prices and market trends using metrics such as transaction volume and active addresses, which provide information about the network's demand. On-chain data can also be used to assess the blockchain's security and network health. Metrics like hash rate and mining difficulty, as explained in the list, are essential for the health of blockchain networks. The last notable use case for on-chain data is compliance and monitoring. The transparency of blockchain allows regulators and law enforcement agencies to monitor transactions effectively. This also helps ensure compliance with financial regulations and detect illegal activities.

In addition to these advantages, on-chain data offers valuable information for sentiment-based financial modeling. Unlike in traditional markets, where transaction-level data is usually proprietary or aggregated, blockchain is accessible in real-time and at an address level. This enables traders and researchers to build models based on actual transactional behavior instead of relying on surveys and other sentiment proxies. The availability of direct sentiment data is crucially important and thus makes analyzing on-chain data important in the overall context of investor sentiment. When analyzing surveys, you cannot know whether people answered truthfully or if surveys can even capture actual human behavior. Someone could say they are optimistic, but in practice, they are still too scared to buy. With on-chain data, that problem disappears as we directly measure investors' actual behavior. Casella and Paletto (2023) point out that this ability to get real-time behavioral data is even more critical in markets like cryptocurrencies, which lack clear fundamentals.

Blockchain technology and, specifically, on-chain data also work well in large-scale computational analysis. According to Sarmah (2018), each transaction can be tracked from the sender to the receiver through UTXO (unspent transaction output) chains or account-based models (in Ethereum), which makes reconstruction of historical transaction flows possible and thus supports rich time-series modeling. This has encouraged researchers to use machine learning (ML) models directly trained on on-chain metrics such as miner flows, exchange inflows/outflows, and active addresses (Kim et al., 2022; King et al., 2024)

However, despite many advantages, using on-chain data also has disadvantages. While many consider blockchain untraceable, transactions can often be traced back to individuals, raising some privacy concerns. Another disadvantage can be the overwhelming volume of data generated by blockchains. Significant resources are required to analyze the data effectively. This might increase inequality between blockchain users. Lastly, integrating into traditional systems can be challenging and costly despite the potential benefits. To conclude, while blockchain and on-chain data offer many benefits, optimizing those advantages may be difficult.

3 Using On-Chain Data for Market Analysis

3.1 Case Study: Historical Market Trends Based on On-Chain Metrics

3.1.1 Methodology for Case Study Simulations

This section explains the methodology used to evaluate the predictive power of three on-chain metrics: Net Unrealized Profit/Loss (NUPL), MVRV Z-score, and Cumulative Value Days Destroyed CVDD.

3.1.1.1 Data Source and Scope

All the trading simulations were done using closing price data for Bitcoin retrieved from Investing.com from 7 December 2013 to 12 April 2025. This range was chosen because it covers three complete Bitcoin cycles. Also, the 7th of December was the first day that Bitcoin had dropped at least 50% from the cycle top of 2013, and this was chosen as a logical entry for the buy-and-hold benchmark.

All return calculations used closing prices. The data has minor inaccuracies from exchange-specific anomalies (e.g., abnormal price spikes). Still, these do not significantly affect the results as the holding periods were long (e.g., 4143 holding days for the benchmark).

3.1.1.2 Buy-and-Hold Benchmark

Different entry prices for the buy-and-hold investor were considered. First, it was evaluated whether using the same entry prices as those triggered by the strategy signals would provide a fair benchmark comparison. However, this would give the buy-and-hold investor an unrealistically good entry, as an average investor usually cannot buy around the bottoms. This is why the point at which the price had dropped at least 50% from the cycle top was chosen. This is a more accurate representation of the average entry. On the other hand, the masses usually buy around the top (allowing large holders to sell), which could mean that even this entry is too good for the average investor. Nevertheless, this entry (7 December 2013) was chosen because we wanted to ensure the on-chain strategies would not get any unfair advantage over the benchmark to ensure reliable results.

These key metrics measured the benchmark's performance:

- Cumulative and annualized log returns
- Simple returns and annualized simple returns
- Annualized volatility
- Sharpe ratio

3.1.1.3 NUPL-Based Strategies

Three rules-based NUPL strategies were developed, each using different exit thresholds depending on risk appetite. All strategies entered when the NUPL ratio dropped below zero, indicating net unrealized losses and extreme fear. The different exit signals were the following:

- **NUPL 1 (Conservative):** Exit when NUPL ≥ 0.67
- **NUPL 2 (Moderate):** Exit when NUPL ≥ 0.70
- **NUPL 3 (Aggressive):** Exit when NUPL ≥ 0.73

Each strategy entered and exited once per cycle, totalling three trades per strategy. However, during the last cycle (2022-2025), none of the strategies reached the exit thresholds, which resulted in identical ongoing trades. This reduces the differences between the strategies as the third trade is the same. It also indicates that, based on this indicator, the current cycle still has room to grow.

For each trade, the following metrics were calculated:

- Holding days
- Log return and annualized log return
- Simple return and annualized simple return

- Annualized volatility
- Sharpe ratio
- Success rate (percentage of individual trades outperforming the benchmark)

In addition to descriptive statistics, two statistical tests were conducted:

- T-tests on Sharpe ratios: To evaluate whether strategy risk-adjusted returns were statistically significantly higher than the benchmark's.
- Two-sample t-tests (unequal variances) on mean log returns: To evaluate statistical differences in average performance.

The exit signals were subjectively chosen, which raises some limitations to the robustness. However, no clear universal rules exist for the exit thresholds, so some subjectivity is required. Choosing three different thresholds allows for evaluating whether this strategy only works when the exit threshold is perfectly chosen and whether the approach works only with hindsight. Although the last identical trades compress performance differences, they still represent the real world, where investors hold ongoing positions if exit signals have not yet been triggered.

3.1.1.4 MVRV Z-Score Strategies

The same methodology was applied to the MVRV strategies. All strategies entered when MVRV Z-score dropped below -0,2, signalling fear and uncertainty. Exit thresholds varied depending on the risk:

- **MVRV 1 (Conservative):** Exit when MVRV Z-score ≥ 5
- **MVRV 2 (Moderate):** Exit when MVRV Z-score ≥ 6
- **MVRV 3 (Aggressive):** Exit when MVRV Z-score ≥ 7

With the MVRV Z-score, there were also three trades per strategy, and the last trade was still ongoing, resulting in identical third trades.

Performance evaluation was conducted the same way as with the NUPL strategies, and the same limitations apply to this metric.

3.1.1.5 CVDD Entry vs. Random Entries (Monte Carlo Simulation)

The long-term bottom detection tool CVDD was tested using a methodology different from the first two metrics. Entry was triggered when Bitcoin's price fell within a 1% distance from the CVDD curve. In each of the three cycles, the CVDD trade was compared against 100 randomly timed entries drawn from within:

- **±50-day range around the cycle bottom**
- **±75-day range around the cycle bottom**
- **±100-day range around the cycle bottom**

The exit prices were the same for all trades, as this simulation aimed to measure the CVDD's ability to predict bottoms. The exit price was always the cycle top, or in the third cycle, the last closing price of the data (12 April 2025).

<u>The performance was evaluated with the following metrics:</u>

- Holding days
- Log returns
- Annualized log returns
- P-values representing the percentage of random entries outperformed by the CVDD-based entry (e.g., 99% means CVDD beat 99 out of 100 random entries)

The p-values were measured for the log returns and the annualized log returns. These Monte Carlo simulations measured whether CVDD's prediction precision significantly exceeds random chance within the three different bottoming ranges. The decision to use the 1% distance from the CVDD curve as the long signal was subjectively chosen and could be influenced by hindsight bias. However, the point was merely to test the indicator's ability to predict bottoms and not necessarily to create a perfect trading system.

3.1.2 Net unrealized profit/loss

The net unrealized profit/loss (NUPL) ratio is one of the most used metrics to gauge investor sentiment among market participants. It estimates the amount of coins currently in profit or loss and can be a reliable indicator of when the market tops and bottoms. High values indicate large amounts of unrealized gains – often signaling euphoria – while low values indicate fear and capitulation.

What makes the NUPL ratio especially compelling is its simplicity and ease of use. Unlike machine learning models, which require high computing power and strong technical knowledge, the NUPL ratio can be calculated using publicly available data. While deep learning models have proved helpful for price prediction (Kim et al., 2022; Casella & Paletto, 2023), the NUPL ratio remains useful precisely because of its accessibility and strong predictive ability, even as a single indicator.

Tamas Blummer, Tuur Demeester, and Michiel Lescrauwaet (2019) calculate the NUPL ratio using market capitalization and realized value. Market capitalization is obtained by multiplying the price of Bitcoin by the number of coins in circulation. Realized value is the average price at which each Bitcoin was last transferred from one wallet to another. That average is then multiplied by the number of coins in circulation. Then, the realized value is subtracted from the market capitalization, giving us an unrealized profit/loss ratio, which is then divided by the market capitalization to finally get the net unrealized profit/loss ratio or, more simply, NUPL ratio.

The formula for the NUPL ratio is:

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

As seen in Figure 1, the NUPL ratio has been proven to be an effective tool in forecasting the optimal prices for buying and selling Bitcoin. You can outperform the market by purchasing in the green zone and selling when the price nears the red zone. This showcases the NUPL ratio's potential as a strategic tool for making investment decisions and timing the market in the constantly fluctuating world of cryptocurrencies.

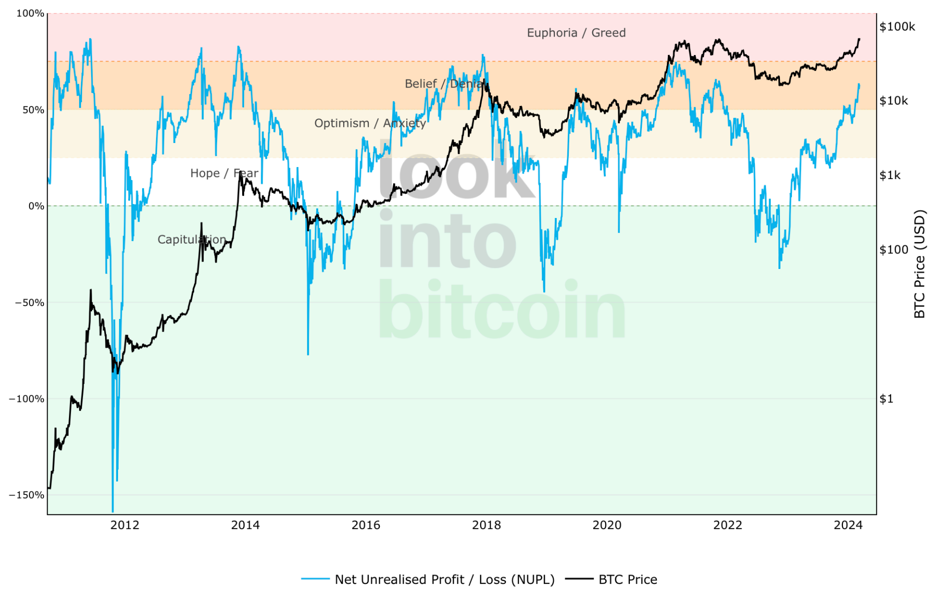

Figure 1. Bitcoin: Net Unrealized Profit/Loss (NUPL) (Look Into Bitcoin, 10.3.2024).

From a behavioral standpoint, the NUPL ratio visualizes cognitive biases like loss aversion, anchoring, and herd behavior. When the ratio is high (over 0,75), it indicates that people do not realize their gains and are instead anchored to their unrealized gains due to emotions like greed and FOMO. Conversely, when the ratio falls to negative, it signals the opposite. Investors are panicking and selling their coins at a loss as they are not capable of seeing any more downside. These behaviors repeat, making the NUPL ratio a valuable proxy for investor sentiment.

This thesis creates three trading strategies using threshold-based signals from the NUPL ratio to measure the real-world applicability of NUPL-based strategies in price prediction. These three strategies are compared against a buy-and-hold strategy between 7 December 2013 and 12 April 2025. Each NUPL strategy goes long when the NUPL ratio falls below the zero level and exits at different upper thresholds. These upper thresholds are 0.67 in the NUPL 1-strategy (conservative), 0.7 in the NUPL 2-strategy (moderate), and

0.73 in the NUPL 3-strategy (aggressive). This is done to test the sensitivity and robustness of NUPL-based strategies. Table 1 presents the comparative performance results across key metrics, including annualized log returns, volatility, Sharpe ratios, cumulative return multiples, t-statistics for Sharpe ratio significance, and success rates.

Table 1. Performance comparison of NUPL-based strategies.

Strategy	Ann. Log Return	Ann. Volatility	Sharpe Ratio	t-stat vs B&H	Return (x)	Success Rate
NUPL 1	0,63	0,63	1,00	29,03***	118x	100 %
NUPL 2	0,68	0,63	1,08	33,34***	179x	100 %
NUPL 3	**0,87**	0,66	**1,32**	**47,66*****	**1234x**	100 %
Buy & Hold	0,42	0,94	0,45	-	120x	-

Note: *Return (x) denotes the multiple of the initial investment value (e.g., 120x means the investment grew 120-fold). Sharpe Ratio is calculated as the annualized log return divided by annualized volatility. The t-statistic tests whether the Sharpe ratio of each strategy is significantly different from that of the Buy & Hold strategy. Success Rate indicates the percentage of trades in which the strategy outperformed Buy & Hold. *** indicates statistical significance at the 1 % level.*

The results show the superior risk-adjusted performance of all three NUPL strategies when compared against the buy-and-hold benchmark. The most aggressive strategy, NUPL 3, achieved an annualized log return of 0,87, significantly higher than the benchmark's 0,42. Also, it beat the benchmark in risk-adjusted returns (1,32) measured by the Sharpe ratio and cumulative returns (1234x). Similar outperformance is seen in NUPL 1 and NUPL 2, although slightly lower in absolute performance, and NUPL 1 fell short of the benchmark in cumulative returns, 118x versus the benchmark’s 120x. However, the Sharpe ratio and annualized log returns are still higher in each strategy than the benchmark.

All NUPL strategies achieved a 100% success rate, measured by the proportion of individual trades that achieved a higher annualized log return than the buy-and-hold strategy. This shows the reliability of the NUPL indicator in each cycle with different risk levels and proves the ratio’s ability to consistently identify favorable entry and exit signals.

To test the results’ statistical significance, t-tests were conducted on the Sharpe ratios of each strategy and compared to the buy-and-hold strategy. All NUPL-based strategies were statistically highly significant, with t-statistics of 29,03, 33,34, and 47,66 for NUPL 1, NUPL 2, and NUPL 3, respectively. Each strategy was significant at the 1% level (***). However, earlier t-tests on mean returns were not statistically significant, likely due to high return volatility and skewed distributions. Still, Sharpe ratio-based tests provide a more robust and informative comparison in this context.

These results also align with previous literature. For example, Łęt, Sobański, Świder, and Włosik (2022) show in their study that on-chain indicators like NUPL ratio and Percent of Addresses in Profit (PAP) can predict cryptocurrency prices and outperform traditional buy-and-hold strategies. They found that strategies incorporating the NUPL ratio had significantly better cumulative and risk-adjusted returns than passive strategies. NUPL-based strategies achieved a 100% success rate in outperforming passive benchmarks. This was even true across different periods, such as the volatile market conditions during the COVID-19 pandemic. These results support the idea that the NUPL ratio can capture investor emotions like fear and greed more effectively than many traditional indicators like MACD and RSI. This is partly because cryptocurrencies lack intrinsic value and are thus more sentiment-driven.

When combined, this thesis's results and Łęt et al. (2022) findings strongly support both hypotheses. In support of H1, these results show that the NUPL ratio aligns with sentiment: extreme values in net unrealized profits correspond to price extremes. Furthermore, Hypothesis 2 is supported by the consistent outperformance of NUPL strategies in risk-adjusted and annualized log returns.

However, this strategy might not work as well in the future as it does now. Still, it illustrates how effectively one indicator can help investors optimize their buying and selling areas. Also, the decision to use the ratios of 67, 70, and 73 as an exit signal was subjectively chosen and may be influenced by hindsight bias. However, as said, the purpose of this strategy was not to create a perfect trading system but to illustrate the predictive capabilities of even a single on-chain metric. Also, given the significant overperformance, it can be argued that even if the strategy were not so optimally built, it would still be a beneficial tool for predictive analysis.

Table 2 provides a breakdown of sentiment zones based on different NUPL ratio thresholds to further clarify the interpretation of NUPL values and their corresponding market phases.

Table 2. NUPL Market Sentiment Zones.

NUPL Market Sentiment Zones

NUPL Value	Market Sentiment	Interpretation
> 0,75	Extreme Greed	Potential market top
0,5 – 0,75	Optimism	Bullish, but risk increasing
0 – 0,5	Hope / Neutral	Recovery or consolidation
< 0	Fear / Capitulation	Potential market bottom

The NUPL ratio works so well because it captures market sentiment effectively. When the net unrealized profit is too high, it indicates over-optimism and greed, and when the indicator goes below zero, it indicates over-pessimism and fear. However, it is essential to note that this methodology has some limitations despite the significant overperformance. Firstly, due to the volatility of cryptocurrency markets, sudden price fluctuations can affect the profitability of the NUPL ratio. However, the ratio's long-term approach may mitigate the adverse effects of market volatility, but it is essential to consider that risk. Additionally, using a single metric, such as the NUPL ratio, to predict the market can be risky. Łęt et al. (2022) note that integrating multiple on-chain metrics can benefit more than just relying on one indicator. This may improve signal quality and reduce false positives – an idea that aligns with limitations later discussed in this thesis.

Furthermore, this approach may only apply to Bitcoin and not other cryptocurrencies. Even though other cryptocurrencies follow the movements of Bitcoin closely, they may have less historical data or consistent investor profiles. This might result in weaker signals and more noise when applying the NUPL ratio to other cryptocurrencies. For example, some coins have heavy token unlocks, low liquidity, or centralized ownership, which can unpredictably affect the NUPL ratio. Additionally, the cryptocurrencies' regulatory landscape is constantly evolving, which may introduce additional uncertainty and potentially impact market sentiment and the cryptocurrency market.

3.1.3 MVRV Z-Score

The MVRV Z-score is a widely used on-chain metric that evaluates whether Bitcoin is undervalued or overvalued relative to its fair value. It uses three components to do that. Two elements are the same as in the NUPL ratio: market value (MV) and realized value (RV). The third is the Z-score, the standard deviation between MV and RV. The metric was first introduced by David Puell and Murad Mahmudov in 2018 and has since gained popularity due to its ability to predict cycle extremes in the price of Bitcoin.

As mentioned, the MVRV Z-score uses the concepts of market and realized value, like the NUPL ratio. Market value is the total current value of Bitcoin calculated by multiplying the current price by the circulating supply, i.e., the number of coins currently in circulation. In contrast, realized value tries to offer a proxy for the average cost basis of all investors by tracking the price of each unit of Bitcoin last moved on-chain. The difference between the NUPL ratio and the MVRV Z-score is in the use of the Z-score in the MVRV Z-score indicator. The Z-score measures how far the current value is from the realized value, which is then adjusted with the standard deviation of market value.

The formula for the MVRV Z-score is:

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

Where:

Market Value (MV) represents Price x Circulating Supply

Realized Value (RV) represents the Sum of coin values at the last on-chain movement

$\boldsymbol{\sigma(MV)}$ represents the standard deviation of Market Value

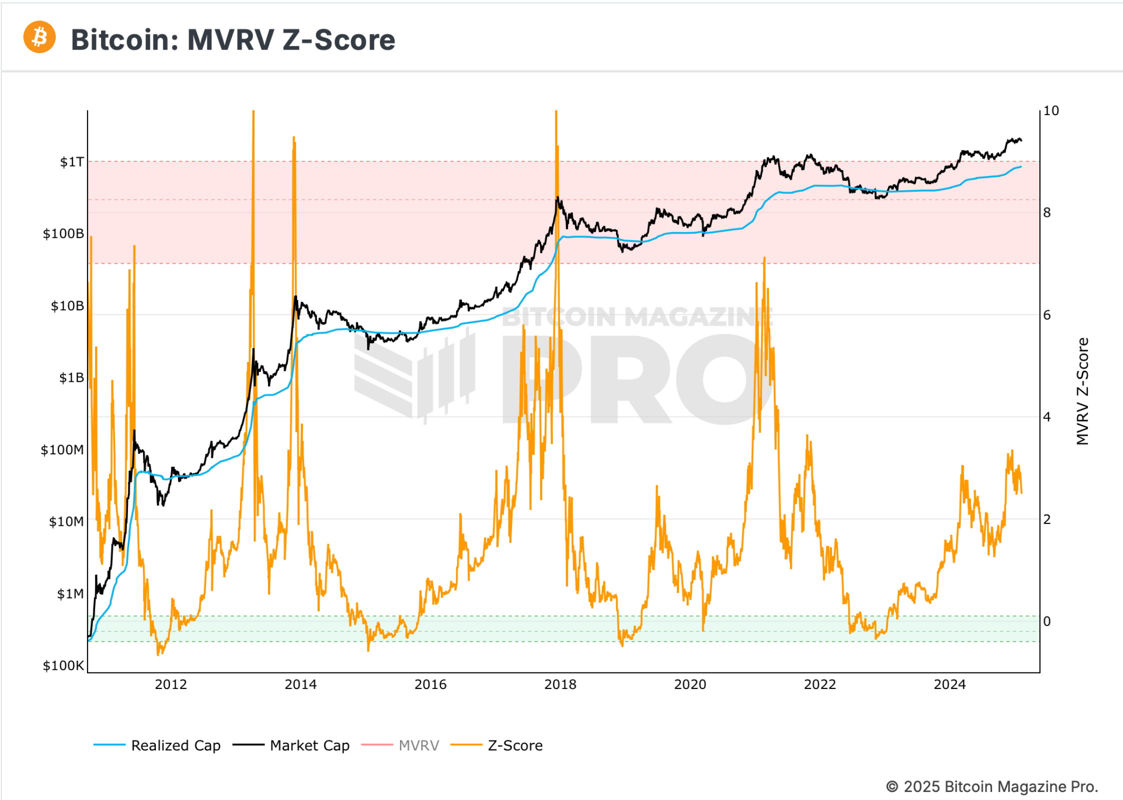

Figure 2. Bitcoin: MVRV Z-Score (Bitcoin Magazine Pro, 8.2.2025).

As seen in Figure 3, the MVRV Z-score has historically been able to predict key turning points in the price of Bitcoin. When the price deviates too far from its RV, it indicates overvaluation (red zone) or undervaluation (green zone).

1) Bull Market Peaks:

- Each Bitcoin cycle saw the MVRV Z-score enter the red zone before a major correction
- The peaks in 2011, 2013, 2017, and 2021 were formed while the MVRV Z-score was in the red zone

2) Bear Market Bottoms:
   - Every major bottom was formed inside the green zone (2011, 2015, 2018, 2022)
   - Investors who bought inside the green zone have seen outsized returns in the following bull markets

Like the NUPL ratio, the MVRV Z-score also captures investor sentiment, which allows it to predict Bitcoin price trends. Since RV reflects the aggregate cost basis of Bitcoin holders, the deviation between MV and RV indicates whether investors are euphoric or fearful.

During bull markets, when Bitcoin's market value rises significantly above its realized price (MVRV Z-score red zone), it indicates that investors are holding a lot of unrealized profits. This typically means heightened speculative activity, FOMO (fear of missing out), and media frenzy. These conditions usually lead to market tops, like in 2011, 2013, 2017, and 2021.

On the other hand, during bear markets, when the market value drops below its realized value (MVRV Z-score green zone), it signals capitulation and extreme fear. During these times, long-term or institutional investors accumulate their positions, buying Bitcoin from panicking retail investors. Historically, periods of extreme pessimism indicated by the MVRV Z-score have marked the most optimal buying opportunities, as seen in the bear market bottoms of 2011, 2015, 2018, and 2022.

Looking closely at previous market cycles, it becomes clear that the MVRV Z-score is reliable at signaling euphoria and capitulation. For example, in 2013, the MVRV Z-score rose to over 8 when Bitcoin peaked at around 1,200 dollars. This was a sign of extreme greed as the price deviated far from the realized price, leading to a market top and an over 80 % drop in price. Almost the same pattern formed during the 2017 bull market when Bitcoin's price traded around 20,000 dollars before falling into a bear market and losing over 80 % in value. In both cases, the MVRV Z-score warned investors of overvaluation and high speculation, allowing them to risk off the market.

The MVRV Z-score was effective not only at providing sell signals during bull markets but also at buying signals around market bottoms. In 2015, MVRV Z-score traded in the green zone for months while the price traded below 300 dollars. This later proved to be the accumulation zone before the next bull run, leading to Bitcoin's price reaching close to a 20,000 dollar price. This same phenomenon occurred during the 2018 and 2022 bear markets, eventually leading to much higher prices. In all these cycles, the MVRV Z-score was able to signal deep undervaluation, giving traders and investors excellent long-term entry points.

To further evaluate the predictive power of the MVRV Z-score, three rule-based strategies were created and backtested across three Bitcoin cycles between 2013 and 2025. These tests were designed to see how well the MVRV Z-score could identify favorable long-term buying and selling opportunities. Each strategy entered a long when MVRV Z-score crossed the -0,2 mark. The exit thresholds were different between the strategies; MVRV 1 (conservative) exited above 5, MVRV 2 (moderate) exited above 6, and MVRV 3 (aggressive) exited above 7. Table 3 below summarizes the results and compares them to the buy-and-hold benchmark.

Table 3. Performance comparison of MVRV Z-Score-based strategies.

Strategy	Ann. Log Return	Ann. Volatility	Sharpe Ratio	t-stat vs B&H	Return (x)	Success Rate
MVRV 1	0,79	0,63	1,26	42,24***	334x	100 %
MVRV 2	0,94	0,67	1,42	51,93***	1630x	100 %
MVRV 3	**1,01**	0,68	**1,49**	**56,49*** **	**3262x**	100 %
Buy&Hold	0,42	0,94	0,45	-	120x	-

Note: *Return (x) denotes the multiple of the initial investment value (e.g., 120x means the investment grew 120-fold). Sharpe Ratio is calculated as the annualized log return divided by annualized volatility. The t-statistic tests whether the Sharpe ratio of each strategy is significantly different from that of the Buy & Hold strategy. Success Rate indicates the percentage of trades in which the strategy outperformed Buy & Hold. *** indicates statistical significance at the 1 % level.*

The results show the superior risk-adjusted performance of all three MVRV strategies when compared against the buy-and-hold benchmark, like the NUPL strategies. The most aggressive strategy, MVRV 3, achieved an annualized log return of 1,01, significantly higher than the benchmark's 0,42. Also, it beat the benchmark in risk-adjusted returns (1,49) and cumulative returns (3262x). All the MVRV strategies achieved higher annualized log returns, Shape ratios, and cumulative returns than the benchmark. They also beat NUPL-based strategies in every metric, although MVRV 2 and 3 had higher volatilities than the respective NUPL strategies.

All MVRV strategies also achieved a 100% success rate, measured by the proportion of individual trades that achieved a higher annualized log return than the buy-and-hold strategy, just like the NUPL strategies. This shows that MVRV strategies have the same reliability as the NUPL strategies, with higher risk-adjusted and cumulative returns.

The statistical significance of the MVRV strategies was also tested. T-tests were conducted on the Sharpe ratios of each strategy and compared to the buy-and-hold strategy. All MVRV-based strategies were statistically highly significant, with t-statistics of 42,24, 51,93, and 56,49 for MVRV 1, MVRV 2, and MVRV 3, respectively. Each strategy was significant at the 1% level (***). Like the NUPL strategies, the earlier t-tests on mean returns were not statistically significant, likely due to the same reasons (high return volatility and skewed distributions).

These results support the hypotheses similarly to the evidence from the NUPL strategies. The fact that MVRV strategies even outperformed NUPL strategies in every key metric further emphasizes MVRV's support for the second hypothesis. It also adds to the overall evidence for on-chain metrics and their price prediction capabilities. The significance of overperformance in risk-adjusted returns also strengthens the argument that investor sentiment, measured through the on-chain data, has a strong power in predicting prices.

The strength of the MVRV Z-score is its ability to turn on-chain investor behavior into a measurable metric. When Bitcoin climbs far above the realized price, it signals confidence and overconfidence. Investors and traders become unwilling to sell their gains into actual profits as they want more. They are confident that the prices will rise and want more from the markets. This leads to substantial unrealized gains and causes the MVRV Z-score to rise. This behavior aligns with classical behavioral finance theories like the disposition effect and FOMO.

On the other hand, when the price falls below the realized price, sentiment is close to despair, anger, and depression. Those traders unwilling to sell their profits into actual gains are now holding the same coins at a loss, and they feel devastated, unable to tolerate more loss. This period typically marks the point when institutional or well-informed investors start accumulating assets from the retail traders. These emotions are precisely what MVRV Z-score visualizes with its red and green zones, making it a valuable tool for timing the market and measuring investor sentiment.

While the MVRV Z-score has proven highly effective and valuable in market timing and sentiment measurement, it has limitations. For example, during the 2021 bull market, it barely breached the red zone, unlike in previous cycles, when it spiked much deeper into the overvaluation zone. This could mean that as Bitcoin matures and grows in value, the parabolic rises might also be less pronounced, and thus, the MVRV Z-score never gets into the red zone before a market top. When looking at the MVRV Z-score, we are still far from the cycle top, while many crypto analysts think we are close to a top or have potentially even made one already. This is something that only time will tell, but given the lower reading of 2021, it could mean that this cycle, the indicator will not rise inside the red zone. For this reason, it might benefit traders and investors to combine on-chain metrics and not only follow a single one. However, the fact that the conservative and moderate MVRV strategies were also able to beat the buy-and-hold strategy shows that it is not necessary to wait for the most optimal exit signal. Regarding the entries, the bottom ranges still play out similarly between cycles. This gives more weight to MVRV Z-score's ability to signal accumulation zones, and thus, those might prove to be more accurate in the future.

3.1.4 Cumulative Value Days Destroyed (CVDD)

The third promising metric chosen for a more thorough examination is the Cumulative Value Days Destroyed (CVDD), created by the famous on-chain analyst Woonomic. This metric is one of the most empirically reliable tools for predicting market bottoms, combining time-based and value-based behavior in the Bitcoin ecosystem. The CVDD metric can be found for free through platforms like Bitcoin Magazine Pro, where investors can visually assess its historical performance in relation to Bitcoin's price.

CVDD consists of a concept called Coin Days Destroyed (CDD), a foundational metric in on-chain analysis that combines the value of coins moved and the length of time they were held. It is calculated by multiplying the value of coins moved by the time they were held in a wallet. For example, if 10 BTC were held in a wallet for 100 days before being transferred, the transaction would "destroy" 1000 coin days. This means that the longer the coins were held before the transaction, the higher their weight in the metric. Giving more weight to longer-held coins creates a proxy for conviction. Movements of long-held coins usually have more meaning, both emotionally and financially, than coins only held for the short term. This makes CDD a valuable sentiment indicator.

CVDD is then created by building on top of CDD by taking the cumulative sum of the USD value of coin days destroyed and then adjusting for the age of the Bitcoin network. From a theoretical standpoint, CVDD illustrates the “thermodynamic” cost of transacting in Bitcoin. Accounting for the time held and value when transacted inserts a time-value dimension into the analysis. One of the main assumptions of the metric is that over time, capital inflows to Bitcoin create this energetic floor below the price where it rarely falls. When this data is aggregated over the network's life, it makes a long-term valuation base, which can usually mark the bottoms of each cycle.

The calculation of CVDD consists of the following steps:

1. Calculate Coin Days Destroyed for every Bitcoin transaction.
2. Multiply the CDD by the USD value of the transaction at the time it was made.
3. Aggregate the value-time destruction across all transactions over the history of the Bitcoin network.
4. Normalize the cumulative sum by the age of the Bitcoin market, expressed in days.
5. Apply a scaling factor (usually 6 million) to adjust the metric in line with Bitcoin’s price data. The factor is not based on a theoretically grounded constant, but it serves a practical purpose in aligning CVDD’s output with historical price levels for comparative analysis.

<u>Mathematically, the metric can be approximated as:</u>

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

Finally, we get a line that acts as the final floor for Bitcoin’s market value, which should be the point at which price bottoms during extremely negative sentiment. In Bitcoin’s history, it has never meaningfully breached that level, providing investors with a robust long-term valuation anchor.

CVDD differs from many other on-chain metrics because it emphasizes long-term holders. The metric and Bitcoin's price tend to meet during extreme market stress when even the long-term holders start to capitulate. This is a significant indication of extreme fear, as even the ones who have held for the entire bear market finally decide to sell. This is usually the point at which market bottoms form, as no sellers are left to sell.

The real value of CVDD can be visualized when it is overlayed with Bitcoin's price action over time. Figure 4 below is freely available for investors in Bitcoin Magazine Pro, and it shows the interaction of the CVDD curve and Bitcoin's price over the past cycles. This figure provides empirical evidence of CVDD's accuracy in predicting key market bottoms in all significant downturns.

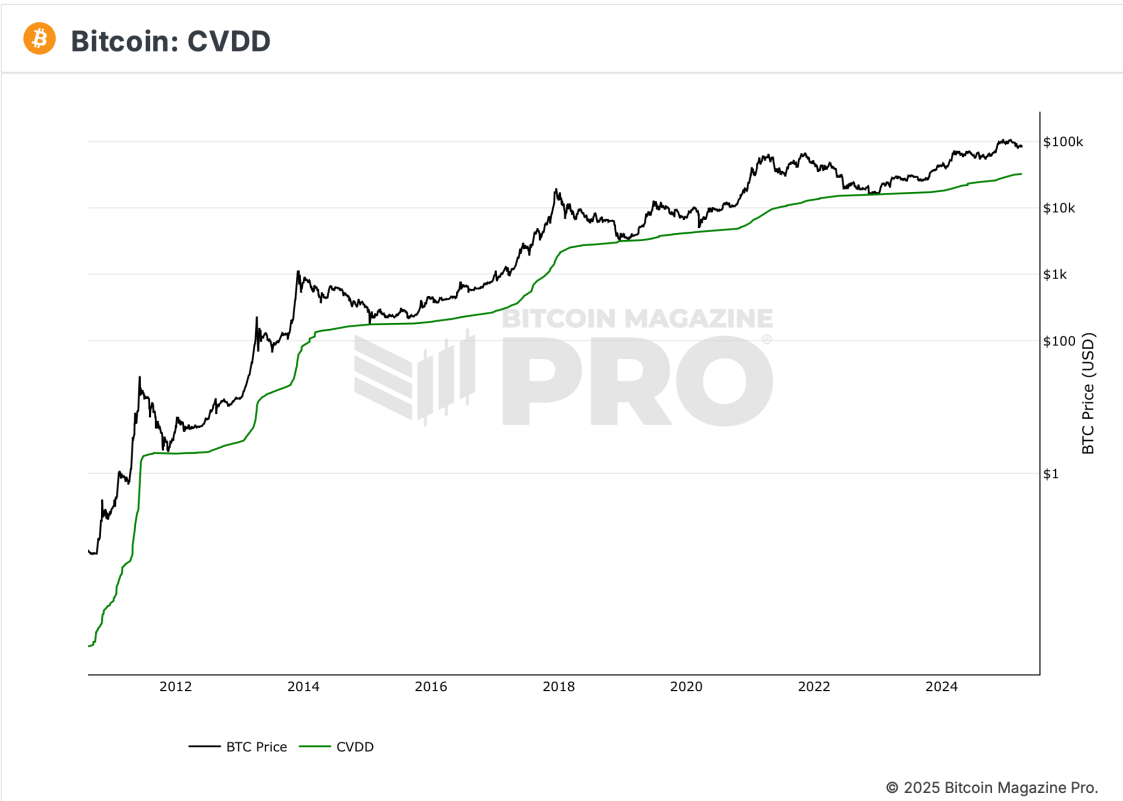

Figure 3. Bitcoin: CVDD (Bitcoin Magazine Pro, 4.4.2025).

After examining the figure, it becomes clear that CVDD has acted as a strong lower bound for Bitcoin's price since its inception. Every cycle bottom formed when the price was almost precisely at the CVDD curve. The only time the price breached the line was in 2022, and even then, it was only by a small margin.

The accuracy of the metric suggests that it is not merely a coincidence but a strong indicator of extreme fear and capitulation. What it indicates is the long-term cumulative behavioral exhaustion of traders, and more specifically, the ones who have held for a long time. When those long-term holders finally give in and sell, it disproportionately adds to the CVDD calculation. Although the CVDD curve rises steadily, Bitcoin’s price usually falls towards it during extreme fear and panic, making CVDD a valuable estimate of deep market undervaluation.

For example, during the 2011 bear market, the CVDD was around two dollars for months while the price traded above ten dollars. The bottom was ultimately only found when the price corrected to approximately two dollars, falling just short of the CVDD curve. This same phenomenon happened during every subsequent cycle bottom (2015, 2019, and 2022).

To assess the effectiveness of CVDD in the bottom prediction, we conducted a series of Monte Carlo simulations across three different Bitcoin cycles. We simulated 100 randomly timed entries between ±50-day, ±75-day, and ±100-day windows around the cycle bottoms in each cycle. The cumulative and annualized log returns of the random entries were then compared to those achieved using CVDD to enter a trade. All trades within a cycle had the same exit price, which was the cycle top or, in the current cycle (2022-), the last closing price of the data. CVDD-strategy entered a trade when Bitcoin’s price was 1% away from or closer to the CVDD line. Table 4 summarizes the results by the percentage of simulations in which the CVDD-based trade outperformed the random entries. It effectively expresses the statistical dominance as a “success rate” p-value.

Table 4. Statistical Comparison Between Random and CVDD-Based Entries**.**

Range	Log Return p-value (CVDD vs. Random)	Annualized Log Return p-value (CVDD vs. Random)
Cycle 1		
±50d	99 %	99 %
±75d	100 %	100 %
±100d	100 %	98 %
Cycle 2		
±50d	100 %	84 %
±75d	100 %	71 %
±100d	98 %	63 %
Cycle 3		
±50d	99 %	66 %
±75d	97 %	76 %
±100d	100 %	80 %
Avg.	**99 %**	**82 %**

Note: *P-values indicate the percentage of 100 randomly timed entries—generated via Monte Carlo simulations within each respective time window—that underperformed the CVDD-based entry in terms of cumulative and annualized log returns.*

The results strongly show the CVDD's ability to predict bottoms. Across all cycles and ranges, CVDD outperformed most randomly timed entries. Regarding cumulative log returns, CVDD achieved an average p-value of 99%, showing its near-perfect ability to predict bottoms. The annualized log returns were slightly lower, having an average p-value of 82%. This indicates that even though CVDD was able to capture the almost exact bottom every time, the holding period was a bit longer than with some of the random entries, reducing the annualized log returns. Still, in Cycle 1 (2013-2017) annualized log returns, CVDD got a 100% p-value in the ±75 range and a close to 100% p-value in the two other ranges. These results strongly support the second hypothesis, showing a clear prediction capability in on-chain data.

It needs to be stated that CVDD does not predict tops or short-term movements. It is designed to detect zones of extreme undervaluation for long-term position building. This means it is not for active traders but rather for investors wanting to find "generational buying opportunities." Also, it is best used when considered as a zone rather than a single point in which to buy. Sometimes, the price falls short, and sometimes, it can breach the curve slightly, but every bottom still forms close to the curve. This makes it a high-probability accumulation zone.

To conclude, CVDD is a key example of how long-term sentiment data can be used to predict prices and, more specifically, market bottoms. While it does not provide traders with short-term signals, its ability to indicate strong accumulation ranges has been validated in every Bitcoin market cycle. Using the value and age of coins moved, it forms the cumulative emotional and financial toll of long-term traders, thus making it one of the most behaviorally backed on-chain metrics.

3.1.5 Summary of Case Studies and Hypothesis Evaluation

This section combines the three case studies (NUPL ratio, MVRV Z-score, and CVDD) and evaluates each regarding this thesis's two hypotheses.

- H1: On-chain metrics reflect investor sentiment
- H2: On-chain metrics can predict Bitcoin prices

While all the metrics try to measure market sentiment and value dislocation, they do it differently. NUPL ratio directly measures unrealized gains and losses, giving us a window into mass herding and extreme emotions. MVRV Z-score builds upon the NUPL ratio by standardizing it and evaluating how market value deviates from the realized price. MVRV Z-score indirectly measures investor sentiment through valuation extremes. CVDD, on the other hand, differs from these two metrics as it measures long-term behavioral conviction by analyzing the time-weighted value of spent coins.

These metrics offer a multidimensional view of sentiment and its relationship to price. The NUPL ratio reacts to real-time sentiment shifts, the MVRV Z-score uses standard deviation to contextualize sentiment, and CVDD shows behavioral exhaustion through coin movements. These different methods allow a more accurate understanding of market cycles and a better view of short-term sentiment shifts. Table 5 summarises the key elements of the metrics and whether they support the two hypotheses.

Table 5. Summary of Case Study Findings.

Metric	Type	H1 Supported (Sentiment)	H2 Supported (Price)	Key Findings
NUPL	Sentiment-based	Yes	Yes	Captured market tops and bottoms. Strong visual signal. Best for contrarian strategies.
MVRV Z-Score	Valuation-based (statistical)	Yes	Yes	Detected valuation extremes. Fast-moving sell signals. Useful for spotting euphoria and capitulation.
CVDD	Time-value behavioral	Partial	Yes	Identified consistent cycle bottoms. Structural, not emotional. Best for detecting undervaluation.

Both the NUPL ratio and the MVRV Z-score strongly support the first hypothesis of on-chain metrics being sentiment measurements. NUPL ratio measures unrealized profits and losses, capturing the mood of the market participants. You are more likely to hold on to profits when greedy or overly optimistic, and vice versa. Its similar cyclical behavior across bull and bear cycles shows how sentiment forms repeatable patterns captured by the NUPL ratio. The fact that the NUPL ratio has consistently aligned with market peaks and bottoms strongly supports the sentiment hypothesis.

Although from a more valuation-based view, the MVRV Z-score also supports the first hypothesis. It measures the degree to which the price has deviated from the average cost basis, indicating over- or undervaluation. Because of the lack of intrinsic value, the deviations from the average cost basis measure simply the sentiment change of the market participant, thus supporting the hypothesis. This, coupled with the fact that MVRV Z-score peaks and bottoms coincided with price extremes, proves its ability to measure sentiment accurately.

In contrast, the CVDD only partly supports the first hypothesis as it does not measure investor sentiment in real time but instead captures long-term structural behavior. CVDD reflects the actions of long-term holders who slowly sell out of their Bitcoins during bull markets and usually stop around market tops. When they stop selling, a base in CVDD is created, to which the price usually falls, indicating deep undervaluation, pessimism, and fear. This means CVDD does not directly measure crowd emotion and instead shows the pause in profit-taking by long-term holders, creating a structural base that indirectly indicates market sentiment when the price falls to it. These observations are supported by the results in Table 4, showing that CVDD could constantly outperform randomly timed entries. This highlights its role as a proxy for long-term fear rather than real-time sentiment. While it does not fully track sentiment, it provides a unique view into sentiment rooted in conviction and capital rotation.

In terms of H2, all three metrics support the hypothesis that you can use on-chain metrics to predict cryptocurrency prices. MVRV Z-score and NUPL ratio were especially effective at predicting market tops and bottoms. However, how those two metrics gave their signals was quite different. The NUPL ratio rose more gradually during the bull market, giving investors time to react and slowly exit the market. By comparison, MVRV Z-score was more abrupt in giving signals; it could rise quickly from mid-range values to the historically overvalued territory (e.g., from 4 to 8) and take only days or weeks. This means that the MVRV Z-score needs to be monitored more closely, but it works better in pinpointing the final-stage price extremes than the NUPL ratio, which is more suited for long-term strategic planning. This was strongly supported by our simulations, where although all NUPL ratio and MVRV strategies outperformed the buy-and-hold benchmark in many metrics like Sharpe ratios, annualized log returns, and success rates, MVRV strategies were still better at pinpointing the final extremes and thus outperforming NUPL strategies in all tested metrics.

CVDD was developed only to predict bottoms, and it does so almost perfectly. It is less reactive and should not be used in short-term trading, but it works well in identifying long-term accumulation ranges. Like the MVRV Z-score and NUPL ratio, it can predict Bitcoin price trends and thus strongly support the second hypothesis. Our Monte Carlo simulations showed that CVDD outperformed 97% to 100% of random entries across all three cycles tested (2015, 2018, 2022) in cumulative log returns, giving even more evidence of its ability to predict accumulation zones.

A key conclusion from the analysis of these three metrics is that different dimensions of sentiment and market structure emerge depending on the metric's design. NUPL ratio can capture short-term sentiment in real-time, MVRV Z-score captures significant value dislocations, and CVDD focuses on structural time-value exhaustion. Even with these differences, what becomes clear is that sentiment leaves traceable footprints in on-chain data that can be measured and then used to predict turning points in Bitcoin's cyclical market.

While this section focuses on Bitcoin, its relevance extends beyond that. The cryptocurrency market often follows Bitcoin's movements, and thus, these indicators not only predict Bitcoin's prices but also the entire market to some extent. Strong historical correlations between Bitcoin and other digital assets demonstrate this relationship. For example, between 2018 and 2024, the Pearson correlation between Bitcoin and Ethereum was over 0,90, indicating an almost identical price action. XRP and ADA also showed moderately strong correlations of 0,57 and 0,60. This suggests that the sentiment derived from Bitcoin will likely represent the entire market to some extent and is not only applicable to Bitcoin.

In conclusion, the case studies offer strong empirical support for both research hypotheses. This was proven by theoretical and historical visual evidence and through different backtested trading strategies and statistical comparisons. While the support was more decisive for H2 than H1, it can still be said with confidence that on-chain data reflects sentiment and can be used to predict price trends in Bitcoin and potentially in other digital assets. These metrics, used in combination, could provide the most successful implementation. For this reason, the next section will analyze multiple on-chain metrics with ML models to achieve even better accuracy and reduce the likelihood of false signals.

3.2 Expanding Predictive Power Through Machine Learning

Applying machine learning (ML) and deep learning (DL) models can significantly increase the accuracy of price predictions. On-chain data combined with ML models can create a powerful tool that is better at predicting prices than each method alone. In cryptocurrencies, it can be hard to capture non-linear relationships, and traditional models like linear regression (LR) fail to capture complex dependencies (Lahmiri and Bekiros, 2019).

Using one indicator, like the NUPL ratio, can benefit investors. However, the predictive power can be increased when different on-chain metrics are combined with machine learning techniques. This section evaluates previous research that uses ML models to analyze on-chain metrics to predict cryptocurrency prices more accurately.

Jagannath et al. (2021) use Long Short-Term Memory (LSTM) networks with on-chain metrics, trying to predict Ethereum (the second-largest cryptocurrency by market cap) price movements. LSTM networks are a DL model designed to process sequential data and are good at capturing temporal dependencies. They find that LSTMs can outperform traditional time-series forecasting techniques when analyzing blockchain data. The ability to process sequential dependencies makes LSTMs especially good at on-chain sentiment-based predictions.

Their model uses different on-chain metrics to forecast Ethereum price trends. The metrics are active addresses, gas fees, and miner revenue, each having a crucial but different role in this predictive model. Active addresses measure network adoption, where an increase in unique addresses can be seen as a sign of investor confidence and usually precedes price increases. Gas fees can also measure increased investor activity, as spikes in gas fees indicate network congestion and typically mean there is a lot of speculative interest. Miner revenue, on the other hand, is more of a supply and demand factor. If the miner revenue is declining, it might result in miners selling some of their coins to combat the decline in revenue. So, a decrease in miner revenue signals a potential selling pressure from miners and, thus, an increase in supply.

Jagannath et al. (2021) show that combining several on-chain metrics can improve prediction accuracy. Also, the study reinforces the importance of blockchain data in price forecasting, showing how Ethereum's on-chain behavior works as a leading indicator. This, combined with the DL model, proves to be a successful way to measure the sentiment of market participants and then use that to predict the price movements. They reveal that the LSTM-RNN model achieves an accuracy of 86,94% when predicting Ethereum prices with on-chain metrics. All the selected inputs (active addresses, gas prices, and miner revenue) have a strong linear correlation with price.

Building on this, Kim et al. (2022) take DL models further by creating a Self-Attention-Based Multiple LSTM (SAM-LSTM) model. This model improves on-chain data price forecasting by combating three key challenges in cryptocurrency price prediction.

First, they use cross-correlation analysis to identify the best on-chain metrics for price prediction. It is crucial first to refine the input variables before price forecasting. They find the most prominent metrics to be active addresses, transaction count, miner revenue, and exchange flows. They filter out low-impact metrics, which reduces noise and overfitting. This approach leads to more accurate predictions.

Second, Kim et al. (2022) introduce a segmentation technique called Change Point Detection (CPD). This technique divides price data into different market phases to avoid reacting to sudden market fluctuations and instead focuses on stable price trend patterns. Previous findings support this idea. Lahmiri and Bekiros (2019) show that cryptocurrency price series exhibit both chaotic and convergent regimes and that deep learning models like LSTM can successfully learn and forecast patterns in these nonlinear environments.

Third, the SAM-LSTM model improves the traditional LSTM model because of its self-attention mechanism. Standard LSTMs give equal importance to all time steps, but the SAM-LSTM model can adjust its focus to the most relevant past events. This is a significant improvement as it can detect long-term dependencies in on-chain market behavior. Kim et al. (2022) find that their model improves Mean Absolute Error (MAE) by 33,79% compared to traditional LSTMs. This further reinforces the benefit of using deep learning in cryptocurrency price predictions.

Kim et al. (2022) show that combining on-chain data and advanced ML models significantly improves predictive accuracy. This also complements the findings by Jagannath et al. (2021), who used LSTMs for Ethereum price forecasting, illustrating that different cryptocurrencies have similar predictive patterns with on-chain data. The SAM-LSTM model is tested across 16 different on-chain metric groups and outperforms traditional LSTM models in 15 of them. They also find that the best results are achieved when combining “price”, “adoption”, and “market” variables, indicating a substantial benefit in combining metrics to increase accuracy. It shows that multivariate on-chain metrics provide better predictive power than univariate approaches. Also, their attention mechanism further improved the model as it could focus on the most relevant historical windows, which could be especially beneficial when predicting regime shifts.

Further supporting these results, King et al. (2024) find that on-chain metrics, such as Bitcoin Difficulty (DIFF), Cost Per Transaction (CPTRA), and Active Addresses (NADDU), correlate well with Bitcoin’s price. These real-time metrics provide valuable inputs for machine learning models trying to predict Bitcoin’s future price movements. Their study goes beyond traditional on-chain metrics by creating more advanced blockchain-based trading signals using “blockchain ribbons”. This method uses moving averages of on-chain metrics like mining difficulty and transaction costs. Their trading strategy demonstrates that specific on-chain metrics are more suited for identifying market accumulation phases than short-term price corrections.

King et al. (2024) use machine learning and algorithmic trading simulations to measure the predictive capability of different on-chain metrics. They find that especially DIFF and NADDU excel in long-term market trends, but CPTRA is more valuable in identifying overheated market conditions. Models using CPTRA-based signals overperform those that use only historical price data. This shows that blockchain transaction costs can give helpful information about the attention and interest of investors toward cryptocurrencies and speculation. They further strengthen the notion that ML models, especially Random Forest (RF) and LSTM networks, work well when trained on structured on-chain data. Both RF and LSTM networks have unique capabilities regarding cryptocurrency price predictions. RF is more capable of identifying key predictors and handling complex interactions between different metrics, whereas LSTM excels at capturing sequential dependencies and high-frequency price fluctuations.

A key takeaway from their study is that on-chain metrics are successful and reliable in making long-term market predictions but less effective in capturing short-term price fluctuations. Also, King et al. (2024) find that long signals are more accurate and profitable than short signals, indicating that additional market signals might be required to increase the accuracy of the trading strategy in corrective market phases. This also aligns with previous research by Kim et al. (2022) and Casella and Paletto (2023, who suggest combining many different sentiment and liquidity data to increase the ML model accuracy. In addition, King et al. (2024) recommend including off-chain data, such as news sentiment and social media analysis. Including off-chain metrics could further improve the accuracy of on-chain metrics, especially in short-term forecasting.

Cryptocurrency forecasting is not just about predicting price movements but also about identifying broader market cycles. Casella and Paletto (2023) demonstrate how on-chain metrics can be used to predict bull and bear market phases. This can offer valuable insights for long-term asset allocation. Many studies have focused on short-term price predictions, but Casella and Paletto (2023) emphasize that predicting longer trends is essential in understanding broader market cycles.

They use deep learning models, such as LSTMs, CNNs, and WaveNet architectures, to forecast on-chain data itself, not just price movements. They find many on-chain metrics that correlate strongly with different market regimes. These include New Addresses, Active Addresses, Block Height, Hash Rate, and Spent Output Profit Ratio (SOPR). New Addresses and Active Addresses indicate how much the network is growing and how active the participants are, Block Height and Hash Rate relate to the blockchain's security and mining activity, and SOPR tells whether investors are selling their coins at a profit or a loss.

Casella and Paletto (2023) find that training ML models with historical on-chain trends enables them to predict whether the market is entering a bullish or bearish phase. Their findings align with King et al. (2024), who used advanced on-chain metrics for market timing. Casella and Paletto (2023) highlight that, as on-chain data is recorded as a time series, it works well in forecasting with DL models. Their model can predict bull and bear market transitions by predicting future values of on-chain metrics like active addresses and SOPR. Their results indicate that not just the price but the directional trends in on-chain metrics themselves can be used to predict market regime shifts.

The effectiveness of machine learning models in price prediction varies depending on the frequency of the data used. Chen, Li, and Sun (2020) show that statistical methods outperform ML models in daily Bitcoin price predictions, but with high-frequency data (5-minute intervals), ML models like Bayesian Neural Networks (BNN) and Support Vector Regression (SVR) came on top. This illustrates the strength of ML models in capturing rapid price movements and volatility, which are characteristics of cryptocurrencies.

Jang and Lee (2017) were among the first to show how Bayesian Neural Networks (BNN) could outperform traditional methods in directional Bitcoin prediction. Even though their model struggled with sudden spikes, it laid the necessary groundwork for BNN in cryptocurrency price prediction. Their work also helped to inspire subsequent research, especially for better-suited models for rapid regime changes.

Jang and Lee (2017) use on-chain data and BNN to predict Bitcoin prices and find that on-chain data is successful at directional prediction. Their BNN model achieves a mean absolute percentage error (MAPE) of 1% when predicting Bitcoin log prices. This is a significant outperformance compared to the SVR and linear models. Their model was also consistent in static and rolling windows, confirming that on-chain metrics provide highly accurate price predictions even in volatile conditions.

Their analysis emphasizes the metrics directly tied to Bitcoin's supply and demand dynamics. Their findings provide empirical evidence that on-chain activity, rather than macro-financial variables, primarily drives recent price volatility. While the paper does not explicitly discuss sentiment, this behavioral dimension can be inferred from their reliance on network-level metrics that reflect participant activity and economic incentives.

In addition to the price prediction capabilities, on-chain data also explained the directional changes in volatility, but the accuracy decreased in sudden changes in volatility. Like Chen, Li, and Sun (2020), Jang and Lee (2017) show that ML models are better suited for price prediction than statistical methods like linear regression. However, their study indicates that BNN is even more successful in price prediction than SVR.

In conclusion, the research indicates that combining on-chain metrics with ML models can be beneficial. This is true especially for models capable of capturing temporal and structural dependencies, which can significantly increase prediction accuracy. These findings support the central thesis and, combined with the results from the case studies, show the significant value of on-chain metrics in cryptocurrency price prediction.

3.2.1 Synthesis of Machine Learning Studies on On-Chain Price Prediction

Although different in scope and methodology, these studies clearly explain how on-chain metrics interact with cryptocurrency prices and investor sentiment. The following table looks at five significant studies central to this thesis. The table asses the studies based on ML technique, on-chain metrics, and whether their findings support either or both hypotheses outlined earlier in this thesis: H1, which proposes that on-chain metrics work as a sentiment indicator, and H2, which believes in statistically significant relationship between price and on-chain metrics.

Table 6. Summary of Machine Learning Studies.

Study	On-Chain Metrics Used	H1 Supported (Sentiment)	H2 Supported (Price Prediction)	Key Findings
King et al. (2024)	DIFF, CPTRA, NADDU, AdCPTRA	Yes	Yes	Metrics tied to investor attention; Random Forest and LSTM compared
Kim et al. (2022)	Active addresses, tx count, miner flows	Indirect	Yes	Segmentation and attention improved accuracy
Casella & Paletto (2023)	SOPR, hash rate, new addresses, block height	Yes	Yes	Forecasted regime shifts; explicitly mention sentiment
Jang & Lee (2017)	Blockchain metrics (e.g. miner revenue, tx volume)	Indirect	Yes	BNN had strong directional predictions; volatility challenges
Jagannath et al. (2021)	Ethereum metrics: active addresses, miner revenue, gas	Indirect	Yes	LSTM outperformed traditional models using ETH data

After reviewing all five studies, it becomes clear that a strong academic consensus supports the second hypothesis. They all concluded that on-chain metrics have predictive value in forecasting price movements in cryptocurrency markets. This is true across different model types like LSTMs, BNNs, Random Forests, and CNNs. For example, Kim et al. (2022) and Casella and Paletto (2023) find that DL models can detect significant trends in on-chain metrics that precede market shifts. In addition, Jang and Lee (2017) prove that BNNs with blockchain data can predict directional cryptocurrency price movements with considerable accuracy. The predictive capabilities of these models, especially when trained on structured or high-frequency on-chain data, indicate a systematic and quantifiable relationship between on-chain metrics and price, validating the second hypothesis.

In contrast, only two reviewed studies support the first hypothesis, which argues that on-chain metrics reflect investor sentiment. King et al. (2024) and Casella and Paletto (2023) directly support the hypothesis by emphasizing that on-chain activity reflects investor sentiment. For example, King et al. (2024) state that their results support the Adaptive Market Hypothesis (AMH), which argues that sentiment affects investor behavior. Casella and Paletto (2023) point out that on-chain analysis is used to determine market trends and sentiment.

Other studies indirectly support the H1 by including behaviorally driven metrics like active addresses and exchange flows. Studies by Jagannath et al. (2021), Kim et al. (2022), and Jang and Lee (2017) focus more on improving price prediction accuracy rather than arguing whether the predictive power of on-chain data comes from sentiment. None of these studies disproves the first hypothesis; instead, they do not take any stance on that.

The variety of different on-chain metrics used across the studies strengthens the case for both hypotheses. Many metrics have behavioral components, like active addresses, gas fees, miner revenue, exchange flows, and difficulty. For example, high gas fees indicate strong demand and increased greed, while exchange flows show the collective interest in buying or selling assets. As these metrics also improve the accuracy of predictive models, they, in turn, strengthen the case for sentiment's importance in forecasting. Furthermore, studies that use segmentation techniques, like the study by Kim et al. (2022), who use change point detection to identify market regimes, show indirectly that the emotional state of the market is not static but evolves in identifiable phases. This indicates that on-chain metrics can serve as a bias-free proxy for these changing emotional states.

This connection between on-chain data and sentiment becomes even more convincing when we consider the nature of cryptocurrencies. Most, if not all, cryptocurrencies lack any intrinsic value, unlike equities with cash flows or commodities with physical utility. This is true even with Bitcoin, but especially with coins like Dogecoin, which was initially created as a joke and achieved a value of over 70 billion dollars despite that. This underscores how the valuation of cryptocurrencies is mainly driven by sentiment. If price is driven by sentiment and on-chain data tracks the footprint of those market participants, then it means on-chain data is not merely a proxy of sentiment but a direct measure of it.

In summary, these studies strongly support the second hypothesis and moderately to strongly support the first. While not all studies directly investigate sentiment's role in on-chain data, the relationship between the two is still evident in many studies. It becomes even more apparent when reviewed in the context of cryptocurrencies and their sentiment-driven nature. These studies show the dual role of on-chain data: it works as a predictive tool and as a sentiment measure. The ability of ML models trained with on-chain data to provide insights that can be acted upon makes this field academically compelling and practically valuable. Since there is a limited body of evidence for this topic, the five studies provide a robust foundation for the ML component of this thesis. They all show that using on-chain metrics to measure sentiment and predict prices is possible and meaningful in the current academic and market landscape.

4 Practical and academic implications

Research on on-chain data has significant implications both in practice and in the academic field of finance. This section outlines how this thesis's findings can influence trading strategies and academic research directions.

One apparent practical implication is the ability to predict the cryptocurrency market more accurately. Using metrics like the Net Unrealized Profit/Loss (NUPL), MVRV Z-score, and Cumulative Value Days Destroyed (CVDD), traders can enhance their market analysis and outperform the markets by identifying different market phases, such as accumulation, euphoria, or capitulation. These metrics essentially detect sentiment extremes that coincide with price peaks and bottoms. Traders and long-term investors can use these metrics to time major price reversals and enhance their returns.

Also, the on-chain metrics can help traders mitigate risks, as traditional risk models often fall short in volatile markets. These metrics give investors a complementary toolset that can help them monitor signals like overheating and widespread panic. It is not only about maximizing returns but also about minimizing risks. Metrics like MVRV Z-score and the NUPL ratio increase as the overall risks in the market increase. Depending on the individual’s risk appetite, they could slowly decrease their risk as these indicators move closer to extremes.

For regulators, the blockchain’s transparency and vast amounts of on-chain data provide a detailed view of market sentiment and systematic behavioral shifts in real time. This allows regulators to be more proactive, as on-chain metrics can serve as early warning signals of overheating markets. By better understanding the different market dynamics, they can act more successfully in times of systemic risk. Better-tailored regulatory oversight can lead to a more stable and reliable market environment.

In an academic context, this thesis adds to the literature of crypto-financial research through the lens of behavioral price formation. It builds around this unique view where sentiment is not some separate psychological construct but is directly measurable from blockchain transactions. This idea both challenges and tries to improve upon traditional asset pricing theories. In a market without apparent intrinsic value, it is reasonable to hypothesize that price is almost entirely a construct of collective sentiment. If that is true, then measuring sentiment directly through on-chain data should be possible. Also, if price is a construct of sentiment, then the best input for an asset pricing model is sentiment itself. From this perspective, we end up with this thesis's two hypotheses: that on-chain data is a direct measure of sentiment and that it can be used to predict cryptocurrency prices.

Finally, this thesis raises broader theoretical questions about the validity of the Efficient Market Hypothesis (EMH). If cryptocurrencies form predictable patterns that can be utilized financially, then the assumption of information efficiency does not hold. Moreover, if predictable patterns exist in cryptocurrencies, they might also exist in the traditional markets. This realization opens the door for cross-market behavioral studies. In fact, it could be argued that without any intrinsic value, the efficient price for most cryptocurrencies should be zero. Thus, any sustained price above zero in crypto markets is direct empirical evidence against the EMH. For this reason, blockchain offers a unique environment to observe how collective sentiment, rather than intrinsic value, drives prices. This raises new questions about how markets function not only in the absence of valuation anchors but also when such anchors exist.

4.1 Limitations and Future Research Directions

Despite the promising results presented in this thesis, potential limitations that might hinder the generalizability and robustness of the findings must also be considered. First, most of this thesis's academic literature and case studies focus only on Bitcoin. Even though Bitcoin is the most stable, liquid, and prominent cryptocurrency, it might not fully represent the entire market, as different cryptocurrencies could behave differently due to many variations in user bases, tokenomics, or market structures. This means that the conclusions made in this thesis might not be universally applicable across the whole crypto ecosystem. However, empirical correlation analysis suggests that some cryptocurrencies strongly correlate with Bitcoin. This was especially true with Ethereum, which had over 0,90 Pearson correlation with Bitcoin. XRP and ADA also had moderately strong correlations with Bitcoin (0,57 and 0,61, respectively), which implies that these results are to some extent generalizable to the entire market, even though the research was mainly focused on Bitcoin.

Another limitation is how different metrics are chosen and interpreted. While indicators like NUPL ratio, MVRV Z-score, and CVDD have successfully predicted market turning points, they are not entirely objective. This field lacks methodologies that are universally accepted for defining signal values. Decisions made by the inventors of different metrics might have been influenced by hindsight bias, reducing the metrics' trustworthiness. Also, new on-chain indicators are frequently created, meaning the ones analyzed here might become obsolete when the market matures and evolves.

In addition, the statistical testing used in case studies has some limitations and could be further enhanced. The thesis does not apply correlations for multiple hypothesis testing, such as a Bonferroni adjustment. For this reason, the risk of false positives is much higher. Also, while the Monte Carlo simulations used random entry testing, more formal validation methods could have been used to increase the robustness of the results. One option could have been walk-forward out-of-sample testing. Future research could address these limitations to enhance the robustness of the results.

The machine learning studies discussed in this thesis also have limitations. While recent studies have proven the benefits of ML models like random forests and neural networks, they still rely on data quality, parameter tuning, and assumptions about stationarity and relevance of input metrics. This is also true with on-chain metrics, even though they have the advantage of being transparent, accessible, and direct behavioral measures. Additionally, most of these models were tested on Bitcoin, making cross-validation across tokens somewhat tricky.

Future research could explore the relationship between on-chain data and prices across different cryptocurrencies, like Ethereum, Solana, and XRP. This would improve our understanding of the whole ecosystem and not just Bitcoin. Additionally, applying unsupervised AI methods for clustering and sentiment classification would be beneficial, as that could lead to discoveries in investor behavior.

Finally, combining behavioral finance, data science, and crypto research could create new sentiment-based asset pricing models in digital assets. These new models could not only improve forecasting accuracy but also redefine valuation models in digital and decentralized markets. Also, discoveries in more sentiment-driven markets like cryptocurrencies could help develop similar sentiment-based asset pricing models in traditional finance. This cross-market perspective could influence how we think about sentiment and its role in asset pricing models beyond the crypto space.

5 Conclusion

This thesis explored whether on-chain data could be used as a sentiment measure and to predict cryptocurrency prices. Through historical case studies of on-chain metrics, custom-built trading strategies, Monte Carlo simulations, and machine learning research it strongly supports both hypotheses. The case studies on Net Unrealized Profit/Loss (NUPL), MVRV Z-score, and Cumulative Value Days Destroyed (CVDD) proved how on-chain metrics could capture shifts in sentiment and predict key market turning points. The NUPL ratio and the MVRV Z-score constantly aligned with market tops and bottoms, while the CVDD could almost perfectly predict every significant bottom.

To validate these observations, we ran Monte Carlo simulations that showed CVDD's ability to outperform 97% to 100% of randomly timed entries in cumulative log returns. We also used rule-based trading strategies to analyze NUPL ratio and MVRV, showing an outperformance against the buy-and-hold benchmark across all cycles in risk-adjusted returns, annualized log returns, and cumulative log returns.

To complement this, the machine learning research demonstrated the predictive accuracy of on-chain metrics, especially when using models like LSTM, Random Forest, and BNN. These models can detect patterns that might be missed by the human eye and, at the same time, strengthen the notion that on-chain data contains behavior that can be used to predict prices. The research suggests that sentiment is not some separate entity but a measurable force embedded directly in on-chain data.

While this thesis primarily focuses on Bitcoin, its findings are still crucial for broader financial analysis, trading strategy, and behavioral finance. Blockchain's transparent and tamper-proof nature offers data for sentiment analysis that is often not found in traditional markets, where such data is usually delayed, distorted, or unavailable. As cryptocurrencies become more popular every year, different methods of analyzing the behavior of market participants become more essential for traders, regulators, and academics.

Although the evidence provided in this thesis strongly supports on-chain price prediction, it needs to be acknowledged that even these methods are nowhere near perfect. As markets evolve, these metrics could lose effectiveness, as what works today might not work tomorrow. Nonetheless, the motivation of this thesis was not to predict price moves perfectly but to gain a better understanding of behavioral patterns and inefficiencies that traditional markets often fail to detect.

Ultimately, this thesis's findings contribute to the growing literature suggesting that investor sentiment can be measured and utilized, especially in markets like cryptocurrencies, where sentiment remains the key driving force. The insights gained through on-chain analysis may change how we think about sentiment, efficiency, and valuation in 21^{st}-century markets.

Future research could build on this by applying similar analysis to a wider range of cryptocurrencies and developing new models based on on-chain behavior. Exploring these dynamics in more detail could help develop and refine sentiment-based valuation models and improve cross-market understanding of sentiment and behavioral finance principles.

References

Alvarez, F., Argente, D., & Van Patten, D. (2023). Are cryptocurrencies currencies? Bitcoin as legal tender in El Salvador. Science, 382(6677), eadd2844. https://doi.org/10.1126/science.add2844

Baker, M., & Wurgler, J. (2007). Investor sentiment in the stock market. Journal of Economic Perspectives. https://pubs.aeaweb.org/doi/pdfplus/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. Economic Modelling, 64, 74-81. https://doi.org/10.1016/j.econmod.2017.03.019

Barberis, N., & Huang, M. (2001). Mental accounting, loss aversion, and individual stock returns (NBER Working Paper No. w8190). National Bureau of Economic Research. https://ssrn.com/abstract=264445

Bernard, V. L., & Thomas, J. K. (1989). Post-earnings-announcement drift: Delayed price response or risk premium? Journal of Accounting Research, 27, 1–36. https://doi.org/10.2307/2491062

Bitcoin Magazine Pro. (n.d.). Bitcoin: MVRV Z-Score. Retrieved February 24, 2025, from https://www.bitcoinmagazinepro.com/charts/mvrv-zscore/

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. Journal of Computational Science. https://www.sciencedirect.com/science/article/pii/S187775031100007X?via%3Dihub

Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., & Felten, E. W. (2015). *SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies*. IEEE Security & Privacy. https://eprint.iacr.org/2015/261.pdf

Brown, G. W., & Cliff, M. T. (2005). Investor Sentiment and Asset Valuation. *The Journal of Business*, *78*(2), 405–440. https://doi.org/10.1086/427633

Burggraf, T., Huynh, T.L.D., Rudolf, M. & Wang, M. (2021), "Do FEARS drive Bitcoin?", *Review of Behavioral Finance*, Vol. 13 No. 3, pp. 229-258. https://doi.org/10.1108/RBF-11-2019-0161

Brunnermeier, M. K., & Pedersen, L. H. (2005). Market liquidity and funding liquidity (NYU Working Paper No. SC-AM-05-06). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1294630

Casella, B. and Paletto, L. (2023). Predicting Cryptocurrencies Market Phases through On-Chain Data Long-Term Forecasting. https://iris.unito.it/bitstream/2318/1902652/3/6.%20ICBC23%20-%20PREDICTING%20BTC.pdf

Chen, Z., Li, C., & Sun, W. (2020). Bitcoin price prediction using machine learning: An approach to sample dimension engineering. Journal of Econometrics, 219(2), 486–500. https://doi.org/10.1016/j.cam.2019.112395

Cryptopedia Staff. (n.d.). The DAO hack and its aftermath. Cryptopedia by Gemini. Retrieved April 28, 2025, from https://www.gemini.com/cryptopedia/the-dao-hack-makerdao

Dale, B. (2024). Axios Crypto Newsletter. Axios. Retrieved from https://www.axios.com/newsletters/axios-crypto-153c0380-ee00-11ef-b86d-6d7d0bcc8610

Da, Z., Engelberg, J., & Gao, P. (2011). In Search of Attention. *The Journal of Finance*, *66*(5), 1461–1499. http://www.jstor.org/stable/41305167

De Long, J. Bradford, Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise Trader Risk in Financial Markets. Journal of Political Economy. https://www.jstor.org/stable/2937765?read-now=1&seq=9#page_scan_tab_contents

Fama, E. F., Fisher, L., Jensen, M. C., & Roll, R. (1969). The adjustment of stock prices to new information.International Economic Review, 10(1), 1-21. Retrieved from https://www.bu.edu/econ/files/2011/01/FFJR.pdf

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. The Journal of Finance, 25(2), 383–417. https://doi.org/10.2307/2325486

Hu, J., Sui, Y., & Ma, F. (2021). The measurement method of investor sentiment and its relationship with stock market. https://www.hindawi.com/journals/cin/2021/6672677/

Investing.com. (n.d.). Bitcoin historical data. Retrieved April 12, 2025, from https://www.investing.com/crypto/bitcoin/historical-data

Jaffe, J. F. (1974). Special information and insider trading. The Journal of Business, 47(3), 410–428. https://doi.org/10.1086/295655

Jagannath, N. et al. (2021). An On-Chain Analysis-Based Approach to Predict Ethereum Prices. In IEEE Access, vol. 9, pp. 167972-167989. https://ieeexplore.ieee.org/document/9650873?denied=

Kim, G., Shin, D.-H., Choi, J. G., & Lim, S. (2022). A Deep Learning-Based Cryptocurrency Price Prediction Model That Uses On-Chain Data. In IEEE Access, vol. 10, pp. 56232-56248. https://ieeexplore.ieee.org/document/9781377?denied=

Jang, H., & Lee, J. (2017). An empirical study on modeling and prediction of bitcoin prices with Bayesian neural networks based on blockchain information. IEEE Access, 6, 5427-5437. https://ieeexplore.ieee.org/document/8125674

Kendall, M. G., & Hill, A. B. (1953). The analysis of economic time-series—Part I: Prices. Journal of the Royal Statistical Society: Series A (General), 116(1), 11-34. Retrieved from http://ww.e-m-h.org/KeHi53.pdf

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

Kristoufek, L. (2013). BitCoin meets Google Trends and Wikipedia: Quantifying the relationship between phenomena of the Internet era. Scientific Reports. https://www.researchgate.net/publication/259154514_BitCoin_meets_Google_Trends_and_Wikipedia_Quantifying_the_relationship_between_phenomena_of_the_Internet_era

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

Lee, C. M. C., Shleifer, A., & Thaler, R. H. (1991). Investor Sentiment and the Closed-End Fund Puzzle. *The Journal of Finance*, *46*(1), 75–109. https://doi.org/10.2307/2328690

Łęt, B., Sobański, J., Świder, P., & Włosik, Ł. (2022). Cryptocurrency market timing with on-chain indicators: Evidence from active strategies. Econstor. https://www.econstor.eu/handle/10419/309769

Lo, A. W., & MacKinlay, A. C. (1987). Stock market prices do not follow random walks: Evidence from a simple specification test. Rodney L. White Center for Financial Research, University of Pennsylvania. Retrieved from https://rodneywhitecenter.wharton.upenn.edu/wp-content/uploads/2014/04/8705.pdf

Murray, C. (2024, December 5). 'Hawk Tuah' creator Haliey Welch criticized for chaotic memecoin launch in latest bizarre internet stunt. Forbes. Retrieved from https://www.forbes.com/sites/conormurray/2024/12/05/hawk-tuah-creator-haliey-welch-criticized-for-chaotic-memecoin-launch-in-latest-bizarre-internet-stunt/

Naeem, M. A., Mbarki, I., & Shahzad, S. J. H. (2021). Predictive role of online investor sentiment for cryptocurrency market: Evidence from happiness and fears. International Review of Economics & Finance. https://www.sciencedirect.com/science/article/pii/S1059056021000083?casa_token=KnWGEXDXGF8AAAAA:Q2wDUiEnisiy_8lw5TnRx2Aas6U59M6DBtctRM-nQSLyRsmia-UPZTyxJOotboumEjRNMlYGfPw

Otaola, J., & Madry, K. (2025, February 16). Argentina's opposition threatens impeachment trial after Milei touts crypto coin. Reuters. Retrieved from https://www.reuters.com/world/americas/argentinas-opposition-threatens-impeachment-trial-after-milei-touts-crypto-coin-2025-02-16/

Q.ai. (2022, September 20). What really happened to LUNA crypto? Forbes. Retrieved from https://www.forbes.com/sites/qai/2022/09/20/what-really-happened-to-luna-crypto/

Sarmah, S. S. (2018). Understanding blockchain technology. Computer Science and Engineering, 8(2), 23–29. https://doi.org/10.5923/j.computer.20180802.02

Seyhun, H. N. (1986). Insiders' profits, costs of trading, and market efficiency. Journal of Financial Economics, 16(2), 189–212. https://doi.org/10.1016/0304-405X(86)90060-7

Shefrin, H., & Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. The Journal of Finance, 40(3), 777–790. Retrieved from https://people.bath.ac.uk/mnsrf/Teaching%202011/Shefrin-Statman-85.pdf

Tamas Blummer, Tuur Demeester, and Michiel Lescrauwaet. (2019). Bitcoin: Net Unrealized Profit/Loss (NUPL). In Lookintobitcoin.com. Visited on 10.3.2024. https://www.lookintobitcoin.com/charts/relative-unrealized-profit--loss/

Tripathi, A., Ahad, M. A. R., & Casalino, G. (2023). A systematic literature review of blockchain scalability, security and consensus mechanisms. Blockchain: Research and Applications, 4(2), 100106. https://doi.org/10.1016/j.dajour.2023.100344

Von Neumann, J., & Morgenstern, O. (1944). Theory of games and economic behavior. Princeton University Press. Retrieved from https://archive.org/details/in.ernet.dli.2015.215284/page/n25/mode/2up

Wang, W., Su, C., & Duxbury, D. (2021). Investor sentiment and stock returns: Global evidence. Journal of Empirical Finance. https://www.sciencedirect.com/science/article/abs/pii/S092753982100058X

Wieczner, J. (2018, January 17). Bitcoin crashes to below $10,000 as cryptocurrency scams scare investors. Fortune. Archived on January 27, 2018. Retrieved February 10, 2018, from https://web.archive.org/web/20180127005114/http://fortune.com/2018/01/17/bitcoin-bitconnect-price-scam/

Woo, W. (n.d.). Bitcoin: CVDD (Cumulative Value Days Destroyed). Bitcoin Magazine Pro. Retrieved April 4, 2025, from https://www.bitcoinmagazinepro.com/charts/cvdd/