Eivind Homb Løkke
Erik Eriksen Grøntvedt
Frida Strandkleiv Øverland

Predicting Bitcoin returns using time series analysis and deep learning

Generating a trading algorithm based on LSTM and Markov switching models

Bachelor's thesis in Business Administration
Supervisor: Mike Denis Blomsø-Becker
April 2024

Bachelor's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Economics and Management
NTNU Business School

NTNU
Norwegian University of
Science and Technology

Eivind Homb Løkke
Erik Eriksen Grøntvedt
Frida Strandkleiv Øverland

Predicting Bitcoin returns using time series analysis and deep learning

Generating a trading algorithm based on LSTM and Markov switching models

Bachelor's thesis in Business Administration
Supervisor: Mike Denis Blomsø-Becker
April 2024

Norwegian University of Science and Technology
Faculty of Economics and Management
NTNU Business School

NTNU
Norwegian University of
Science and Technology

Preface

An exciting, but challenging time comes to an end with this bachelor’s thesis as we finish up our bachelor’s degree here in Trondheim. We are three Business administration students who have had the honor of studying at NTNU Business School. Throughout this bachelor program we have learnt a lot from various fields, from leadership to statistics. Choosing business analytics as our major we have had the opportunity to combine these insights in business, economics and management with data science. This is also reflected in this thesis as we combine data science with finance, more specifically investing in cryptocurrency.

We will soon go our separate ways, and we look forward to pursuing master’s degrees in which we can continue to build on the knowledge acquired throughout the last 3 years.

We would also like to give a special thanks to our supervisor Mike Denis Blomsø-Becker. He has helped us a lot throughout this last school year with this thesis, but also as our lecturer in the core Business Analytics courses here at NTNU Business School. We are grateful for the time you have spent teaching us, giving us feedback and offering support.

- Frida Strandkleiv Øverland, Erik Eriksen Grøntvedt, Eivind Homb Løkke

Abstract

In this Bachelor thesis we develop algorithms based on predicted returns for trading the cryptocurrency Bitcoin, challenging the Efficient Market Hypothesis (EMH). We achieve this by creating models for sequence learning like Markov Switching models and a Long short-term memory (LSTM) model. These models rely primarily on historical data and a Crypto Fear and Greed Index (CFGI) to predict Bitcoins returns from the year 2020 to 2023. The predicted returns from our models are then used to create buy and sell signals for Bitcoin which we use to generate different trading algorithms.

According to EMH we should not be able to create such well-performing trading algorithms since the markets are efficient and the available information would already be reflected in the Bitcoin price. When comparing the algorithms to a buy and hold strategy, we observe promising results from some of the models. We have also compared the different algorithms to each other by looking at return and risk and we see that the Markov Switching model based on CFGI and the LSTM model perform better than the Markov Switching model based on moving averages. Although our trading algorithms show potential to outperform the Bitcoin market one should keep in mind that they are based on a limited time frame, and results may vary depending on the observed period.

Sammendrag

I denne Bacheloroppgaven utvikler vi algoritmer basert på predikert avkastning for handel av kryptovalutaen Bitcoin, og med dette utfordrer hypotesen om effisiente markeder. Vi oppnår dette ved å lage modeller for sekvenslæring som Markov Switching-modeller og en Long short-term memory (LSTM) modell. Disse modellene er først og fremst avhengige av historiske data og en Crypto Fear and Greed Index (CFGI) for å forutsi Bitcoins avkastning fra året 2020 til 2023. Den predikerte avkastningen fra modellene våre brukes deretter til å lage kjøps- og salgssignaler for Bitcoin som vi bruker til å generere forskjellige handelsalgoritmer.

Ifølge hypotesen om effisiente markeder burde vi ikke klare å lage slike velfungerende handelsalgoritmer siden markedene er effisiente og den tilgjengelige informasjonen allerede vil reflekteres i Bitcoin-prisen. Når vi sammenligner algoritmene med en kjøp og hold-strategi, får vi lovende resultater fra noen av modellene. Vi sammenligner også de ulike algoritmene med hverandre ved å se på avkastning og risiko. Her ser vi at Markov Switching-modellen basert på CFGI og LSTM-modellen presterer bedre enn Markov Switching-modellen basert på glidende gjennomsnitt. Selv om handelsalgoritmene våre viser potensial til å overgå Bitcoin-markedet, bør man huske på at de er basert på en begrenset tidsramme, og resultatene kan variere avhengig av den observerte perioden.

Table of contents

Table of Figures

1. Introduction

In recent years, the rapid development of new, mainly digital securities and asset classes has created an enormous interest for speculation among retail and institutional investors alike. Cryptocurrencies, which are highly speculative assets, are one of the most volatile asset classes that currently exist. Cryptocurrencies are characterized by major fluctuations in price based on world events, media coverage, public opinion, etc. Cryptocurrencies differ from traditional securities such as stocks, mainly due to the lack of tangible value. Owning a share of Apple Inc. will entitle you to future dividend payments, as well as capital gains from price increases. Your position in Apple Inc. may be affected by news about production or sales volume of iPhones, as well as quarterly and yearly earnings reports which indicate the financial health and development of the company. Cryptocurrencies on the other hand, are not tied to such tangible metrics. You are not entitled to some future payments, such as with stocks or bonds, and the security has no inherent value, other than what the market agrees upon. In this way, cryptocurrencies may resemble other currencies, such as USD or EUR. However, it is debatable whether cryptocurrencies fulfill the three properties that modern currencies need (McLeay, et al., 2014). Mainly, the role of money as a predictable store of value is contesting cryptocurrencies' position as legal tender in most countries.

This puts cryptocurrencies in a unique, but difficult position. It is not equity, nor is it legal tender, but it exhibits properties of both. How then, can you make accurate assumptions and predictions about future price developments? This topic has been of increasing interest amongst market participants, researchers, and academics. Traditional methods for valuating stocks and bonds such as DCF's are insufficient, as are variables used to predict currency values. Therefore, we look to statistical analysis, with emphasis on time series analysis as a means of predicting future prices of Bitcoin, the first and most well-known cryptocurrency (Ashford & Adams, 2022). In this thesis we will try to acquire insights and knowledge around the following question: Can you reliably generate excess returns by trading Bitcoin based on historical data?

This thesis is structured as follows:

Chapter 1: Introduction

The first chapter gives the reader a short introduction to the thesis's topic and structure.

Chapter 2: Literature Review

The second chapter contains earlier studies which explores the same topics our thesis is based upon. We look at how market sentiment impact Bitcoin's predictability, and how different machine learning models are used to predict its price and volatility.

Chapter 3: Theory

The third chapter explains economic theory relevant to our bachelor thesis. This includes both the efficient market hypothesis, the fear and greed index and the bitcoin's ever-changing role in the financial market.

Chapter 4: Data and Methodology

The fourth chapter presents the dataset used in the thesis. This includes our data collection, as well it's preparation and formatting. The chapter also details the models utilized in our analysis, and our reasoning for using them.

Chapter 5: Results

The fifth chapter covers the discoveries from our three models. With this, we interpret our findings and compare the models in relation to the thesis objective.

Chapter 6: Conclusion

In our final chapter, we emphasize our result, summarizing and discussing the findings. We also reflect on our experiences and offer recommendations for further research.

2. Literature Review

Despite its recent introduction, there have been numerous attempts at modeling statistical time series models based on market conditions, aiming to outperform the bitcoin index. In our literature review, we seek to summarize and analyze prior research that discusses time series models in predicting bitcoin's returns. Through this effort, we intend to highlight disagreements and contested claims, clarifying the current state of research on bitcoin's financial predictions.

2.1 Bitcoin Predictability and Market Sentiment

Eom et al. (2019) carried out a study to examine how bitcoins statistical characteristics and the predictability of its return and volatility relate to investor sentiment. The research utilized an Autoregressive Model Framework to analyze predictability with the Google Trend Index (GTI) as a proxy for investor sentiment. Additionally, daily bitcoin and euro exchange rates were used as a base to evaluate bitcoin's behavior against traditional assets. The study results indicate thar investor sentiment improves the predictability of bitcoin's volatility and as well as it provides useful insight into bitcoins unique behavior compared to conventual assets.

A study by Mokni et al. (2022) could challenge these beliefs. The study aimed to apply a quantile-based analysis to examine the relationship between bitcoin and investor sentiment, during the COVID-19 pandemic. The research suggests a positive significant positive relationship between investor sentiment and bitcoin's returns, as when investor sentiment increased, the bitcoin returns also tended to increase. However, investor sentiment alone was not able to reliably predict bitcoin's return or volatility. The study differs as it uses The Crypto Fear and Greed Index (CFGI) as a proxy for investor sentiment instead of the GTI. A study by Gaies et al. (2023) also used the CFGI as a measure for investor sentiment. By applying a bootstrap rolling window Granger causality test, their research revealed both negative and positive correlations between market sentiment and bitcoin prices, further providing evidence for a connection between the two.

The predictive power of the CFGI is further explored in the study by Łet et al. (2023). The researchers aimed to simulate an active cryptocurrency investment strategy for the 20 largest non-stable cryptocurrencies, using the CFGI to determine if they could outperform a passive buy-and-hold strategy. The active strategy viewed low values of CFGI as an investment opportunity, while high values were seen as indicators to sell. Performance is measured by how much return is gotten in return for risk taken, using both the Sharpe ratio, which compares profit to risk, and the average profit over time. The study concludes that strategies based on the CFGI performed significantly better compared to passive buy-and-hold strategies, further expanding the understanding that market sentiment can be used as a base for investment strategies. It's important to note that the study examines 20 of the largest non-stable cryptocurrencies, meaning that it does not solely confirm the strategy’s effectiveness for bitcoin alone.

This is instead examined in a study by Huang (2024). Similarly, the study sets out to investigate the CFGI as an indicator for market sentiment and predictive parameter for bitcoin prices. Through machine learning models such as linear regression, polynomial regression, random forest and XGBoost Huang are able to predict bitcoin prices with varying success. The XGBoost model proved to be the best for predicting bitcoin prices, further underscoring the significant correlation with the CFGI. The predictions from the model should be viewed with caution as high MSE values could suggest potential overfitting. The two texts mainly focus on different subjects, investment strategy and price forecasting, respectively.

In discussions about the ability to predict prices to outperform the market, the Efficient Market Hypothesis (EMH) is also often brought up. Bitcoins volatility structure and its price fluctuations are investigated by Çelik (2020). By applying the Fractal Market Hypothesis (FMH) as an investigative instrument, models like R/S, DFA, Periodogram and GPH examine daily bitcoin price behavior from April 2013 to January 2019. The results revealed fractal behavior in the market, indicating inefficiency, thus allowing potential predictability in the price movements. It is worth noting that the study covers a relatively short period of time, which could affect the results. Furthermore, the paper solely focuses on price data, ignoring market sentiment, widely understood as an important factor of bitcoin's price behavior based on the earlier studies covering the CFGI.

2.2 Machine Learning Models and Predicting Bitcoin's Price and Volatility

Ibrahim et al. (2020) utilized various time series models to compare their ability to predict the direction of Bitcoin price movements within a 5-minute timeframe. Notable models in the study included Autoregressive Integrated Moving Average, Random Forest (ARIMA), and Multilayered Perceptron (MLP) Deep Neural Network. Each model aimed to predict if Bitcoin's price would go up or down, with the goal of maximizing the chances of making profitable returns. To assess their accuracy, the models were measured against a momentum strategy. The strategy predicts whether the price will move in the same direction as the previous interval, acting as a benchmark for comparison. The multilayer perceptron deep neural network came out as the most accurate model, achieving a 54% prediction accuracy. The results underscore the potential of machine learning's ability to predict price movements for Bitcoin. However, the study might be subject to limitations. Firstly, the trading data utilized in the study only came for two different crypto exchanges, which may not represent all market dynamics and thus limiting the models' validity. Secondly, the focus short-term price movements may overlook market trends and other factors that could affect Bitcoin's price.

A similar conclusion is reached by Mudassir et al (2020). While Ibrahim et al. focused on predictions within a 5-minute timeframe, Mudassir et al. implemented machine learning-based models for one, seven, thirty and ninety days to forecast Bitcoin prices. The study utilized models such as Artificial Neural Network (ANN), Stacked Artificial Neural Network (SANN), Support Vector Machined (SVM) and Long Short Term Memory (LTSM). Conversely, the performance of the models is here evaluated by Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Absolute Percentage Error (MAPE). While all models performed satisfactorily, the LTSM provided the best overall performance. On average the models scored up to 65% accuracy for next-day predictions, and values close 62% and 64% for the seventh-ninetieth-day forecasts. While the studies pursued different objectives, predicting the direction of the price movement and predicting the actual prices, respectively, we see that incorporating wider periods could indicate a more precise model. Both the studies lack inclusion of external factors. By focusing on data tied to bitcoin's historical prices, external factors such as macroeconomic indicators and market sentiment get overlooked, thus potentially making the models less precise.

McNally et al. (2018) further researched time series data analysis of cryptocurrency. Their study similarly investigated the accuracy with which bitcoin prices can be predicted using machine learning models. Instead of focusing on a short-term perceptive, the study focuses on a broader and more general perspective of price movement. Furthermore, the study is more concise with only implementing Bayesian Optimized Recurrent Neural Network (RNN) and LTSM as predictive models. Despite this, they still experienced success with implementing neural network-based models in predicting bitcoin prices, with their LTSM model also achieving an accuracy if 52%, further highlighting the evidence for the use of machine learning in forecasting bitcoin prices movement.

Chkili (2021) explores models for describing price volatility for bitcoin, aiming to identify the most accurate. Instead of predicting concrete prices, as earlier studies have attempted, the study instead aims to analyze if past price fluctuations have a long-term impact on bitcoin's future volatility. There were two types of models utilized in the study to examine the claim, a Long Memory Model, specifically a Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity (FIGARCH) model, and a Markow Switching GARCH Model. The former model analyzed the data for long memory, while the latter model analyzed the data for regime changes. The models were evaluated using various information criteria and the log-likelihood, where the FIGARCH model exceeded in both. Indicating it being more precise for modeling the volatility, suggested by explicit accounting for long memory properties present in the in the conditional volatility of bitcoin prices. The models primarily analyze historical price data. External factors could influence the volatility without being accounted for in the models. Malladi & Dheeriya (2020) also aimed to analyze the bitcoin's returns and volatility. Instead of looking at past price shifts, the study instead sought to examine the correlation between returns and volatility, and stock indices, gold prices, and fear indicators. The study utilized different time series models such as an Autoregressive-moving-average model with exogenous input model (ARMAX), GARCH model, Vector Autoregression (VAR) model, and Granger causality test. None of the models found a significant correlation between Bitcoin's return and volatility, and the stock market and gold. This could indicate the external factors could have less of an impact on bitcoin' return and volatility than initially believed.

3. Theory

3.1 The Efficient Market Hypothesis

The efficient market hypothesis (EMH) is credited to Nobel Prize winner Eugene Fama, following his publication of "Efficient Capital Markets: A Review of Theory and Empirical Work" in 1970. EMH views the overall market average as virtually impossible to beat. While an investor may achieve short-term gains from a few stocks, maintaining success over longer periods of time is highly unlikely. The theory is based on the assumption that current stock prices represent the accumulation of information that is freely and widely available to all investors (CFI Team, u.d.). As the prices are constantly reflected in the current available information, the stocks will theoretically never be anything but their fair market value, thereby making market analysis and timing strategies mostly ineffective (Baldridge, 2022).

To better illustrate the theory and the concept of efficient markets, the well-known story involving a finance professor and a student who discover a $100 bill on the ground is often brought up. As the student reaches to pick it up, the professor says, "Don't bother - if it were truly a $100 bill, it wouldn't still be there." This line of thought, according to the EMH, also translates to the financial market. Given the market's efficiency, finding undervalued stocks to make quick profits is challenging. Even when prices are volatile or people act irrationally, the market will simply adjust to new information too quickly (Burton, 2003). Essentially, no one can consistently beat the market to make extra profits, just as you can't expect to find free money lying around.

EMH has attracted both support and criticism from both sides. Research by Morningstar found that only U.S. small growth and emerging markets funds were able to outperform index funds, and even then, only about half the time (Jackson, 2023). Another study by Morningstar revealed that less than one-quarter of the most successful active fund managers demonstrated the ability to consistently surpass index fund performance (Greengold, 2023). Examining the research, one might start to question the utility of detailed market analysis and the expertise of well-educated professionals, as their success could merely reflect luck. Conversely, investors such as Seth Klarman, Jim Simons, and Benjamin Graham have consistently been able to do this, and have outperformed the market. Despite EMH suggesting such outcomes should be unattainable, these examples exist (CFI Team, u.d.). EMHs statement that stocks theoretically always have their fair market value is also contested, especially by examples like Tesla's stock price. Despite the downward trajectory more recently, the second quarter earnings of 2023 continued to show that it was one the most overvalued stocks in the market (Trainer, 2023).

3.2 The Fear & Greed Index

When discussing the current consensus of bitcoin in the crypto market, the Crypto Fear & Greed Index (CFGI) is often used as a proxy. The index is mainly used to predict and assess bitcoin's future performance, thereby also acting as a pointer for investments (Tambe & Jain, 2023). It is based on fear and greed indexes commonly prevalent in stock market analyses. As crypto holds no intrinsic value, the CFGI is solely based on past performances and current market opinion (Johnson, 2023).

CFGI works on a scale from 0 to 100, where zero indicates "extreme fear", while 100 measures "extreme greed". Like traditional fear and greed indexes, fearful investors may opt to sell their holdings, leading to a potential decrease in crypto asset prices and creating a strategic buying opportunity for discerning investors. Conversely, greedy investors tend to overvalue the market and thereby also drive prices beyond sustainable levels, eventually leading to a market correction (Jansen & Nikiforov, 2006).

The index is built upon six key factors. The metrics are described below, along with their proportionate effect on the index brackets (Alternative, u.d.):

- Volatility – Current performance to its average over the last 30 and 90 days (25%).
- Volume – Volume over the last 30 and 90 days (25%).
- Social media – Consensuses of the market on social media (15%)
- Surveys – Weekly polls with 2,000-3,000 investors to understand market views (15%).
- Dominance – Market cap share within the entire cryptocurrency market (10%)
- Trends – Search trends pertaining to bitcoin and its associated topics (10%)

3.3 Bitcoin in the financial market

Since the introduction of bitcoin, the pioneering cryptocurrency has remained the most recognized and actively traded among the thousands of others that have emerged since (Forbes, 2024). Created as “an electronic payment system based on cryptographic proof instead of trust” (Ashford & Adams, 2022), bitcoin eliminates the need for traditional intermediaries, embracing decentralization and transparency in the form of Blockchain technology (PwC, u.d.).

This is in stark contrast to traditional currencies like the USD and NOK, which rely on banks for transaction management and regulation. Compared to assets with intrinsic value, such as gold and real estate, bitcoin is viewed as significantly more volatile, entirely driven upon the price another party is willing to offer. As it is mainly dependent by the ever-changing market sentiment, it stands out in the financial landscape (Huang, 2022), exemplified by May 2021 when its price dropped by 30% in a single day (Pratley, 2021).

Despite this, Bitcoin is still an investment that can grant impressive returns. In February 2024 Forbes shared an article where bitcoin was suggested to be the number one cryptocurrency one should buy (Adams & Dammeyer, 2024). BTC has also had historically large returns. If one were to buy a bitcoin in 2014, they would have had to spend about 450 US dollars, in the beginning of 2024 that same bitcoin has a price of more than 40 000 US dollars (Yahoo Finance, 2024).

4. Data and Methodology

In this thesis we will use a dataset containing historical prices of Bitcoin. The historical data used is downloaded from yahoo finance and it is measured in the currency USD (Yahoo Finance, 2024). It has a timespan from September 2014 until 1 of April 2023 and we use the daily close prices. Giving us over 3100 data points to work with. Looking at the figure below we see that there has been a large increase in the BTC price during this time span, there have also been some fluctuations. If we look to the Efficient Market Hypothesis described earlier, we should not be able to create a well-functioning trading algorithm using only historical data.

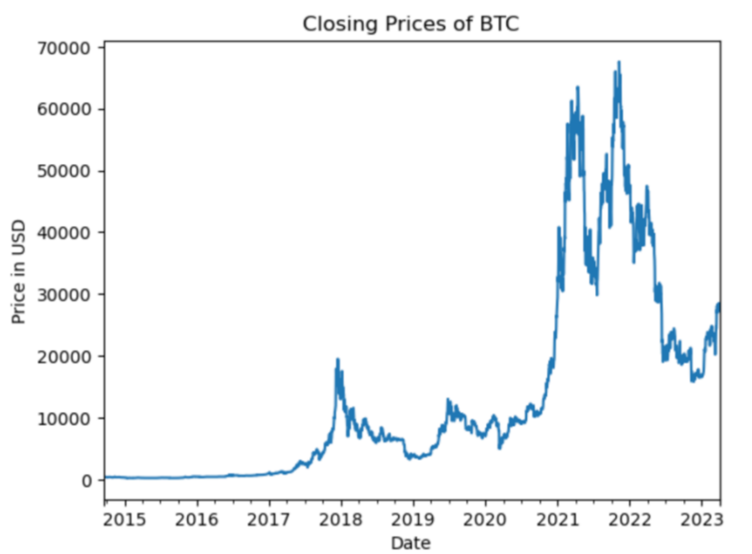

Figure 1: BTC price over time based on data from Yahoo Finance

The rules we set for our trading algorithm will be quite rigid, mainly due to the limited nature of this thesis. Firstly, we do not include any transaction costs in our trading algorithms. This will impact the returns positively and will need to be taken into account when assessing the efficiency of the different models. Some algorithms may yield higher returns, but will have a substantially higher number of trades, thereby decreasing total profits in a real-world scenario. Secondly, we do not allow for a continuous system of holdings, meaning that the algorithm cannot chose to sell parts of its holdings, but rather it must either sell everything and keep it in the bank account, or go all in when buying back in (a binary system). In one model, we also add the ability to short sell bitcoin when this is favorable. Short selling means that we bet on the price of bitcoin going down, as opposed to when we go long and bet that the price will go up. There are several ways to short sell bitcoin. Most of the short methods are by using derivatives (securities that depend on the underlying value of an asset), which means that they are not effective hedging mechanisms. Hedging is what short positions often are used as, but in our case, we do not want to hedge our bets (we only want to go all in). Since Bitcoin is such a volatile security, it is our assumption that adding short positions will greatly benefit the cumulative returns generated. But in other models we do not allow for short sales, which means that the model can only be all in on Bitcoin (1), hold the current position (0) or have all the money in a bank account (-1). We will compare two models with the same parameters with and without short sales. The model starts by holding bitcoin, and thus identifies exit points before entry points.

4.1 Markov Switching Models

Markov Switching is a time series model that model changes in regime in which the data operates. Markov Switching is a non-linear model, and thus differs from models such as ARIMA or VAR. The model uses a maximum-likelihood estimation of the parameters to make assumptions about when a regime shift happens (Hamilton, 1989). By regime shift, we mean structural changes in the behavior of the data, based on complex factors such as market sentiment, outlook or risk appetite. One example of regime shift is for instance going from a low volatility state to a high volatility state. The Markov Switching model captures these shifts as an unobservable state variable based on the immediate past value (Kuan, 2002). The Markov Switching model is especially suitable for economic time series data, such as GDP, market indices or specific assets. In general form a two-state model can be written as:

$$Z_t = \begin{cases} \alpha_0 + \beta z_{t-1} + \varepsilon_{t,} & s_{t=0} \\ \alpha_0 + \alpha_1 + \beta z_{t-1} + \varepsilon_{t,} & s_{t=1} \end{cases}$$

Where $|\beta| < 1$ and ε_t are independent and identically distributed variables with mean zero and variance σ_ε^2 (Kuan, 2002). This is an AR(1) process, with mean $\frac{\alpha_0}{(1-\beta)}$ when $s_{t=0}$, and switches to the secondary AR(1) process with mean $\frac{(\alpha_0+\alpha_1)}{(1-\beta)}$ when $s_{t=1}$. In other words, depending on the "state variable" s_t, the model admits two structures at two levels, and the switching between regimes (different distributions with different means) is governed by the state variable (Kuan, 2002). The state variable follows a first order Markov Chain (the probability of a certain state being true is only based on the state attained in the previous period),

$$P = \begin{bmatrix} \mathrm{P}(s_t = 0 \mid s_{t-1} = 0|) & \mathrm{P}(s_t = 1 \mid s_{t-1} = 0|) \\ \mathrm{P}(s_t = 0 \mid s_{t-1} = 1|) & \mathrm{P}(s_t = 1 \mid s_{t-1} = 1|) \end{bmatrix} = \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix}.$$

p_{ij} $(i, j = 0{,}1)$ denotes the transition probabilities of a state j in s_t being true (1), given that $s_{t-1} = i$ (Kuan, 2002). This is the Markov Switching model. In our case, we will use models with both two states and several states, and consequently the model and transition matrix must be expanded according to the possible states. For simplicity, we only explain the two-state model mathematically. As with most time-series models, it is important to work with stationary data when employing Markov Switching. This means removing trends and seasonality in the data, of which there are a lot in Bitcoin historical prices. Luckily for us, we can still use a Markov model when a unit root is present in the data, we can express the Markov Trend model as: $y_t = (\alpha_0 t + \alpha_1 \sum_{i=1}^{t} t s_i) + \beta_1 y_{t-1} + \cdots + \beta_k y_{t-k} + \sum_{i=1}^{t} \varepsilon_t$,where the function within the parenthesis is a trend function, the second term is the same as with the other model, and the final term is the stochastic trend (random component) (Kuan, 2002). This model will yield changes in the slope of the line depending on when the change to $s_i = 1$ happens and changes the slope by α_1 upwards or downwards. When the change in $s_i = 1$ is consecutive, the function will change the slope until $s_1 = 0$ again (Kuan, 2022). In figure 2 from Kuan (2002), we see two Markov Trend lines where the black boxes illustrate periods in which $s_i = 1$. The left model is a trend model in which $\alpha_0 > 0,\ \alpha_1 > 0$, and the right model is a trend model where $\alpha_0 > 0,\ \alpha_1 < 0$. In practice, employing a trend model does not require major changes, using the appropriate statistical tools in python.

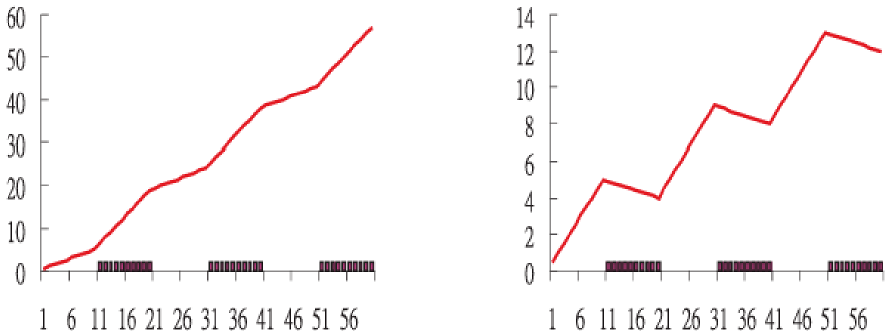

Figure 2: Left: markov model with $\alpha_0 > 0$ and $\alpha_1 > 0$. Right: A Markov model with $\alpha_0 > 0$ and $\alpha_1 < 0$

The Markov Switching model has two ways of being estimated, but we will only focus on "Quasi-Maximum likelihood estimation". We assume a vector of parameters: $\theta = (\alpha_0, \alpha_1, \beta_1, \dots, \beta_k, \sigma_\varepsilon^2, p_{00}, p_{11})'$.

Then we denote the information set up to time t as $Z^t = \{z_t, z_{t-1}, \dots, z_1\}$. We need to evaluate optimal forecasts based on different information sets. We differentiate between *prediction probabilities* $P(s_t = i \mid Z^{t-1}; \theta)$ (based on prior information), *filtering probabilities* $P(s_t = i \mid Z^t; \theta)$ (based on past and current information) and *smoothing probabilities* $P(s_t = i \mid Z^t; \theta)$ (based on full sample information) (Kuan, 2022). The quasi-maximum likelihood function is obtained as a byproduct of deriving the algorithms of these probabilities. The probabilities are derived by Kuan (2002) the following way: We assume normality, and denote the density of z_t conditional on Z^{t-1} and $s_t = i$ $(i = 0,1)$ is:

$$f(z_t | s_t = i, Z^{t-1};\ \theta) = \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}} \exp\{\frac{-(z_t - \alpha_0 - \alpha_1 i - \beta_1 z_{t-1} - \dots - \beta_k z_{t-k})^2}{2\sigma_\varepsilon^2}\}$$

We have previously stated the prediction probability as $P(s_t = i \mid Z^{t-1}; \theta)$, and we can derive the density of z_t conditional on Z^{t-1} alone from the previous density as:

$f(z_t | Z^{t-1}; \theta) = P(s_t = 0 \mid Z^{t-1}; \theta)\, f(z_t | s_t = 0, Z^{t-1}; \theta) + P(s_t = 1 \mid Z^{t-1};\ \theta)\, f(z_t | s_t = 1, Z^{t-1};\ \theta)$ (Kuan, 2002).

For $i = 0{,}1$ the filtering probabilities of s_t are:

$$P(s_t = i \,|Z^{t-1};\, \theta) = \frac{P\,(s_t = i\,|Z^{t-1}; \theta)\, f(z_t|s_t = i, Z^{t-1}; \theta)}{f(z_t|Z^{t-1}; \theta)}$$

And the relationship between the filtering probabilities and the prediction probabilities is:

$P\,(s_{t+1} = i\,|Z^t; \theta) = p_{0i}P(s_t = 0\,|Z^t;\, \theta) + p_{1i}P(s_t = 1\,|Z^t;\, \theta)$ (Kuan, 2002). In this relationship, $p_{0i} = P(s_{t+1} = i\,|s_t = 0)$ and $p_{1i} = P(s_{t+1} = i\,|s_t = 1)$ are transition probabilities (Kuan, 2002). From these functions we derive the quasi-log-likelihood function:

$\mathcal{L}_T(\theta) = \frac{1}{T}\sum_{t=1}^{T} ln f(z_t|Z^{t-1}; \theta)$ (Kuan, 2002). We compute smoothing probabilities (probability of being in a particular state at a past time) $P(s_t = i\ \ |Z^t; \theta)$ the following way:

$P(s_t = i \mid s_{t+1} = j, Z^T;\, \theta) = P(s_t = i \mid s_{t+1} = j, Z^t;\, \theta) = \frac{p_{ij}P(s_t = i\,|Z^t;\, \theta)}{P(s_t + 1 = j \mid Z^t;\, \theta)}$ (Kuan, 2002). For $i, j = 0{,}1$ the smoothing probabilities can be expressed as:

$P(s_t = i\,|Z^T;\, \theta) = P(s_{t+1} = 0\,|Z^T;\, \theta)\, P(s_t = i \mid s_{t+1} = 0, Z^T;\, \theta) + P(s_{t+1} = 1\,|Z^T;\, \theta)\, P(s_t = i \mid s_{t+1} = 1, Z^T;\, \theta) = P(s_t = i\,|Z^t;\, \theta) * (\frac{p_{i0}P(s_{t+1}=0\,|Z^T;\, \theta)}{P(s_{t+1}=0\,|Z^t;\, \theta)} + \frac{pi1\, P(s_{t+1}=1\,|Z^T; \theta)}{P(s_{t+1}=1\,|Z^t;\, \theta)})$ (Kuan, 2002). We iterate the equations backward to get the smoothing probabilities for $t = T - 1, \dots, k + 1$, which are, as the quasi-log-likelihood function above, functions of θ (Kuan, 2002).

We will use the Markov Switching model to estimate different market regimes and the transition between them. We will explore different indicator variables and regimes, mainly the *Crypto Greed and Fear* index as well as identifying regimes based on moving averages. We will create trading algorithms with the goal of beating a “buy and hold” strategy for a given time period of bitcoin.

The first Markov Switching model we estimate will identify trading signals based on moving averages. A moving average is a way for statisticians and investors to identify the current trending direction of a time series model, usually pertaining to a specific security. The moving average will give an indication of whether the security is on a downwards, upwards or sideways trend. This can provide valuable insight into which course of action the investor should take. The moving average can be calculated in several ways, but will focus on simple moving average, otherwise known as a one-sided moving average for prediction of future returns. Other types of moving averages include two-sided moving averages, exponential moving averages, weighted moving averages etc. These could be explored in further studies. The simple moving average is calculated in the following way:

$$z_t = \frac{1}{k+1}\sum_{j=0}^{k} y_{t-1}$$

Where $t = k+1, k+2, \ldots, n$ (Hyndman, 2011). The Markov Switching model will identify entry and exit points based on the moving average we calculate. The buying signal will be times where the 50-day moving average crosses above the 200-day moving average. Similarly, the selling signal will be times where the 50-day moving average goes below the 200-day moving average. We see the behavior of the 50-day and 200-day moving averages for the Bitcoin price in the timeframe in figure 3.

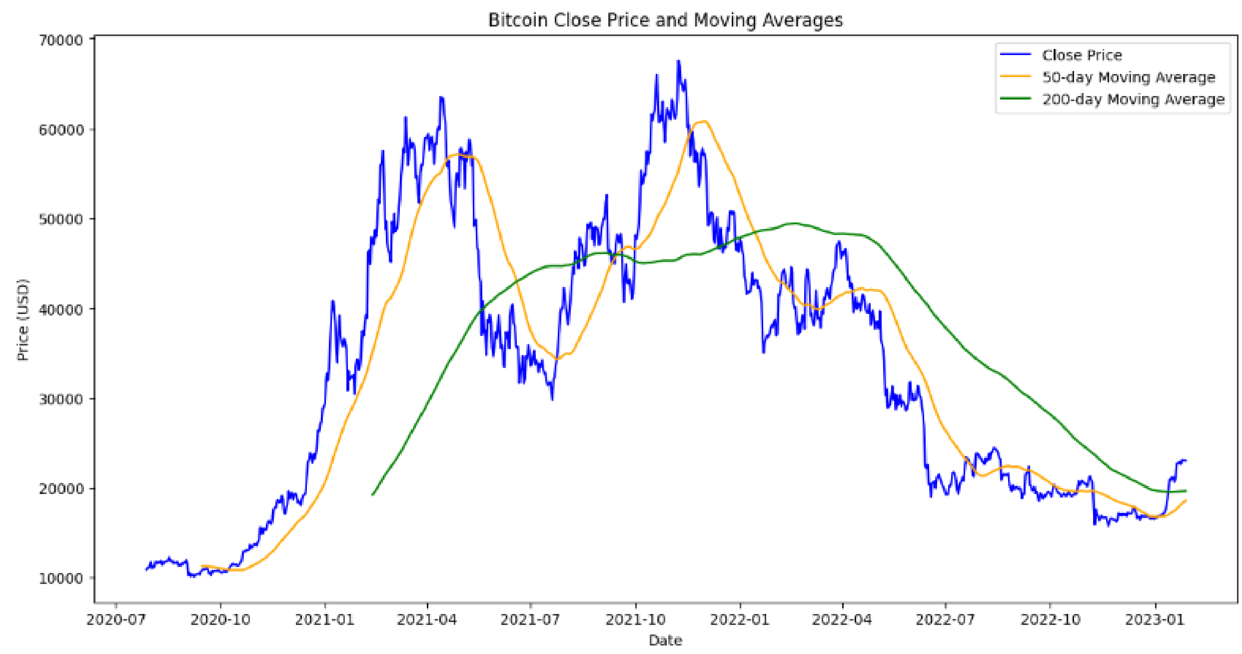

Figure 3: 50-day and 200-day moving average for Bitcoin prices

The second Markov Switching model we will estimate will be gathered from the Crypto Fear and Greed (CFGI) dataset that we detailed above. Bitcoin, and cryptocurrency in general is highly volatile, and a large part of the reason is due to investor psychology and behavioral tendencies. In times of high uncertainty, we postulate that there is a buying opportunity, and in times of excessive speculation we postulate that it is a good time to sell. We will identify the five different regimes (Extreme Fear, Fear, Neutral, Greed and Extreme Greed) using the Markov model with a time trend, since we can observe that Bitcoin had a rising (but highly volatile) trend during the period 03.02.2018-31.03.2023. Adding the time trend to the Markov model allows us to bypass the stationarity demand usually present in time series modelling. In figure 4 below we can observe how the price movement of Bitcoin relates to the identified market regime. What is most noticeable is that leading up to the all-time highs of March and October of 2021, we observe a period of greed and extreme greed in the market. During the fall from March to June

2021 we observe a period of extreme fear. We also notice this during the fall after the high of October 2021, a trend that persisted until around February 2021.

Figure 4:Identified regimes using the Crypto Fear and Greed Index

The Markov Switching model based on CFGI with short positions allowed is created in the following way. Firstly, we apply the “value” column as dependent variable. Then we initialize the model using the dependent variable and specify the trend and the number of regimes (which in this case is five). Afterwards, we fit the model to the data, which then estimates the parameters that best explain the observed behavior. The smoothing probabilities are then calculated to identify the most probable regime for each observation. Next, we translate the numerical regime identifiers to the labels in the CFGI (Extreme Fear – Extreme Greed). These identifiers are used to generate the trading signals, in which hold is set as default. We specify that the regime states “Fear” and “Extreme Fear” are buying signals (1), while the state “Extreme Greed” is a selling signal (-1). Finally, we calculate the logarithmic returns of Bitcoin to quantify results over time, as well as the returns generated from the trading signals.

For the Markov Switching model based on CFGI without short positions allowed, we change the sell signal from (-1) to (0).

4.2 LSTM Model

In this subchapter, we will introduce a LSTM neural network, a variant of RNN. We will also discuss the process of defining the LSTM model in Python and how we have used it to create a trading algorithm.

RNN stands for Recurrent Neural Network is a type of neural network that uses a feedback loop. This loop allows the output of a step *n-1* to be used again as input that will help predict step *n*. A RNN can therefore be used with sequential data like time series data. One can unfold the RNN so that we have input and output for each data point, this makes it easier to understand what happens (Hochreiter, 1998) (Sherstinsky, 2020). This can be seen in the figure below. The more data points one has the more one unfolds the RNN and it becomes harder to train because of the exploding- and the vanishing gradient problem (Hochreiter, 1998). To avoid these problems, we use a type of RNN called Long short-term memory (LSTM).

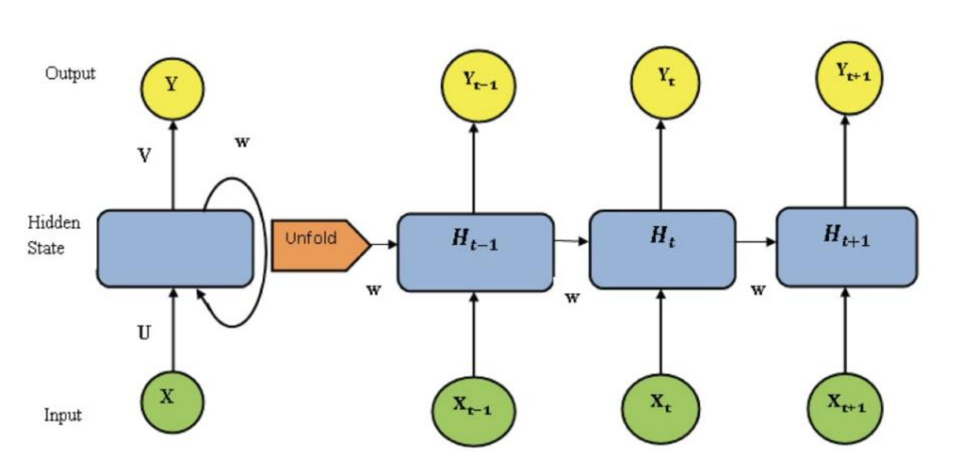

Figure 5: Unfolded RNN (Ancy & Latha, 2023)

The LSTM has two paths, one for the long-term memories and one for the short-term memories. It is based on a more complicated unit (as seen in the figure 6) than a traditional RNN. The cell state represents the long-term memory, and the hidden state represents the short-term memory. In the LSTM, the long-term memory path can only be modified by multiplication and addition, not weights or biases. This makes it so that no exploding- or vanishing gradient problem will occur. The short-term memory however will be impacted by weights and biases.

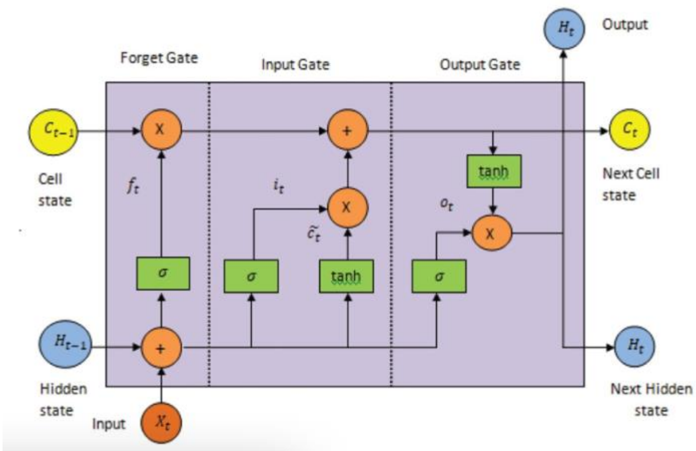

Figure 6: Architecture of LSTM cell (Ancy & Latha, 2023)

In the first stage, also called the *Forget Gate*, the hidden state and the input are multiplied by their weights and added together. Then a bias is added, and it's put into the sigmoid activation function, which leaves us with a number between 0 and 1. This number represents the percentage of long-term memory that will be remembered, we therefore multiply it with the cell state and go into the next stage termed the *Input Gate*. Here we see that the hidden state and the input are combined and put into a Tanh activation function leaving us with a number between –1 and 1 representing potential long-term memory based on the short-term memory and the input. In this stage the hidden state and input are also combined and put into a sigmoid activation function, this gives us the percentage of the potential long-term memory to be remembered. We multiply the potential long-term memory to be remembered by the percentage and add this to the exiting long-term memory creating a new long-term memory (next cell state). In the last stage, the *Output Gate*, we put the new long-term memory as input for the Tanh activation function leaving us with potential short-term memory. The hidden state and the input are also combined and put into a sigmoid activation function which also here gives us a percentage. This time, the percentage is of potential short-term memory to be remembered. We multiply them to find the new short-term memory, which is the output of the LSTM unit. This output can either be used again or if it is the final short-term memory of the unfolded LSTM, it is the model's output (Ancy & Latha, 2023).

For this thesis we will create an LSTM model in python based on Bitcoin's historical returns from the last 10 years. We first start by downloading the data defined earlier using yahoo finance, then we add a column with the calculated returns. We decided to predict returns instead of prices because of the high increase in BTCs price. When creating an LSTM model, one uses most of the data to train the model and the rest to test how the model performs. If we train our model on the lower prices from the beginning of the time period, we might not get the results wanted when predicting the higher prices from 2021/2022 (as seen in figure 2). By using returns we avoid some of this problem, because the price does not increase/decrease much from day to day, even though it does so over a time.

When creating and LSTM model in python we also want to split the data into X and y values, where the X values are sequences of historical BTC returns, and the y values are the next day returns that correspond to those specific sequences. For this LSTM model we decided to use a sequence length of 60, meaning that we use the previous 60 days (about 2 months) to predict the next day's return. We also want to split the data into so-called training and test sets. We train the model using the training set (70% of data) and then we test the model's performance using the test set (30% of data). This way we can compare the predicted returns to the actual returns for the test set data.

We will now use the Keras package to define the LSTM model. It contains 2 LSTM layers with 50 units each, these parameters are chosen so that the model will learn complex patterns without overfitting. To further prevent overfitting, we added dropout layers with a rate of 0.2, meaning that 20% of the units will be randomly dropped (Keras, u.d.).We then compile and fit the model, while creating a validation set of 20%. This is data that the model will not be trained on, but it will be used to evaluate the training process after each epoch (Keras, u.d.). When this is done, we may predict the return values.

Using the predicted returns from the LSTM model we create a trading algorithm similar to the one used for Markov Switching. We define a signal for each day in the test data which is 1 if the predicted return is positive and –1 if the predicted return is negative. We then use this signal to decide if one should buy or sell. Finishing off with calculating the cumulative returns and standard deviation, these will be used to look at how well the algorithm works and to compare with others.

5. Results

We compare the trading strategies with the buy and hold strategy, as well as with each other in different environments (with or without short). To make the strategies' returns comparable, we have chosen an interval within the smallest dataset, namely the CFGI dataset, which has data spanning from 28.07.2020 to 29.01.2023.

5.1 Moving average

The Moving Average models perform as we could expect, rather poorly compared to a buy and hold strategy. The holding period return was 116,31% for the period, while the model with short capabilities had a return of 93,52%. As we can see from figure 7, the deviation from the buy and hold strategy happens quite late in the time period, and the trade is a short sale. After this trade, the trading system gains a higher return with a negative correlation with the bitcoin price. However, after bitcoin recoups its fall in market value during late 2022 and early 2023, the model fails to identify the exit of its short position, and thus wipes out all the excess returns generated.

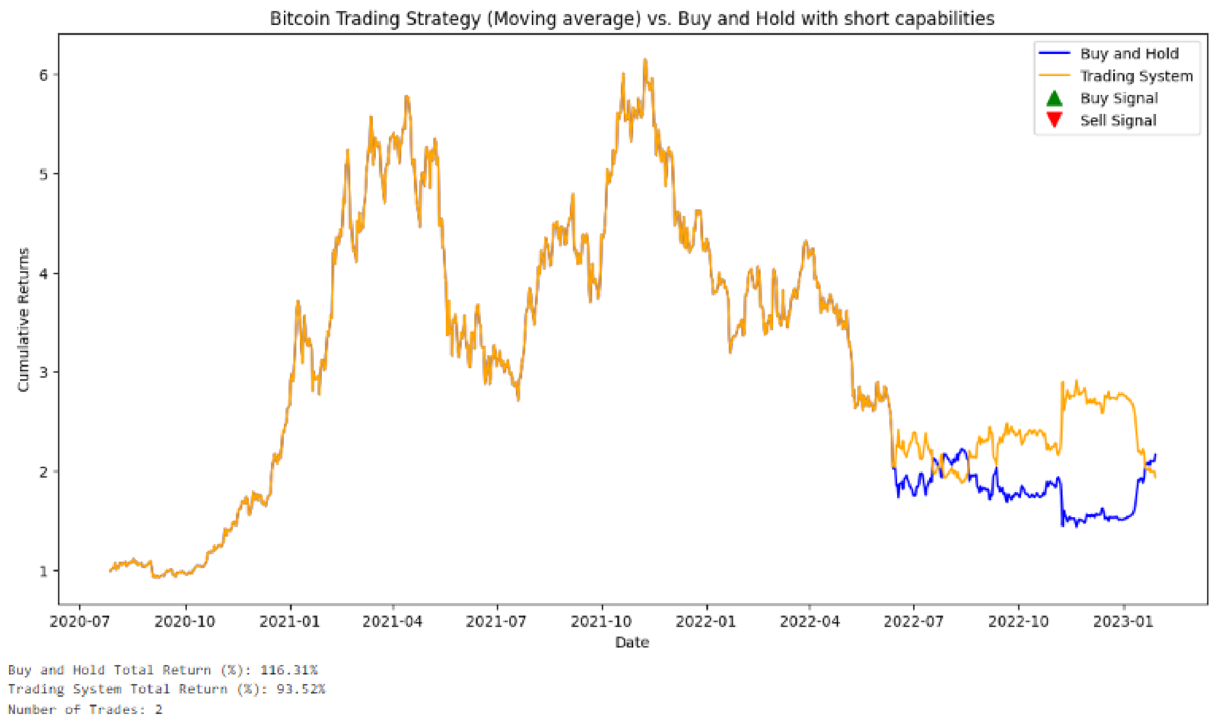

Figure 7: Result of Markov Model based on Moving Averages with short positions allowed

The model that does not allow for short sales surprisingly performs slightly better than the model that does allow for short sales, generating a return of 104,6%. What we observe from figure 8 is that the model only performs one transaction, selling after bitcoin had dropped quite a bit in a short timespan. The model then holds all its holdings in a risk-free asset (such as a deposit yielding 0%). It correctly identifies a decline in Bitcoin price, however the excess returns are also wiped out during the rise in late 2022 and early 2023. Both these models performed worse than simply buying bitcoin and holding it for the period.

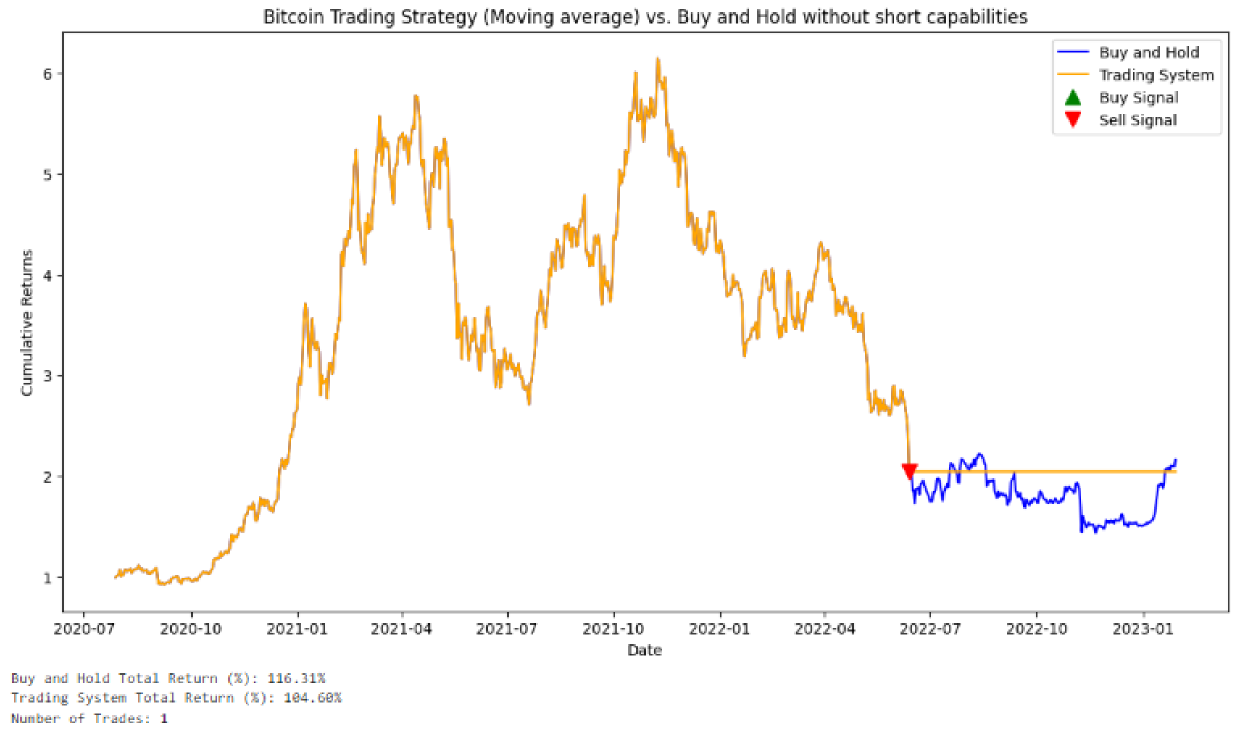

Figure 8:Result of Markov model based on Moving Averages without short positions allowed

Comparing the figure that shows the behavior of the 50-day and 200-day moving averages, and the models based on these moving averages, we notice that the Markov models perform poorly in identifying the buy- and sell signals. A model based purely on these signals (and not a Markov model interpreting signals) would start by selling bitcoin during the downtrend after the initial all time high of early 2021. It would then buy during the uptrend leading up to the all-time high of late 2021 and sell again in the downtrend early 2022. This we can observe from figure 3. This model would have three transactions, close to the amount that the Markov models perform. This model might have performed better than the Markov models for this time period, however, it is uncertain how much fault we can attribute to the model, as Bitcoin price is highly variable, and the results will vary depending on the time period. The Markov model might also need a longer time period to appropriately identify the patterns and make trading decisions based on these.

The annualized standard deviation of Bitcoin for this period was 69,44%. The annualized standard deviation for the model allowing for short sales was also roughly identical, while the model not allowing for short sales had a standard deviation of 63,3%, slightly less than the other two. An investor that has the risk appetite to invest in Bitcoin, might not be tempted by the slight reduction in risk for the almost 12% reduction in returns.

The findings from the first model suggest that employing moving averages in trading decisions does not generate excess returns, which strengthens the claim that Bitcoin is at least weak-form efficient. In other words, one cannot generate excess returns based on historical data, and the current price reflects all historical information. However, it should be noted that the data generated from creating 50- and 200-day moving averages is significantly reduced when operating with such a small time period. Since the model loses 200 of the total 917 datapoints due to the nature of moving averages, it would be more accurate to measure such a model over the course of a significantly longer time span. This is an interesting research topic for further studies.

5.2 Crypto Fear and Greed Index

Both Markov models based on the CFGI perform surprisingly well, which contradicts the findings from the simple moving average models. We note from figures 9 and 10 that the holding period return was 116,31%, while the returns for the models with and without short were 725,36% and 364,87% respectively. The algorithm that allowed short sales performed nine transactions, while the one that did not allow short sales performed one. This strengthens our initial theory that employing short selling in the strategy could provide more opportunity for higher returns, due to Bitcoins volatile nature. The model that allows for short sales provides almost exactly double the returns as the other one, and over 6 times the market return. The second model still manages an impressive tripling of the market return, simply by timing a favorable exit.

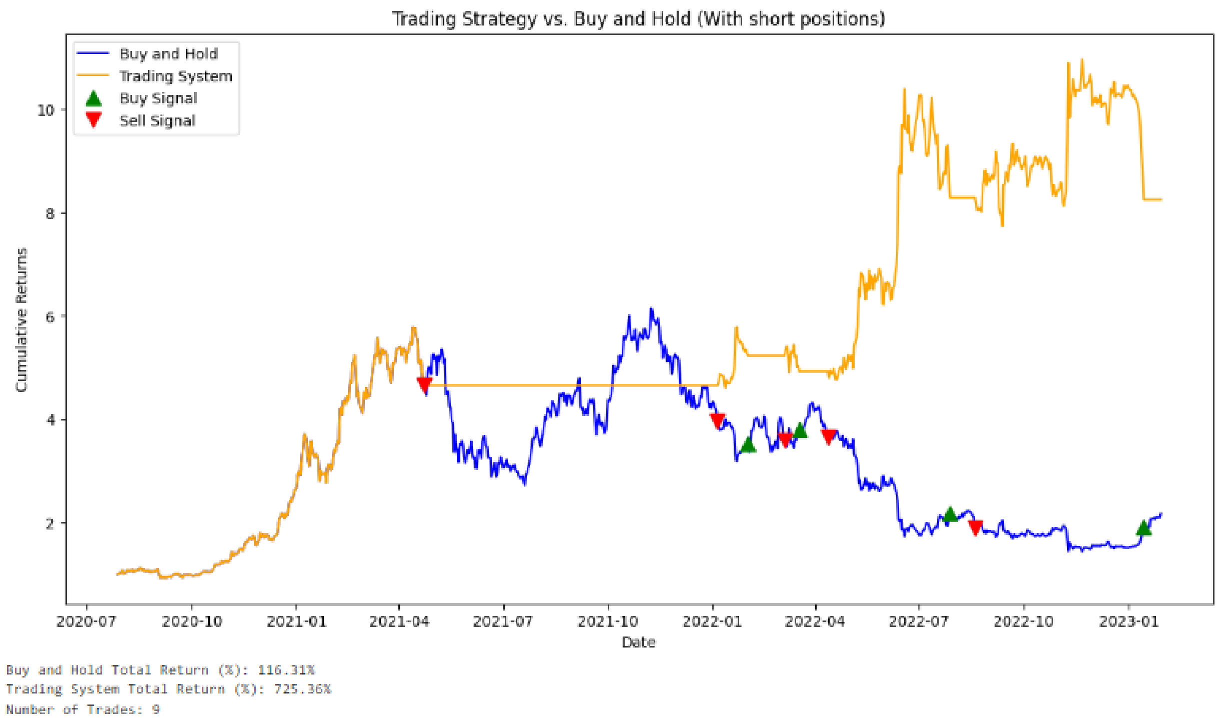

Figure 9:Results of Markov Model based on CFGI with short positions allowed

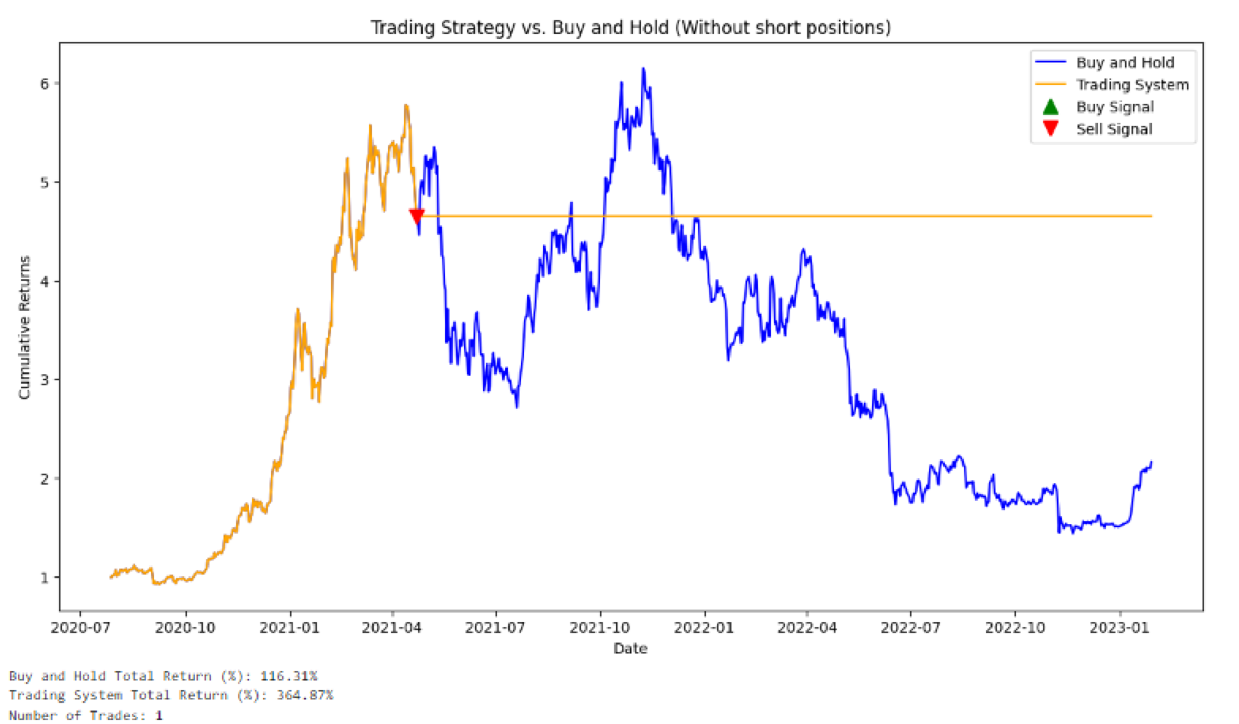

Figure 10:Results of Markov model based on CFGI without short positions allowed

There are, however, several reasons why we should take these returns with a grain of salt. The period chosen is extremely volatile, with two all-time highs and subsequent crashes, which gives ample opportunity for trading with short capabilities. This might not be the case in the future, which would mean that the algorithm may perform worse. Also, the standard deviation for the short algorithm is 52,23%, well above the risk appetite of many institutional and retail investors. However, we should mention that the annualized standard deviation for Bitcoin in the time period was 69,44%, meaning that if an investor was holding bitcoin in this time period, they would actually reduce risk significantly by employing the model. We also notice that the model makes the largest returns when bitcoin is trending downwards, which indicates that the model prefers returns generated from short sales. This may perhaps be the case due to these down periods lasting longer than the upwards trends, generating more reliable returns. The investor mimicking the model would have to hold short periods for long amounts of time, which means that they expose themselves to a large downside risk, since the losses are theoretically infinite. There may also be problems with implementing instruments that allow for such short sales with such a large timespan, and the transaction costs related might be significant.

The algorithm that does not allow for shorting has an identical approach as the first model for the first trade. After selling before the large fall in mid-2021, it holds its returns for the rest of the period. We notice a substantially lower standard deviation, at 37,89%. This model also makes fewer trades, which indicates that the true spread between the models is slightly lower, due to more transaction costs in the first model. However, it is unlikely that transaction costs would make this gap significantly smaller. It is also worth noting that the model does not identify the bottom of the dip between the two all-time highs, which would generate a substantial higher return than simply exiting the market entirely from the first all time high. This would also imply higher risk, since it also assumes that the model can identify a favorable exit point before the large fall in late 2021 and early 2022.

The annualized standard deviation for the period for Bitcoin is 69,44%. In comparison, the annual standard deviation of the SPY exchange traded fund (which tracks the S&P 500) for the period was 19,13% (Yahoo Finance, 2024). The annual expected return for Bitcoin and SPY in the period is 80,49% and 10%, respectively. We can calculate Sharpe ratios for both securities by finding the average risk-free rate for the period, which is 0,853% (UST, 2024). This equates to an impressive Sharpe ratio of 113,84% for bitcoin and 47,85% for SPY. This is an indication that, for the period, the risk adjusted return you could expect from bitcoin was well over double the risk adjusted return you could achieve from the broader market index S&P 500. We note that for most of the time period we chose, the risk-free rate was near zero, which would generate a larger Sharpe ratio due to higher risk premiums. However, even with a 5% risk free rate, the Sharpe ratio of Bitcoin for this period would be 107,91%.

The findings from the Markov models based on the Crypto Fear & Greed Index deviate from the findings we had in the moving average models. The excess returns generated based on historical data in the CFGI model suggests that the current market price of Bitcoin does not reflect all historical information, which indicates that Bitcoin is weak form inefficient. However, if we were to choose a different time-period, the results may have been substantially different, perhaps even negative. It is difficult to ascertain how much of the excess returns we can attribute to the performance of our models, and how much is “luck” related to the specific time period we chose. This is an interesting point which could be an object for further studies.

5.3 LSTM

This section is dedicated to the results of the LSTM model and in the figure below you can see the trading algorithm based on the LSTM model. The systems return outperforms the buy and hold return, with cumulative returns of 150,22% and 116,31%, respectively. The trading algorithm for the LSTM model also outperforms the Markov Switching model based on moving averages in terms of returns. The LSTM based algorithm performs 16 trades over the almost 3-year period, leaving us with an average of about 6 trades pr year. Around April 2022 the system's cumulative return drops after selling and we see a similar return for both the system and the buy and hold strategy all the way until September of that same year. Here it first sells and then shortly after buys, making the cumulative return for the algorithm end up higher than the buy and hold return. The system starts by shorting BTC, and we see a tendency for it to first sell and then buy again shortly after.

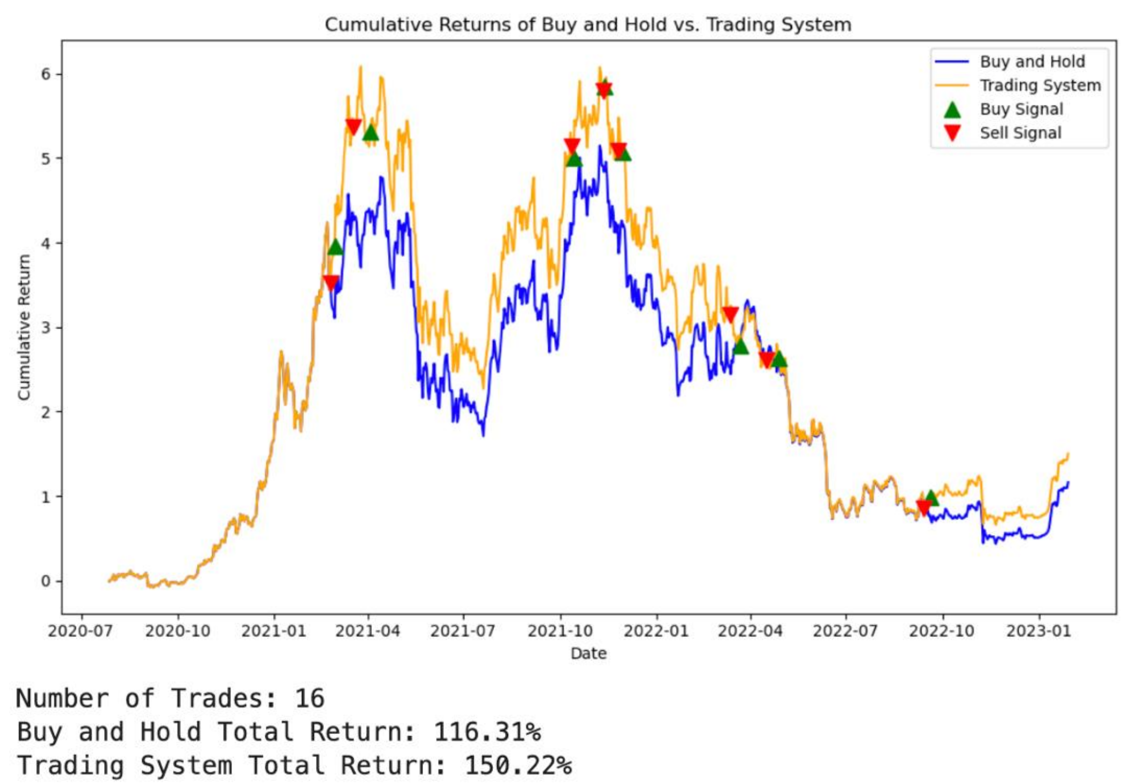

Figure 11:Buy and Hold vs. LSTM Trading system

Even though the system ends up with a higher cumulative return after the trades, it does keep quite close to the buy and hold over the period. This leaves both strategies with similar risk here in the form of annualized standard deviations of about 69% for both strategies, note that this is also higher than the Markov models.

The LSTM model might be overfitted and we see poor results from the loss calculations during training in figure 12. The validation loss is always close to zero, and after epoch 10 we can see a small increase in validation loss together with a decrease in training loss, indicating that the model might be overfitted and therefore not performing how we wanted.

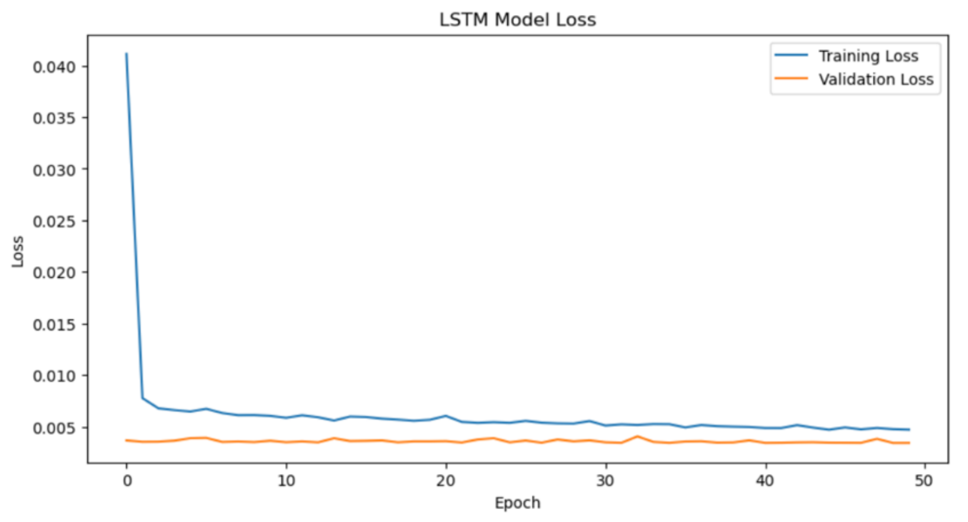

Figure 12:Training loss vs. Validation loss

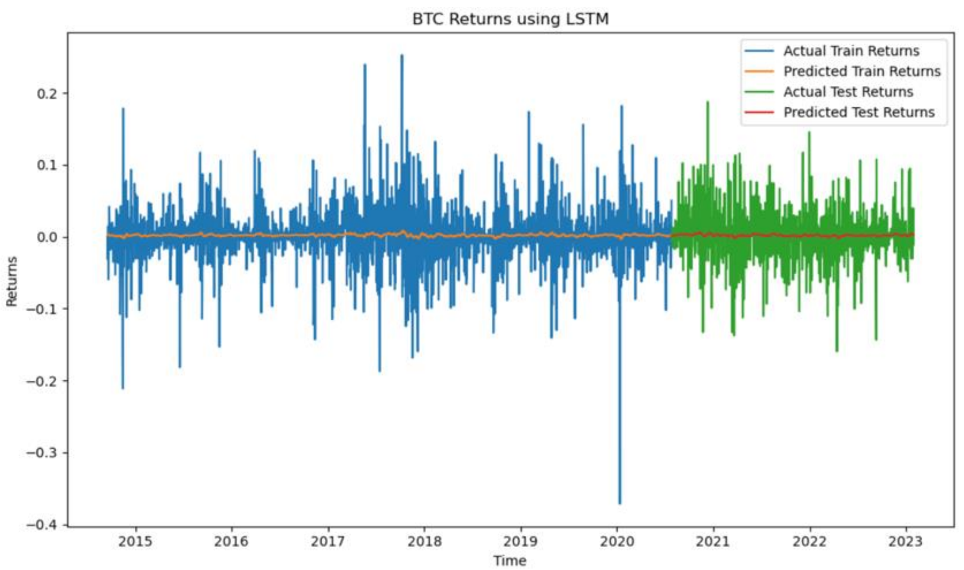

Figure 13: LSTM model predictions

The predicted returns are also relatively small compared to the actual returns. This can be seen in the chart above. The predicted returns used for the algorithm are from the test set (depicted as red in the chart, while the actual return from this period is depicted in green). Even though they are quite small, the predicted returns often have the correct sign. Our algorithm is based on signals that indicate if the returns are positive or negative, not the size of them. Therefore, the algorithm might perform well even though the predicted returns are quite small.

6. Conclusion

This thesis has examined various models for trading Bitcoin, focusing on their effectiveness over a range of trading strategies and market conditions. We examined the performance of three distinct models, using two methodologies: Markov Switching models based on moving averages and the Crypto Fear and Greed Index (CFGI), and Long Short-Term Memory (LSTM) neural network. We created two models for each methodology and dataset, one that could perform short sales and one that could not. Each model's efficacy was measured against the traditional buy-and-hold strategy, providing insights into their practical application in the volatile cryptocurrency market. The models were also compared to each other, both by means of risk (measured by standard deviation) and return (measured by returns and risk adjusted returns through the Sharpe ratio).

Our analysis revealed that while simple historical data-based models like the Moving Average underperform, sentiment-driven and advanced predictive models like CFGI and LSTM showed promising results. Particularly, the CFGI model, when incorporating short selling based on market sentiment indicators, significantly outperformed the passive strategy, underscoring the potential of reactive models in leveraging market psychology. The LSTM model further demonstrated the capability of machine learning technologies to adaptively predict price movements with a reasonable degree of accuracy. These findings suggest that the choice of trading model, the sophistication of the data analysis techniques employed, and the ability to capture and react to real-time market dynamics are crucial for success in cryptocurrency trading. However, the practical applicability of the models remains to be seen. We must also note that while the models performed well overall, variables such as the time period chosen may have had a significant impact on the result. This indicates that further research must be done before concluding on the efficiency of the Bitcoin market.

The comparatively poor performance of the moving average model may be explained by the fact that such a model requires larger datasets, due to the nature of the calculation. The model lost about 22% of its total data, which gives the model less data to make conclusions based on. This model could perform better, given a larger dataset. Also, other models such as exponential moving averages and Time Series Forecast (TSF) could be applied to this data. Further studies could also look at a larger time period, as well as training and applying the models to different time periods and comparing the results. Finally, a model with added complexity, implementing transaction costs, macro-and microeconomic data and other market mechanisms could provide a larger insight to the efficiency of the Bitcoin asset.

7. Bibliography

McLeay, M., Radia, A. & Thomas, R., 2014. Money in the modern economy: an introduction. *Bank of England Quarterly Bulletin Q1 2014*, p. 13.

Kuan, C.-M., 2002. *Lecture on the Markov Switching Model,* s.l.: Institute of Economics, Academia Sinica.

Hyndman, R. J., 2011. Moving Averages. *International encyclopedia of statistical science*.

Trainer, D., 2023. *Forbes.* [Internett]
Available at: https://www.forbes.com/sites/greatspeculations/2023/07/20/teslas-2q23-earnings-confirm-my-thesis-that-tsla-is-one-of-the-most-overvalued-stocks/
[Funnet 20 July 2023].

Johnson, J., 2023. *Social Science Research Network.* [Internett]
Available at:
https://deliverypdf.ssrn.com/delivery.php?ID=474090073005018102082066117110026086022027028059062003011089119000073000030003026000041101048107026028021110089001094117105028028085086079040081097084007100097071088014094025109100126019076065097092122069121118
[Funnet 22 March 2024].

Baldridge, R., 2022. *Forbes.* [Internett]
Available at: https://www.forbes.com/advisor/investing/efficient-market-hypothesis/
[Funnet 23 March 2024].

Burton, M., 2003. *The Efficient Market Hypothesis and Its Critics.* [Internett]
Available at: https://www.princeton.edu/~ceps/workingpapers/91malkiel.pdf
[Funnet 22 March 2024].

Forbes, 2024. *Forbes.* [Internett]
Available at: https://www.forbesindia.com/article/explainers/bitcoin-price-history/92523/1
[Funnet 11 April 2024].

Jansen, I. P. & Nikiforov, A. L., 2006. Fear and Greed: a Returns-Based Trading Strategy around Earnings Announcements. *Journal of Portfolio Management,* I(1), pp. 3-5.

Greengold, R., 2023. *Morningstar.* [Internett]
Available at: https://www.morningstar.com/funds/large-growth-active-fund-managers-have-faced-an-uphill-battle
[Funnet 25 March 2024].

Huang, R., 2022. *Forbes.* [Internett]
Available at: https://www.forbes.com/sites/rogerhuang/2022/12/04/does-bitcoin-beat-inflation/
[Funnet 24 March 2024].

Jackson, R., 2023. *Morningstar.* [Internett]
Available at: https://www.morningstar.com/etfs/actively-managed-funds-measured-up-well-market-rebound
[Funnet 22 March 2024].

Pratley, N., 2021. *The Guardian.* [Internett]
Available at: https://www.theguardian.com/business/nils-pratley-on-finance/2021/may/19/the-30-one-day-fall-in-bitcoins-value-looks-like-a-turning-point
[Funnet 22 March 2024].

Tambe, N. & Jain, A., 2023. *Forbes.* [Internett]
Available at: https://www.forbes.com/advisor/in/investing/cryptocurrency/fear-and-greed-index-crypto/
[Funnet 25 March 2024].

Ancy, J. & Latha, T., 2023. Stock market prediction based on deep hybrid RNN model and sentiment analysis. *Automatika*, pp. 981-995.

Sherstinsky, A., 2020. Fundamentals of Recurrent Neural Network (RNN) and Long Short-Term Memory (LSTM) network. *ScienceDirect* , 29 January.

Keras, u.d. *Model training APIs.* [Internett]
Available at: https://keras.io/api/models/model_training_apis/
[Funnet 07 March 2024].

Keras, u.d. *Dropout layer.* [Internett]
Available at: https://keras.io/api/layers/regularization_layers/dropout/
[Funnet 07 March 2024].

The Pennsylvanina State University, u.d. *Applied Time Series Analysis.* [Internett]
Available at: https://online.stat.psu.edu/stat510/lesson/11/11.2
[Funnet 24 January 2023].

PwC, u.d. *Making sense of bitcoin, cryptocurrency and blockchain.* [Internett]
Available at: https://www.pwc.com/us/en/industries/financial-services/fintech/bitcoin-blockchain-cryptocurrency.html
[Funnet 07 February 2024].

Adams, M. & Dammeyer, L., 2024. *Top 10 Cryptocurrencies Of February 2024.* [Internett]
Available at: https://www.forbes.com/advisor/investing/cryptocurrency/top-10-cryptocurrencies/
[Funnet 07 February 2024].

Yahoo Finance, 2024. *Bitcoin USD (BTC-USD).* [Internett]
Available at: https://finance.yahoo.com/quote/BTC-USD/history/
[Funnet 24 January 2024].

Hochreiter, S., 1998. The Vanishing Gradient Problem During Learning Recurrent Neural Nets and Problem Solutions. *International Journal of Uncertainty Fuzziness and Knowledge-Based Systems,* 01 April , pp. 107-116.

CFI Team, u.d. *Corporate Finance Institute.* [Internett]
Available at: https://corporatefinanceinstitute.com/resources/career-map/sell-side/capital-markets/efficient-markets-hypothesis/
[Funnet 23 March 2024].

CFI Team, u.d. *Corporate Finance Institute.* [Internett]
Available at: https://corporatefinanceinstitute.com/resources/career-map/sell-side/capital-markets/efficient-markets-hypothesis/
[Funnet 23 March 2024].

Eom, C., Kaizoi, T., Kang, S. H. & Pichl, L., 2019. Bitcoin and investor sentiment: Statistical characterisitcs and predictability. *Physica A,* Volum 514, pp. 511-521.

Alternative, u.d. *Alternative.* [Internett]
Available at: https://alternative.me/crypto/fear-and-greed-index/
[Funnet 4 April 2024].

Ashford, K. & Adams, M., 2022. *What Is Bitcoin And How Does It Work?.* [Internett]
Available at: https://www.forbes.com/advisor/investing/cryptocurrency/what-is-bitcoin/
[Funnet 07 February 2024].

Mokni, K., Bouteska, A. & Nakhili, M. S., 2022. Investor sentiment and Bitcoin relationship: A quantile-based analysis. *North American Journal og Economics and Finance,* Volum 60.

Gaies, B., Nakhli, M. S., Sahut, J.-M. & Schweizer, D., 2023. Interactions between investors' fear and greed sentiment and Bitcoin prices. *The North American Journal of Economics and Finance,* Volum 67.

Łet, B., Sobański, K., Świder, W. & Włosik, K., 2022. Inverstor Sentiment and Efficiency of the Cryptocurrency Market: The Case of the Crypto Fear and & Greed Index. *Eurasian Business and Economics Perspectives ,* Volum 25, pp. 271-287.

Huang, Q., 2024. Bitcoin price prediction based on fear and greed index. *SHS Web Conferences,* Volum 181.

Çelik , I., 2020. Can Bitcoin be a stable investment?. *Financial Studies,* 24(2), pp. 19-36.

Malladi, R. K. & Dheeriya, P. L., 2020. Time series analysis of Cryptocurrency returns and volatilities. *Journal of Economics and Finance ,* 45(1), pp. 75-94.

Ibrahim, A., Kashef, R. & Corrigan, L., 2021. Predicting market movement direction for bitcoin: A comparison of time series modeling methods. *Computers & Electrical Engineering,* Volum 89.

Chkili, W., 2021. Modeling Bitcoin price volatility: long memory vs Markov switching. *Eurasian Econ Review,* Volum 11, pp. 433-448.

McNally, S., Roche, J. & Caton, S., 2018. Predicting the Price of Bitcoin Using Machine Learning. *Euromicro International Conference on Parallel,* Volum 26, pp. 339-343.

Mudassir, M., Bennbaia, S., Unal, D. & Mohammad, H., 2020. Time-series forecasting of Bitcoin prices using high-dimensional features: a machine learning approach. *Neural Comput & Applications.*

Appendix

Please see attached ZIP-file

Appendix 1: PDF-file of Markow switching models code

Appendix 2 (dataset.csv): CSV file of CFGI dataset

Appendix 3: PDF-file of LSTM model code

Appendix 4: Declaration of AI

Norwegian University of
Science and Technology