2-2024

Bitcoin Short-term Price Prediction Using Time Series Analysis

Alanood Alkamali
aa9681@mail.rit.edu

Bitcoin Short-term Price Prediction Using Time Series Analysis

by

Alanood Alkamali

A Thesis Submitted in Partial Fulfilment of the Requirements for the Degree of Master of Science in Professional Studies: Data Analytics

Department of Graduate Programs & Research

Rochester Institute of Technology – Dubai

February 2024

RIT

Master of Science in Professional Studies:

Data Analytics

Graduate Thesis Approval

Student Name**:** Alanood Alkamali

Thesis Title**:** Bitcoin Short-term Price Prediction Using Time Series Analysis

Graduate Committee:

Name: **Dr. Sanjay Modak** **Date:**

Chair of committee

Name: **Dr. Ioannis Karamitsos** **Date:**

Member of committee

Acknowledgments

First, I would like to express my deepest appreciation to my supervisor and mentor Dr. Ioannis Karamitsos for his constant support and guidance throughout this journey. His experience and knowledge were invaluable in shaping this research.

I extend my thanks to Dr. Sanjay Modak, Head of Committee at RIT Dubai for his professional advice and insightful feedback. Furthermore, I am grateful to RIT for providing the necessary academic resources.

Abstract

This thesis explores the application of Autoregressive Integrated Moving Average (ARIMA) model to predict Bitcoin prices, a prominent and volatile cryptocurrency. The research falls within the context of financial forecasting, focusing specifically on the cryptocurrency markets.

The primary research question is: “Can time series analysis be used to predict the future price of Bitcoin?” To answer this question, historical Bitcoin daily price data from 17/09/2014 to 17/09/2023 was obtained from Kaggle and analyzed.

The study employs ARIMA modeling techniques to capture the autocorrelation, seasonality, and trend present in Bitcoin price time series. As a prerequisite for ARIMA modeling, the data was transformed using a logarithmic function to stabilize the variance, then differenced by an order of 1 to make it stationary.

The findings reveal that ARIMA can in fact predict Bitcoin prices. The best model in terms of lowest error rate is ARIMA(4,1,1), which achieved an RMSE of 0.03099 and MAE of 0.02121. However, the lowest MAPE that could be achieved using historical data alone was 123%. This indicates that traditional time series techniques are limited by their use of only past values to predict future ones, especially considering that cryptocurrency prices are influenced by various other features and external factors.

Future research should explore correlations of Bitcoin with other currencies, include additional factors, and investigate hybrid models like ARIMA with CNN or LSTM for improved predictions.

Keywords: Bitcoin, cryptocurrency, machine learning, ARIMA, time series analysis.

Table of Contents

List of Figures

List of Tables

Chapter 1. Introduction

"There are 3 eras of currency—commodity based, politically based, and now, math based."

- Chris Dixon

1.1 What is cryptocurrency?

Cryptocurrencies, or digital currencies, are peer-to-peer decentralized virtual currencies. They operate on the principles of cryptography, which is a method that encrypts transaction data in a secure network to control and confirm the transfer of the currencies. Most types of cryptocurrencies keep track of transactions by using a distributed decentralized database utilizing a network of computers, instead of a single server, to store the data. This distributed database is called a blockchain and will be explained in more detail later. Cryptocurrencies are not issued by any central bank, not controlled by any government, and do not require intermediaries. They enable direct peer-to-peer trading and instant payments (Tredinnick, 2019).

There were many attempts to create a digital currency in the 1980s but none of them were successful. In 1990, the first digital currency, DigiCash that used an anonymous and secure payment system was created but it was never widely adopted (Lipovyanov, 2019). From the late 1990s onwards, PayPal and similar companies emerged as online payment services that link a user's bank account or credit card to their account (PayPal, 2023). These companies continue to be major players in online transactions and international trade. However, they do not transact in digital currency but instead use traditional fiat currencies to send and receive online payments.

Moreover, PayPal is a centralized system controlled by a single company and is regulated by financial institutions as a money transmitter. This means that if the company were to fail, all PayPal payments would be unavailable. On the other hand, cryptocurrencies are decentralized and therefore more secure as there is no central entity that can be infiltrated and hacked.

1.2 What is a blockchain?

A blockchain is a record of transactions that is encrypted and publicly accessible. All records in a blockchain are grouped in blocks that are then ordered chronologically and linked cryptographically to the previous blocks in a chain of sorts. The transaction records in a blockchain are permanent and cannot be modified, which makes it a secure and trustworthy platform. In the case of Bitcoin, a blockchain has all the Bitcoin transactions that were ever made and is continually growing as more of the Bitcoin is mined. It is also decentralized as no one person or group has control over it. Therefore, the blockchain operates in a direct, peer-to-peer way with no central authority needed to process the transactions or information transfer (Lipovyanov, 2019).

1.3 What is Bitcoin?

Bitcoin is the first decentralized digital currency and is now the most valuable and widely recognized cryptocurrency (CoinMarketCap, 2023). It is considered a form of asset and is used as an online payment tool that operates independently of financial institutions. As Satoshi Nakamoto (2008) notes in the original Bitcoin paper, the goal of the cryptocurrency is to create "a purely peer-to-peer version of electronic cash" that would facilitate online payments without the need for or intervention of a financial institution or a payment processor.

The process of mining Bitcoin involves the creation of blocks that contain information about transactions conducted within the network. To elaborate, as transactions get verified, new blocks are created and then added to a chain of blocks, forming what is commonly known as a blockchain. Using powerful computers, miners essentially compete with each other to record transactions by solving complex mathematical problems. The first miner to solve a problem is then rewarded with a block of Bitcoins, which is the set of transactions that was verified by that miner (Tredinnick, 2019).

Bitcoin's limited supply of only 21 million coins has also contributed to its appeal as a scarce and valuable asset. Since it is a decentralized monetary system, there is no human intervention in the cryptocurrency's issuance. Instead, the release of new Bitcoins into circulation is controlled by a special algorithm that sets a clear timetable for its issuance. The frequency at which blocks are generated is consistent at six per hour. The number of newly mined Bitcoins is reduced by 50% every 210 thousand mined Bitcoin blocks, resulting in a four-year issue cycle. Accordingly, the number of issued Bitcoins will never exceed 21 million coins (Meynkhard, 2019).

Additionally, the adoption of Bitcoin by some central banks and governments as a reserve currency and the growing use of blockchain technology, which Bitcoin utilizes, in the financial sector have also helped in boosting its popularity. Bitcoin has been considered as a potential hedge against geopolitical risk as well, with evidence suggesting that it can serve as a safe haven asset during times of market turbulence (Urquhart & Zhang, 2018).

1.4 Top Cryptocurrencies by Market Capitalization

Rank	Name	Symbol	Market Cap	Price	Circulating Supply	Volume (24h)
1	Bitcoin	BTC	$725,728,703,583	$37,112.57	19,554,525 BTC	$18,700,878,486
2	Ethereum	ETH	$242,692,081,919	$2,018.31	120,244,917 ETH	$10,110,357,972
3	Tether	USDT	$88,785,306,625	$1.00	88,783,826,002 USDT	$37,657,856,122
4	BNB	BNB	$34,436,765,495	$227.01	151,699,279 BNB	$774,938,656
5	USD Coin	USDC	$32,448,973,709	$0.6030	53,816,975,568 XRP	$1,202,960,232
6	XRP	XRP	$24,510,301,243	$1.00	24,511,709,915 USDC	$4,653,614,228
7	Cardano	ADA	$23,224,187,547	$54.85	423,425,777 SOL	$1,181,995,134
8	Dogecoin	DOGE	$13,395,379,809	$0.3795	35,298,243,162 ADA	$246,248,845
9	Polygon	MATIC	$11,148,971,990	$0.07851	142,003,346,384 DOGE	$738,932,521
10	Solana	SOL	$9,232,112,901	$0.1042	88,564,691,648 TRX	$266,541,190

Table 1 Top 10 Cryptocurrencies by Market Capitalization as of 27 November 2023 (CoinMarketCap, 2023)
Market Cap is the total market value of a cryptocurrency in circulation. Market Cap = Current Price x Circulating Supply. Volume is how much of a cryptocurrency was traded in the last 24 hours. Circulating Supply is the amount of cryptocurrency that is in circulation in the market.

1.5 Advantages and Disadvantages of Cryptocurrency

As with any innovative concept, cryptocurrency has both positive features and potential drawbacks. Below is an overview highlighting the advantages and disadvantages of cryptocurrency.

1.5.1 Advantages and Features

- Decentralized Nature: The cryptocurrency trading network is decentralized and not controlled by any government or central authority, making it accessible without any rules or restrictions.
- Low Transaction Costs: Since cryptocurrencies are decentralized, there is no need for banks to verify and process transactions, which results in significantly lower transaction costs.
- Worldwide Use: There are no foreign exchange fees when trading with cryptocurrency since they are global currencies.
- Irreversible: Cryptocurrency transactions are irreversible and require approval from both parties.
- Security: Through the use of a blockchain, each Bitcoin has a unique identity that can be traced back to its origin so any fake or stolen currency can be easily identified. Additionally, the record cannot be changed without changing every previous transaction, making the blockchain robust.
- Transparency: All transactions are recorded and available on the network in real time for all users and they cannot be modified as they are encrypted, which also reduces the risk of fraud.
- Speed of transactions: Cryptocurrency transactions are processed in near real time.
- No Inflation: Most cryptocurrencies, namely Bitcoin, have a limited supply so they cannot be devalued in the traditional way that fiat currencies can be devalued when governments print more money.
- Foreign Trade: Lower transaction costs of cryptocurrencies increase foreign trade volumes and allow for easier and lower tariffs on international trade.

1.5.2 Disadvantages and Risks

- High Volatility: The value of cryptocurrencies is volatile and unpredictable. It is affected by various factors and could plummet following an exposed scam, suspected hack attempt, or crash in a digital exchange.
- Sustainability Issues: The mining reward's system encourages increased use of computing power that is energy-intensive. This high electricity consumption exacerbates the sustainability challenges that outweigh the benefits enjoyed only by a small group of people.
- Illegal Activities: Anonymity makes it difficult to track transactions, which can be used by criminals to carry out illegal activities and money laundering e.g. The Silk Road incident.
- Acceptance: Cryptocurrency is not widely accepted by the public, governments, or banks. This is because most people do not understand how it works and are afraid of losing control of their money.
- Cybersecurity Risks: There is also the risk of a cryptocurrency exchange poorly managing funds and losing customers’ money like what happened with FTX in 2022 (Reiff, 2023), or being vulnerable and a target for hackers like the infamous Mt. Gox 2011 and 2014 incidents (Frankenfield, 2023), and more recently the Poly Network exploit attack in 2021 in which more than $600 million worth of cryptocurrency was stolen then returned 15 days later (Ponciano, 2021).

1.6 Biggest Crypto Exchanges by Trading Volume

Cryptocurrencies are traded on online exchanges that accept credit card payments, money transfers, and other payment methods in exchange for cryptocurrencies. According to CoinMarketCap (2023), the biggest cryptocurrency exchanges by trading volume in November 2023 are:

1. Binance
2. BIKA
3. Upbit
4. IndoEx
5. OKX

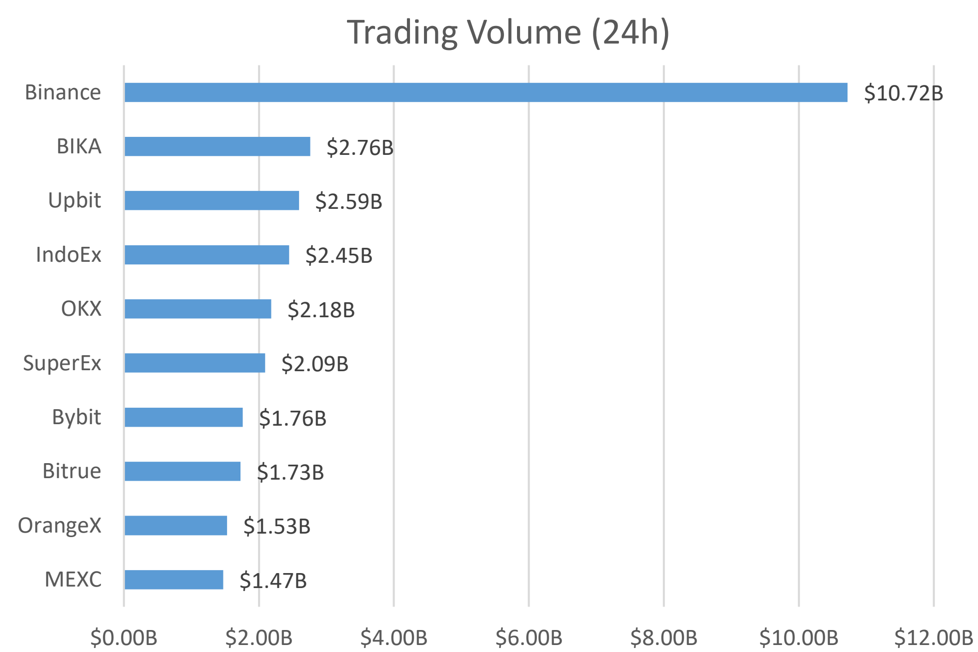

Figure 1 Top Cryptocurrency Exchanges by Trading Volume as of 27 November 2023 (CoinMarketCap, 2023)

1.7 Research Aim and Objectives

Predicting the price of cryptocurrency is essential for researchers to study the high volatility of these prices. As well as for investors and traders due to the potential significant profits it can yield. However, the volatility of cryptocurrency markets, particularly Bitcoin, makes it challenging for investors and traders to make informed decisions. This thesis explores the possibility of predicting Bitcoin prices using traditional time series analysis. The use of deep learning models requires a great deal of computing power and longer training times, making them infeasible for the broader range of researchers and analysts. Classical models, on the other hand, have a simpler structure that can be tuned, implemented, and interpreted easily. They are faster to train and less computationally intensive, which is useful when quick predictions are needed.

Primary research question:

- Can time series analysis be used to predict the future price of Bitcoin?

Secondary research questions:

- Which is the best combination of p, d, q that produces the lowest error values?
- What level of prediction accuracy can be achieved using only historical price data?
- What are the challenges and limitations of using traditional time series analysis to predict cryptocurrency prices?

The research aims to:

- Review the literature on time series analysis and Bitcoin price prediction.
- Develop a time series forecasting model for Bitcoin prices using historical data.
- Evaluate the performance and accuracy of different p, d, q values of the ARIMA model.

The introduction in Chapter 1 has established the context and significance of the research. Subsequently, Chapter 2 provides a comprehensive literature review of similar studies. While Chapter 3 explains the research methodology and Chapter 4 presents the data analysis and empirical findings. Chapter 5 critically analyzes and evaluates these findings. Finally, Chapter 6 concludes the thesis by highlighting the study results and limitations.

Chapter 2. Literature Review

Many studies explored the potential of machine learning techniques predicting the price of Bitcoin and other cryptocurrencies using various methods and a variety of features and technical indicators. Only a few considered the use of ARIMA time series to predict Bitcoin while the rest utilized deep learning methods.

Wirawan et al. (2019) found that the ARIMA method achieves high accuracy in short-term predictions, specifically the first and second future periods. Using training data from May 2013 to May 2019, the authors created three prediction scenarios of test data to predict the prices for the next 1-7 days. Consequently, it was found that the accuracy level decreased as more periods are predicted, and that (4,1,4) produced the lowest Mean Absolute Percentage Error (MAPE) in 2 out of 3 scenarios: 3.24 and 2.92.

Alahmari (2019) also concluded that shorter period predictions using ARIMA produces lower error values; that is, daily predictions produced the lowest Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE) than weekly and monthly prediction for three cryptocurrencies: Bitcoin, XRP, and Ethereum.

Hua (2020) applied ARIMA and Long Short-Term Memory (LSTM) methods to predict Bitcoin prices. The author did not provide any evaluation metrics for the ARIMA(1,1,0) model used. However, from the available figures in the paper, it is evident that “next single step” prediction followed the actual prices more closely than “next 5 steps”. On the other hand, LSTM reported to have an average error rate of 0.48 and a standard deviation of 2.09.

Fiaidhi et al. (2020) used 5 months data from August 2019 to December 2019 to predict Bitcoin closing price for the first seven days of 2020. The authors found that (8,1,0) produced the lowest MSE value of 170962.195. They concluded that prediction using only historical closing price results in high MSE values because of the volatility of Bitcoin prices. However, the results also show that the ARIMA model is useful in predicting sub-periods of time, in other words, shorter time periods.

In another study (Darley et al., 2021), ARIMA was again confirmed to be an effective tool for short-term forecasting of Bitcoin prices when using historical prices only. ARIMA(6,1,12) was found to be the most suitable model and gave near precise prediction values with an accuracy of 99.94% for the first seven days, 99.59% for the next 14 days, and 95.84% for the next 30 days. Once more, this study reinforces that the prediction accuracy for ARIMA decreases when used for longer forecast periods. It is worth noting that the authors in this paper used the terms "MAPE" and "accuracy" interchangeably even though they are opposites, and that the percentages labelled as "MAPE %" in the table could mean accuracy and not percentage error.

A different study (Hafid et al., 2022) used data that consisted of historical Bitcoin close prices and volume from February 2021 to February 2022, with a time step of 15 minutes, as well as technical analysis indicators. They employed the most common classification models such as Logistic Regression (LR), Support Vector Machine (SVM), Random Forest (RF), and Voting Classifier (VC) to predict the direction of the market by giving buy and sell signals. Evaluation metrics used were accuracy, precision, recall, and a k-fold comparison of the accuracy of the different proposed models. The RF model was chosen because it could handle the nature of Bitcoin data, and it performed the best with an accuracy of 88.4%.

Ranjan et al. (2022) employed statistical and machine learning models on Bitcoin price data, daily and 5-minute interval. The models were evaluated based on their accuracy, precision, recall, and F1-score. It was found that LR, with an accuracy of 64.8%, and Linear Discriminant Analysis (LDA) models showed better performance when dealing with daily data that has a lot of features. However, for high-frequency data, the Boost model performed better than all other machine learning algorithms, with an accuracy of 59.4%.

Nagula & Alexakis (2022) suggested the use of a hybrid model to predict the price of Bitcoin using data from February 2014 to September 2021, and 119 technical and fundamental features. Their research concluded that the hybrid model had superior performance compared to the Deep Cross Networks (DCN) regression model, with a 58% decrease in MAE and a 29% increase in directional hit rate.

In another study (Liu et al., 2021), the authors employed the Stacked Denoising Autoencoders (SDAE) to forecast Bitcoin's price for the next day. They evaluated the performance of SDAE by comparing it to benchmark methods such as the Back Propagation Neural Network (BPNN) and the Support Vector Regression (SVR). They used 1356 days of data along with some 40 features, taking into consideration various aspects of the cryptocurrency market, public interest, and macroeconomic conditions. SDAE achieved the lowest MAPE of 0.1019 and RMSE of 160.63, and the highest Directional Accuracy (DA) of 0.5985, while BPNN performed the worst.

Jay et al. (2020) used many technical features in their stochastic Multi-Layer Perceptron (MLP) and LSTM models, along with tweets volume and google trends. Using data ranging from 2017 to 2019, they trained the models on three cryptocurrencies' prices (Bitcoin, Ethereum, Litecoin) for the previous seven days to predict the price on the eighth day. The models were evaluated using Stochastic MAPE, MAE, RMSE, and MSE. Their study found that Stochastic Neural Networks outperformed regular neural networks.

Additionally, Mittal et al. (2019) explored the relation between Bitcoin prices with tweets volume and sentiment, and number of searches on Google. They used the average Bitcoin price for each day from April 2014 to January 2019, employing Linear Regression, Polynomial Regression, Recurrent Neural Network (RNN), and LSTM. The study assessed the RNN, LSTM, and ARIMA models for predicting Bitcoin's price using only historical price data first. The accuracy rates for predicting the direction of daily price change were 43.78% for RNN and 42.98% for LSTM models. In contrast, ARIMA performed the worst, with an accuracy rate of 38.02%. However, there was no mention of which ARIMA model was used. Furthermore, since tweet sentiment showed poor correlation with Bitcoin price, the authors applied LSTM, RNN, and Polynomial Regression on tweet volume and Google trends and were able to predict the direction of the price with an accuracy of 77.01% and 66.66%, respectively, using Polynomial Regression.

While Ji et al. (2019) applied deep learning methods such as Deep Neural Network (DNN), LSTM, Convolutional Neural Network (CNN), Deep Residual Network (ResNet), and combinations of these models to predict Bitcoin's price. In the study, LSTM performed slightly better than the other models for regression, while DNN performed better for classification. The performance of all proposed models was similar, except for ResNet as the data was insufficient for it to be effectively trained.

Moreover, LSTM was found to achieve a marginally better accuracy of 52% compared to RNN and ARIMA, while RNN achieved the lowest RMSE of 5%. ARIMA appeared to perform the worst in terms of both accuracy and RMSE but it is imperative to note that it was used to forecast 30 days into the future and the accuracy of the model was worse by only 2.7% compared to LSTM. The study suggests that non-linear deep learning methods are more accurate at predicting Bitcoin prices than the traditional ARIMA time series model when considering longer time periods (McNally et al., 2018).

A different study (Demir et al., 2018) examined the relationship between Bitcoin and economic policy uncertainty (EPU). Aiming to predict daily Bitcoin returns, the authors used the Bayesian Graphical Structural Vector Autoregressive model (BGSVAR) as well as the Ordinary Least Squares (OLS) and the Quantile-on-Quantile Regression (QQ) estimations. They found that as EPU increases, Bitcoin returns decrease. However, this effect is positive and significant at the lower and higher quantiles, which means that Bitcoin could be used as a hedging instrument against uncertainty in bullish market conditions.

Saad & Mohaisen (2018) used past prices data from April 2016 to December 2017 in addition to other related features to identify patterns and predict future Bitcoin prices. The prediction model was trained on Linear Regression, RF, and Gradient Boosting (GB). Each model was evaluated on accuracy, RMSE, and MAE. Linear Regression achieved the highest accuracy score of 99.4% and the lowest RMSE (0.0113) and MAE (0.006) scores.

Furthermore, Singh & Agarwal (2018) implemented multiple regression models such as Linear, Ridge, Lasso, Polynomial Regression as well as SVR, and K-Nearest Neighbour-based (KNN) Regression. Using various transactional features and data of eight years, they found that KNN was the most suitable model with an MSE of 0.00021.

Another study (Sin & Wang, 2017) explored the relationship between around 200 features of Bitcoin and its price fluctuations on the following day using a method called Genetic Algorithm based Selective Neural Network Ensemble (GASEN). The model was then compared to other trading strategies; Single MLP and previous day trend following. GASEN resulted in 85% returns, while the MLP model resulted in 53%, and trend following in 38% returns.

Azari (2019) used data for three years, from September 2015 to 2018, and found that ARIMA(8,8,1) achieved the lowest Residual Sum of Squares (RSS) of 0.002. At the same time, it produced an MSE that is almost 100x the minimum achievable MSE, through ARIMA(4,2,1), of which specific value is not reported in the study. The author concluded that using historical prices in prediction results in large MSE values due to the volatile nature of Bitcoin prices, specifically during the 2017-2018 period.

Abu Bakar & Rosbi (2017) applied ARIMA(2,1,2) on monthly Bitcoin prices, referred to as exchange rate in the paper, from January 2013 until October 2017. The model produced an R-squared value of 0.44443 and MAPE of 5.36%.

Benzekri & Özütler (2021) tested the model ARIMA(1, 1, 0) using quarterly data on the last two quarters of 2020 only. As a result, the model showed high accuracy with 4.24% MAPE and 0.46 RMSE.

Mudassir et al. (2020) modeled Bitcoin prices using Artificial Neural Network (ANN), Stacked Artificial Neural Network (SANN), SVM and LSTM utilizing regression and classification methods. The authors divided the forecasting periods into three intervals. It was evident that all regression models resulted in higher error rates when predicting daily closing price in the third interval, as it included the highest Bitcoin price volatility occurring after April 2017. Daily prediction in the third interval resulted in ANN producing the lowest MAE and RMSE values of 39.50 and 74.10, respectively. While SVM had the lowest MAPE value of 1.44. Meanwhile, predicting prices using the third interval for seven, thirty, and ninety days resulted in higher error values overall for all models as compared to predicting daily prices. This further confirms that Bitcoin prices are best predicted in shorter time periods. On the other hand, The SANN classification model resulted in the highest accuracy of 60%, while LSTM had the highest F1-score of 0.66, using the third interval.

Roy et al. (2018) used data from July 2013 to August 2017 to predict the next 10 consecutive days' Bitcoin prices. They compared ARIMA with AR and MA models and found that ARIMA resulted in 90.31% accuracy, while AR resulted in 89.24% and MA in 87.58% accuracy. The authors did not state which p, d, q variables were used in the ARIMA model.

Fernandes et al. (2021) applied deep learning prediction models; RNN, LSTM, and Gated Recurrent Unit (GRU) and used historical Bitcoin prices as well as sentiment data. However, after pre-processing, the sentiment data proved insufficient and gave invalid outputs so it was dropped. All three models produced high error values. The lowest MSE of 42147659.55 was achieved by LSTM, and the lowest MAE of 861.00 was achieved by GRU.

Additionally, Munim et al. (2019) analyzed daily prices from January 2012 to October 2018 using ARIMA and Neural Network Autoregression (NNAR) models. As a result, ARIMA(4,1,0) resulted in RMSE value of 0.038, MAPE of 0.379, and Mean Absolute Scaled Error (MASE) of 0.969. Notably, ARIMA was shown to be more accurate than NNAR in the test-sample forecasts, which include extremely volatile periods of Bitcoin price.

Alessandretti et al. (2018) predicted the daily price of 1,681 cryptocurrencies between January 2016 and April 2018, excluding Bitcoin, using three prediction methods. Two methods were based on Gradient Boosting Decision Trees (GBDT), while the third one was based on LSTM where the model used previous prices to predict future ones. Comparing the return on investment for each method, LSTM seemed to perform better for long-term predictions i.e., 50 days of data, while the GBDT methods performed better for short-term predictions, 5 or 10 days.

Chen et al. (2020) did dimension engineering on Bitcoin price data before leveraging any machine learning models. Statistical methods, LR and LDA, were utilized for the data with high-dimension features to predict daily Bitcoin prices. On the other hand, machine learning methods, RF, XGBoost, Quadratic Discriminant Analysis (QDA), SVM, and LSTM, were utilized for data with low-dimension features to predict 5-min interval prices. As a result, the statistical methods of LR and LDA performed better on daily price data, with an average accuracy of 65%. XGBoost turned out to perform the worst with 48.3% accuracy, while SVM achieved an accuracy percentage of 65.3% that is almost similar to the statistical methods. For the data with 5-min intervals, the machine learning methods performed better than the statistical methods. The authors suggest that ARIMA has the best precision of 100% on linear data, while their LR and LDA methods outperform it on accuracy.

Moreover, Peng et al. (2018) used data of three cryptocurrencies, Bitcoin, Ethereum and dash market price in USD, for the period between January 4th, 2016, and July 31st, 2017. Their low-frequency data constituted daily prices, while the high-frequency data constituted hourly prices. Overall, the results show that the RMSE and MAE of SVR-GARCH were the lowest compared to all nine GARCH benchmarks used in their paper, across all currencies and datasets.

Poongodi et al. (2020) applied the ARIMA model on Bitcoin data from April 2013 to July 2017. The authors compared the actual closing price vs. predicted closing price on daily basis and found that the model had an accuracy of 49%. However, they did not mention which parameters achieved such accuracy.

Yenidoğan et al. (2018) utilized ARIMA and PROPHET models to predict Bitcoin prices, using historical data between May 2016 and March 2018 in minute format. Along with the cryptocurrency's price data, the authors included additional variables to help improve the prediction accuracy. The prediction period used in both models was 90 days. From the results, it appears that ARIMA produced RMSE of 593.80, MAPE of 0.056 on the test set, and 68% precision. While PROPHET produced RMSE of 245.09, MAPE of 0.020, and 94.5% precision. There was no indication of which ARIMA model was chosen and evaluated in this paper, and it would have been more insightful if the authors applied their model on shorter time periods.

Moreover, when predicting the next 30 days using data between April 2013 and October 2017, Karakoyun & Çıbıkdiken (2018) found that ARIMA(4,2,1) produced RMSE of 1146.067 and MAPE of 11.86%. While LSTM produced RMSE of 93.27 and MAPE of 1.40%, confirming further that ARIMA should be used for shorter prediction periods instead.

The key takeaways from the literature review are:

- ARIMA is effective for short-term Bitcoin price prediction, especially for daily and weekly forecasts. The studies reported high accuracy in predicting the immediate future (1-7 days), whereas the accuracy decreased with longer forecast periods.
- Different studies used different ARIMA parameters, it is necessary to experiment with different configurations to find the most suitable model for Bitcoin price data.
- The evaluation metrics varied across studies, with most of them using accuracy rate and RMSE to evaluate their models, along with MAPE, MAE, and MSE.
- The literature mostly focused on long-term Bitcoin price prediction, which in our opinion does not suit the nature and sensitivity of Bitcoin prices. Especially since there are various influential factors that cannot all be considered in the prediction model. Therefore, historical prices seem to be the most reliable and consistent feature.

This thesis reviews the literature on time series analysis and Bitcoin price prediction, aims to develop a time series forecasting model for Bitcoin prices using historical data, and evaluates the performance and accuracy of different p, d, q values of the ARIMA model.

It attempts to answer the primary question, which is: can time series analysis be used to predict the future price of Bitcoin? Along with the secondary questions: which is the best combination of p, d, q that produces the lowest error values? What level of prediction accuracy can be achieved using only historical price data? What are the challenges and limitations of using traditional time series analysis to predict cryptocurrency prices?

Chapter 3. Research Methodology

The main method that will be used in this thesis is the ARIMA time series forecasting model, in order to predict Bitcoin prices using only historical data. The choice to use ARIMA is because it is easier to train, faster to implement, and easily interpretable, which makes it practical and feasible for inexperienced researchers, investors, or hobbyists. The downside is that ARIMA often performed worse than deep learning models, according to the literature. However, the majority in the literature that used ARIMA did not report the specific parameters that were tested in their studies, and it was often applied on longer time periods.

This thesis aims to use ARIMA time series on short-term data, determine the best combination of p, d, q parameters, as well as evaluate and compare the prediction accuracy and performance of different ARIMA models. This chapter will define and explain time series, its components and characteristics, provide some examples of time series data, explain the ARMA and ARIMA models, and the tools used in these models. Lastly, in order to assess the performance of any machine learning model, certain evaluation methods must be carried out, which will also be explained.

3.1 Time Series Analysis

Time series analysis is a statistical method that is used to analyze data points collected sequentially over a period of time. It is used in various fields such as finance, economics, and engineering to analyze patterns and trends in data over time. Time series is often used for financial data and is well suited to the dynamic nature of financial markets (Carmona, 2014).
Time series forecasting uses historical values of data to predict the future values. The use of time series in forecasting has various applications. These include predicting future retail sales, forecasting weather and temperatures, and estimating stock prices in financial markets. As well as predicting cryptocurrency prices, which is the focus of this research.

3.2 Time Series Components

The main components or characteristics of time series data are **trends**, **seasonality**, **cycles,** and residual **random/white noise**.

- **Trends** show the overall direction of the data, whether it is increasing, decreasing, or remaining stationary over a long period of time. For instance, the global temperature of the planet showing an upward trend since the late 19th century.

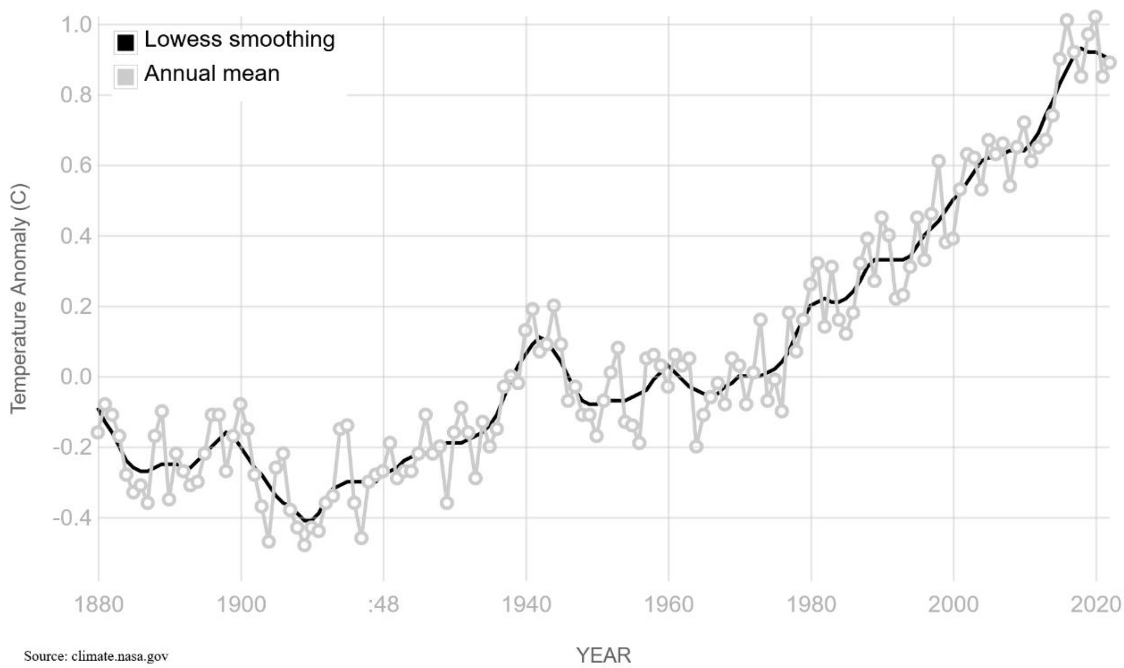

Figure 2 Global Land-Ocean Temperature Index (NASA GISS, 2022)

- **Seasonality** refers to recurring, predictable patterns that happen at fixed intervals within one year. They could occur within a week, a month, a year, or any other interval not greater than a year. These patterns are fixed in timing, direction, and volume. An example of seasonality is the increase in US natural gas consumption in winter and summer.

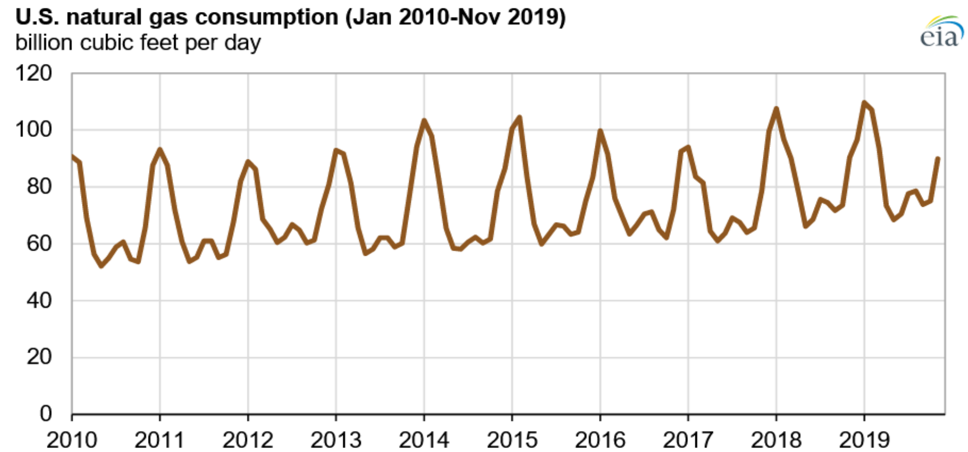

Figure 3 U.S. Natural Gas Consumption from January 2010 to November 2019 (US eia, 2020)

- **Cycles** are patterns that last for more than one year and repeat over several years. They do not have a consistent period i.e., one cycle might have a duration of 18 months, followed by the next cycle lasting 22 months. Economic expansions and recession, or business cycles that fluctuate between growth and decline are examples of cyclical patterns.

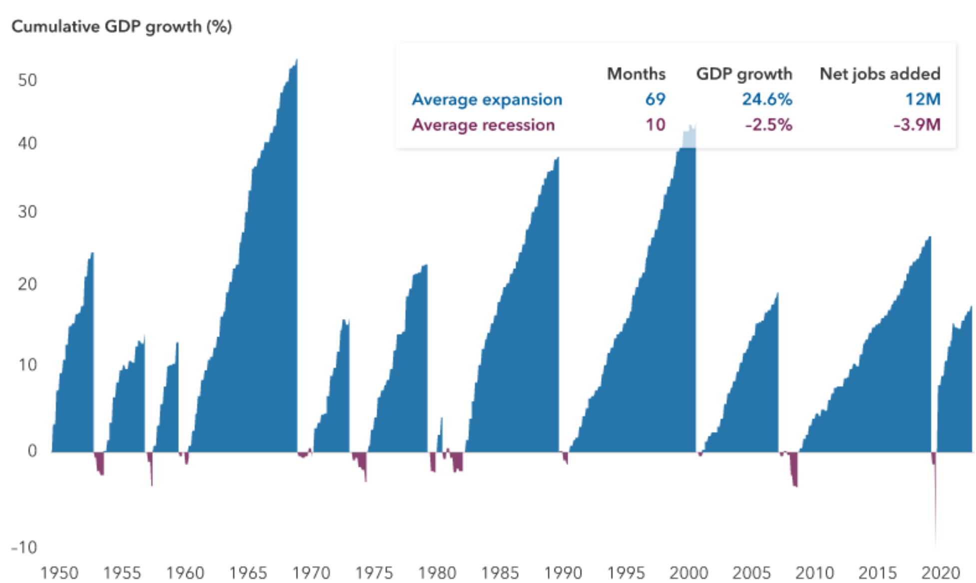

Figure 4 Cumulative GDP Growth during Expansions and Recessions (Franz, J. & Spence, D., 2023)

- Time series data also include the **noise** element, which is the remaining fluctuations in the data that are random and cannot be explained by trends, seasonality, or cycles. They are unpredictable and inconsistent.

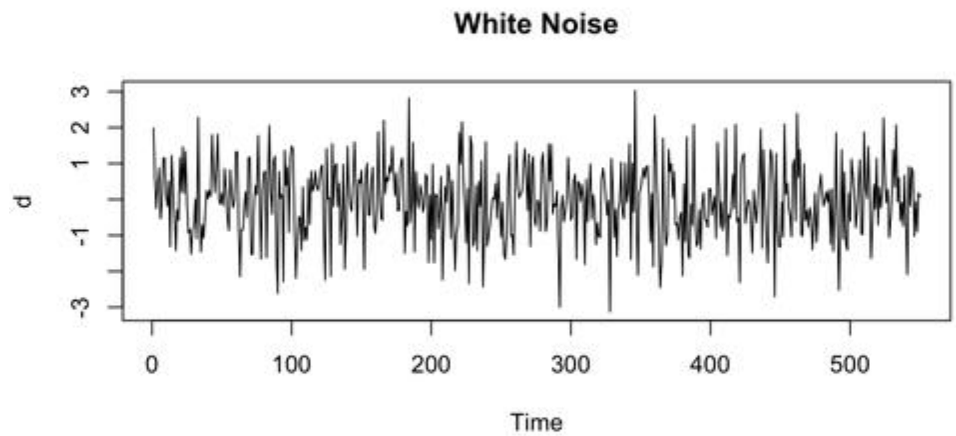

Figure 5 White Noise Example

3.3 Types of Time Series

There are several types of time series data:

1. **Univariate and Multivariate**: A univariate time series has one single variable that is recorded over a period of time, while in a multivariate time series, multiple variables are considered simultaneously over the same time period.
2. **Stationary and Non-Stationary**: A stationary time series has constant statistical properties such as a mean and variance that remain constant over time. On the other hand, a non-stationary time series has changing statistical properties over time, characterized by trends or seasonality.
3. **Continuous and Discrete**: A continuous time series is where the observations are collected and recorded continuously over time without interruptions, while a discrete time series is where the observations are recorded at distinct and separate points in time.
4. **Gaussian and Non-Gaussian**: In a Gaussian time series, the data follows a normal distribution. Meaning that the majority of the data is close to the mean, and the further away from the mean a value is, the less likely it is to occur. In a non-Gaussian time series, the data does not follow a normal distribution but instead may be skewed or heavy-tailed.

Figure 6 of the UAE's annual GDP growth is an example of a univariate, non-stationary, discrete time series.

Figure 6 Annual Percent Change in UAE GDP (World Bank – processed by Our World in Data, 2023)

Figure 7 of the UAE production trend is an example of a multivariate, discrete time series, with some stationary and non-stationary variables.

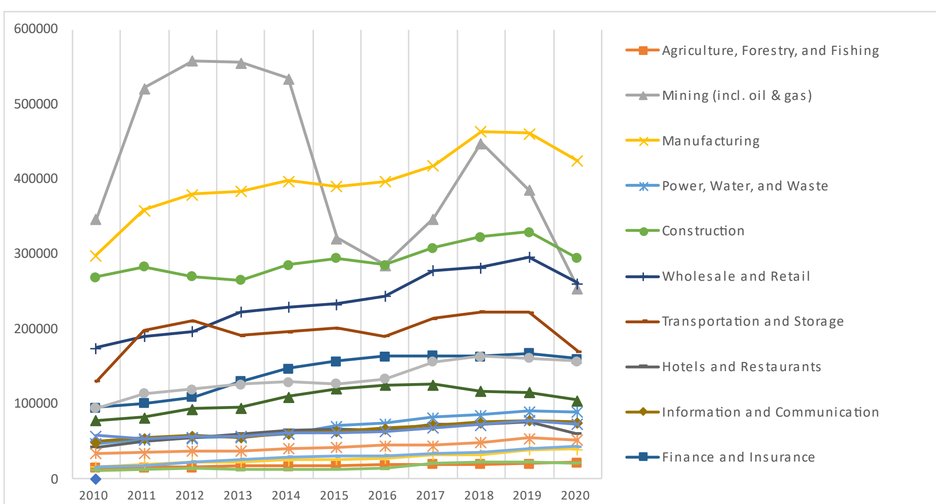

Figure 7 UAE Production Trend by Economic Sector (FCSC, 2021)

Figure 8 of Bitcoin prices is an example of a univariate, non-stationary, continuous timeseries.

Figure 8 Bitcoin price fluctuations 2015-2024 (Investing.com)

3.5 ARMA Model

Two common techniques used in time series analysis are ARMA and ARIMA. The Autoregressive Moving Average (ARMA) assumes the data is stationary, and is a combination of two other simpler methods:

- Autoregressive (AR): This method predicts future values based on past values. For example, if a stock price has been increasing for the past few days, then it is likely to increase tomorrow as well.
- Moving Average (MA): This method checks the error of past predictions. If the model has been continuously overestimating the stock's price for the past few days, then it would adjust the prediction for tomorrow accordingly.

3.6 ARIMA Model

The Autoregressive Integrated Moving Average (ARIMA) model is an extension of ARMA that includes an extra step called 'differencing' to make the time series stationary in case it was not. A time series is stationary if its statistical properties, like the mean and variance, do not change over time. If a time series is not stationary, exhibiting trends or seasonality, it would be difficult for the model to differentiate between the underlying structure of the data and the random noise. The ‘differencing’ step helps to make the time series stationary by removing these trends or patterns, ensuring the model can understand the structure of the data and make accurate forecasts. The ‘integrated’ part (I) in ARIMA refers to this differencing step.

The ARIMA model is defined by three parameters: *p*, *d*, and *q*.

- *p* is the number of past observations that the model will consider. If p is 3, then the model will consider the prices from the past 3 days. This is the AR component and is expressed as:

$$X_t = \phi_1 X_{t-1} + \phi_2 X_{t-2} + \ldots + \phi_p X_{t-p} + Z_t$$

$$X_t = \sum_{i=1}^{p} \phi_i \, X_{t-i} + Z_t$$

- *d* is the degree of differencing, which is the number of times needed to subtract the previous day’s price from today’s price to make the data consistent at all points in time. If the time series is already stationary, then d=0. Otherwise, it could be 1; immediate previous observation, or 2; the past two observations etc. This is the *I* component.
- *q* is the number of past errors in prediction. If the predictions were inaccurate by a certain amount in the past few days, the model will use this information to adjust the next prediction. This is the *MA* component and is expressed as:

$$X_t = Z_t + \theta_1 Z_{t-1} + \theta_2 Z_{t-2} + \ldots + \theta_q Z_{t-q}$$

$$X_t = Z_t + \sum_{i=1}^{q} \theta_i \, Z_{t-i}$$

The ARMA(p, q) model combines the AR(p) and MA(q) models:

$$X_t = \sum_{i=1}^{p} \phi_i X_{t-i} + Z_t + \sum_{i=1}^{q} \theta_i Z_{t-i}$$

And the ARIMA(p, d, q) model includes differencing:

$$\nabla^d X_t = \sum_{i=1}^{p} \phi_i \nabla^d X_{t-i} + Z_t + \sum_{i=1}^{q} \theta_i Z_{t-i}$$

Where:

- X_t is the time series X at a time t, and X_{t-i} is X at a previous point in time.
- Z_t is the error terms or white noise, and Z_{t-i} is Z at a previous point in time.
- ϕ are the parameters of the autoregressive part of the model.
- θ are the parameters of the moving average part of the model.
- $\nabla^d X_t$ is the differenced time series (*d* times).

3.7 Time Series Stationarity

It is important for a time series to be stationary for models like ARMA and ARIMA to produce accurate results, particularly when dealing with financial data like Bitcoin prices. As mentioned earlier, a stationary time series has statistical properties that do not depend on the time at which the data point is observed, having a constant mean and variance that do not change over time. There are some specific methods that were developed to assess the stationarity of a time series. The most significant and widely used method is the Augmented Dickey-Fuller Test.

3.7.1 Augmented Dickey-Fuller (ADF) Test

The Augmented Dickey-Fuller (ADF) test is a statistical test used to assess the stationarity of a time series. The null hypothesis of the test assumes that the time series has a unit root, indicating non-stationarity i.e. it has some time-dependent structure. The alternate hypothesis, which is rejecting the null hypothesis, suggests that the time series is stationary (Dickey & Fuller, 1981). If the p-value produced from the test is less than a specific significance level e.g. 0.05, the null hypothesis will be rejected, meaning that the time series is stationary. Otherwise, if the p-value turns out to be more than 0.05, that would mean the time series is non-stationary and differencing must be used.

3.8 ACF and PACF

The Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) are tools used to identify the relationship between data points in a time series. They help in understanding the patterns and correlations within the data, which is essential for building accurate forecasting models.

- The **ACF** measures the correlation between observations of a time series at different points in time i.e., how much a data point in a time series is related to other data points separated by a certain number of periods (these periods are called 'lags'). It can, for example, indicate how many lags (previous days) should be considered when predicting tomorrow's price. In the context of ARIMA, it helps identify the value of *q* that is needed.
- The **PACF** is similar to ACF, but it only considers the correlation between points separated by multiple lags, without the effect of the lags in between. For example, it can show how much today's price is related to the price 4 days ago, without considering the prices of the 3 days in between. In the context of ARIMA, it helps identify the value of *p* that is needed.

3.9 Evaluation Methods

Time series model performance can be evaluated using several metrics such as the Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE).

- Root Mean Square Error (RMSE) measures the differences between values predicted by a model and the values observed. It is the square root of the average of squared differences between predicted and actual observation. It is beneficial in penalizing large errors, because squaring the errors gives more weight to large errors.

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(Y_i - \hat{Y}_i)^2}$$

- Mean Absolute Error (MAE) is the average of the absolute differences between the predicted and actual observations. It measures the average magnitude of the errors in a set of forecasts, without considering their direction. It is less sensitive to outliers compared to RMSE, making it a good measure of accuracy when the dataset contains outliers.

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} \mid Y_i - \hat{Y}_i \mid$$

- Mean Absolute Percentage Error (MAPE) is the average of the absolute percent difference between the actual and the predicted values. It expresses the size of the error as a percentage.

$$\text{MAPE} = \frac{100\%}{n}\sum_{i=1}^{n} \mid \frac{Y_i - \hat{Y}_i}{Y_i} \mid$$

Where, for all metrics,

- Y_i is the actual value.
- $\hat{Y}_i$ is the predicted value.
- n is the total number of observations.

In model prediction, the goal is to minimize the error, so lower values of all these metrics indicate better model performance.

3.10 Methodology Outline

To summarize the methodology to be used in this thesis:

1. Graphically visualize the time series data to better understand its structure.
2. Test the data for stationarity.
3. Calculate the autocorrelation (ACF) and partial autocorrelation (PACF) coefficients and present them in graphs to further examine the time series data.
4. Use differencing to convert the data to stationary if it turns out to be non-stationary.
5. Apply ARIMA model with different combinations of p, d, q values.
6. Report the process of finding the appropriate ARIMA combination.

7. Evaluate the different parameters using RMSE, MAE, and MAPE, where the combination with the lowest error values is considered to be the most accurate and the best for prediction.

Figure 9 illustrates the outline of the methodology.

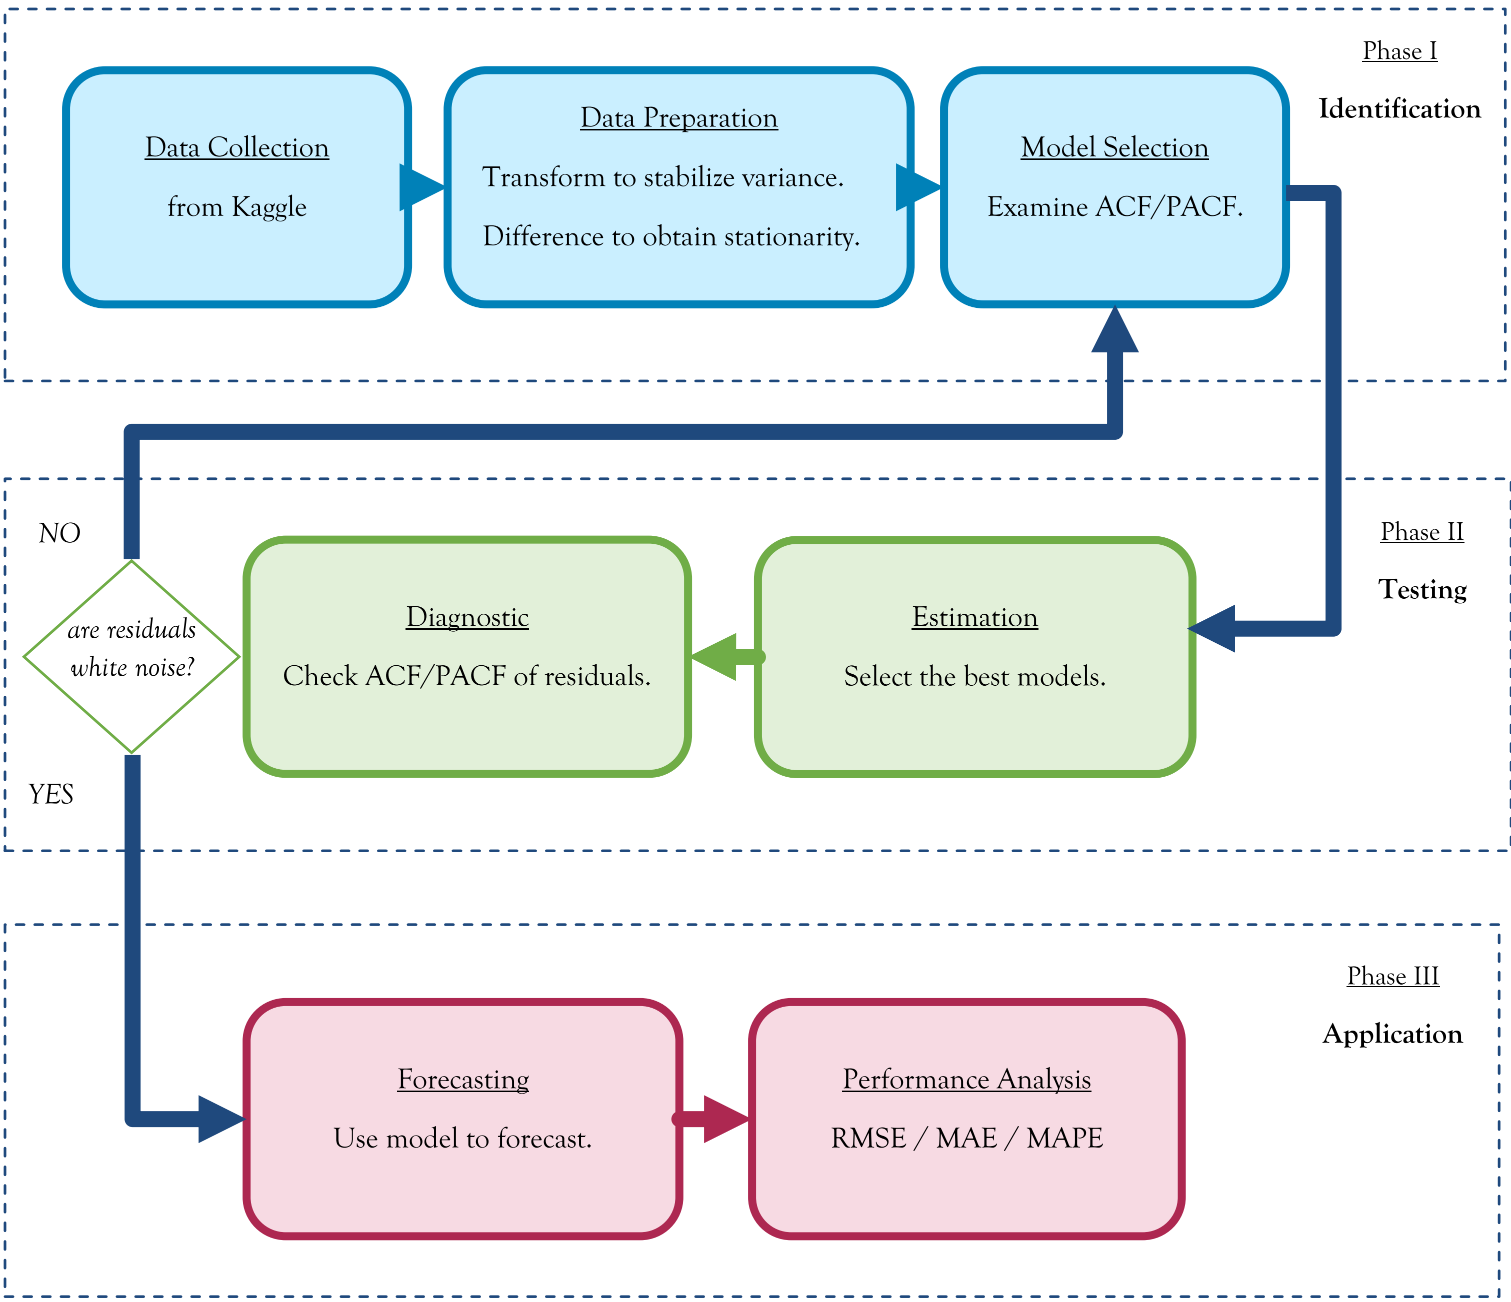

Figure 9 ARIMA Time Series Methodology Outline

Chapter 4. Findings and Data Analysis

In order to forecast future Bitcoin prices, historical prices for the time period 17/09/2014 to 17/09/2023 obtained from Kaggle will be used as the main data set: https://www.kaggle.com/datasets/arslanr369/bitcoin-price-2014-2023.

The dataset contains 3,228 observations for the following 7 variables:

- Date: the date of each observation.
- Open: Bitcoin's opening price for the day.
- High: the highest price of the day.
- Low: the lowest price of the day.
- Close: the closing price for the day.
- Adj. Close: the adjusted closing price.
- Volume: the number of transactions within the day.

4.1 Data Exploration

The first step to start the analysis is to load the data into the programming environment, and then load the necessary libraries that will help in the analysis. It is also worth noting that there were no missing values in the dataset.

The only variables of interest in this dataset are *Date* and *Adj. Close*, where they represent the date (time) of each data point and their respective adjusted closing price, which is the target variable for forecasting. The other variables were removed to simplify the data and the visualization, which is an important step to better understand the data. The graph below obtained from the data shows the fluctuations in the price of Bitcoin.

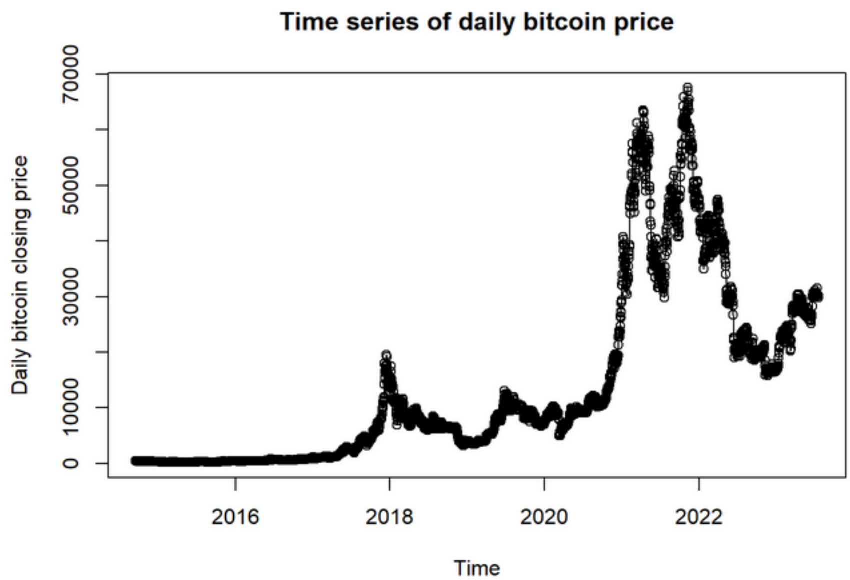

Figure 10 Daily Bitcoin Price Time Series

It is evident from **Figure 10** that Bitcoin prices have a general upward trend but with sharp fluctuations and no obvious seasonality or predictable pattern. There are periods of relative stability e.g. between 2016 and late 2017, and 2019 and 2020. There are also periods of high volatility such as late 2017 where the price experienced a significant spike, as well as between late 2020 and mid 2021 where the price sharply increased and then started to decline after reaching its peak. From an initial observation, this time series appears to be non-stationary as both the mean and variance kept changing over time. This indicates that there might be a need to transform the data by differencing to make it stationary, but first, a formal stationary test must be performed to confirm the assumption.

Another data exploration process is checking for correlations. **Figure 11** shows that there is a strong correlation between the current day's closing price and the previous day's closing price.

This means that if Bitcoin price was high today, it is likely to be high tomorrow as well and vice versa, suggesting a dependency of the price on its previous day's value.

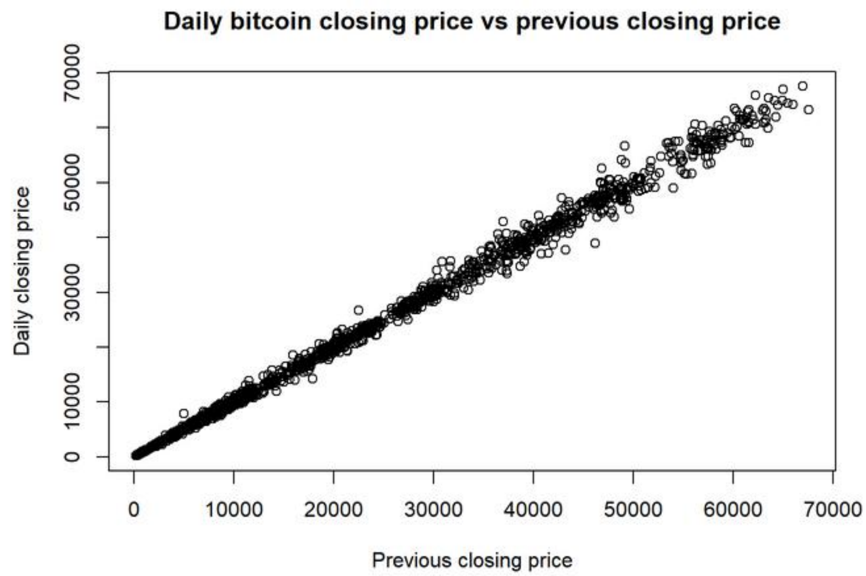

Figure 11 Correlation of Daily Closing price with Previous Closing Price

4.2 Stationarity Check

As mentioned, the time series should be stationary in order to get the most accurate results. The graphs below show the components that make up the Bitcoin time series. It is evident from **Figure 12** that the data has a strong upward trend, but it could however shift downward in the future. In addition, there is strong and clear seasonality in the time series, which appeared only after decomposing. These characteristics indicate that this time series data is not stationary.

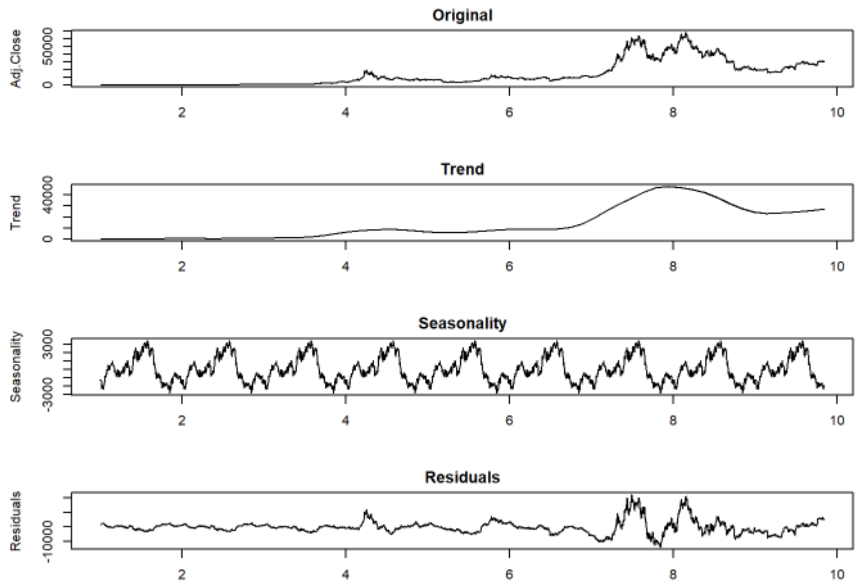

Figure 12 Time Series Decomposition

To formally test the observation that this time series is stationary, the Dickey-Fuller test (ADF) must be performed.

```
par(mfrow=c(1,1))
adf.test(bitcoin)
```

```
##
##  Augmented Dickey-Fuller Test
##
## data:  bitcoin
## Dickey-Fuller = -2.2649, Lag order = 14, p-value = 0.4662
## alternative hypothesis: stationary
```

Figure 13 Augmented Dickey-Fuller Test on Bitcoin Data

The ADF test shown in **Figure 13** reveals that the time series is in fact non-stationary, since the **p-value 0.4662** is higher than 0.05, so there is not enough evidence to reject the null hypothesis.

4.3 Autocorrelation Check

For further examination of the data, the ACF and PACF of data points is carried out. The ACF plot in **Figure 14** shows a gradual decline, indicating a strong correlation at the beginning that weakens as time goes on. This confirms the correlation graph above that showed a strong positive correlation between the current day's closing price and the previous day's closing price. The PACF plot shows a significant spike at lag 1, which also confirms the observation from the previous graphs that there is strong influence of previous day's price on the current price.

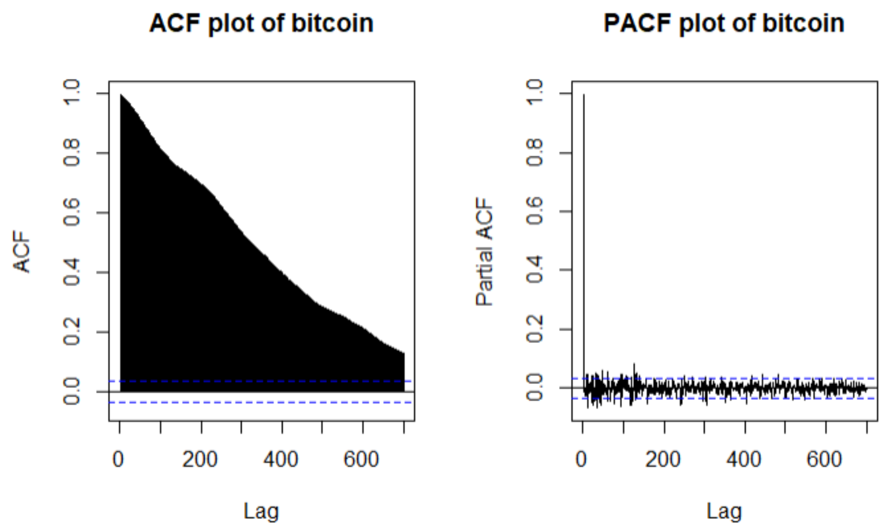

Figure 14 Autocorrelation and Partial Autocorrelation Plots

4.4 Normality Check

The next step in time series analysis is to check the normality of the residuals, as many prediction models like ARIMA require the residuals or errors to be normally distributed. However, before checking for normality, log transformation of the data was performed to stabilize the variance and make the data more suitable for modeling, since the raw data of Bitcoin prices exhibited periods of high volatility. Log transformation also helps reduce the impact of extreme values or outliers and managing drastic price changes.

```
bitcoin.log = log(bitcoin)
qqnorm(bitcoin.log)
qqline(bitcoin.log, col = 2)
```

Figure 15 Log Transformation of Bitcoin Data

4.4.1 QQ Plot

For a perfect normal distribution, the points in a QQ plot should lie along the straight line (in red). However, the QQ plot in **Figure 16** shows that the log-transformed Bitcoin prices closely follow the line but deviate at the tails, meaning that extreme price changes are more common in this data. So, the log transformation did not improve the normality of the residuals.

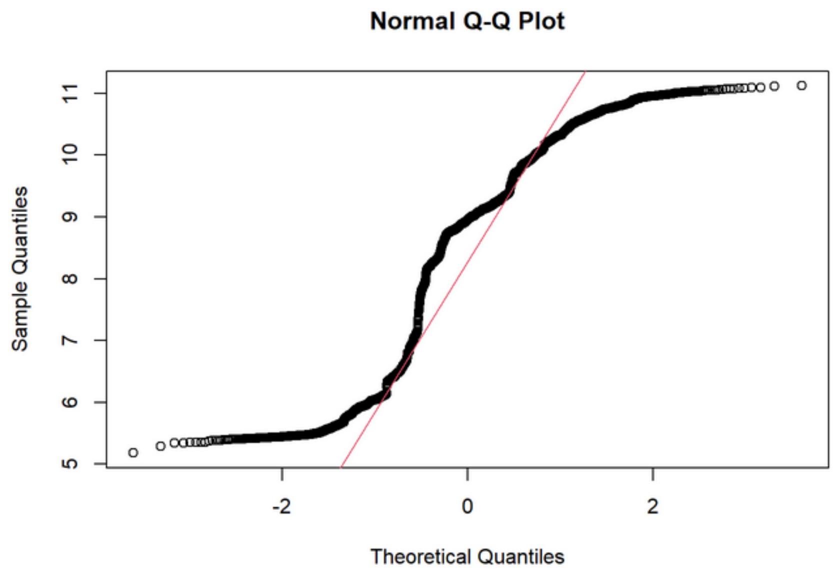

Figure 16 QQ Plot Normality Plot

4.4.2 Shapiro-Wilk Test

Another test for normality is the Shapiro-Wilk test, which in this case resulted in a p-value of practically zero (less than 2.2e-16). This means that the null hypothesis is rejected, also suggesting that the data does not follow a normal distribution.

```
shapiro.test(bitcoin.log)
```

```
##
##  Shapiro-Wilk normality test
##
## data:  bitcoin.log
## W = 0.91085, p-value < 2.2e-16
```

Figure 17 Shapiro-Wilk Normality Test

4.5 Differencing

The next step is to convert the data to stationary using differencing. As was apparent from the autocorrelation graphs in **Figure 14**, consecutive observations (lag 1) directly influence each other, and so the first difference of the log-transformed data was used to make the data stationary.

```
diff.bitcoin.t = diff(bitcoin.log ,differences = 1)
plot(diff.bitcoin.t,type='l', ylab='Daily bitcoin closing price', main = "Time Series plot of the first difference")
```

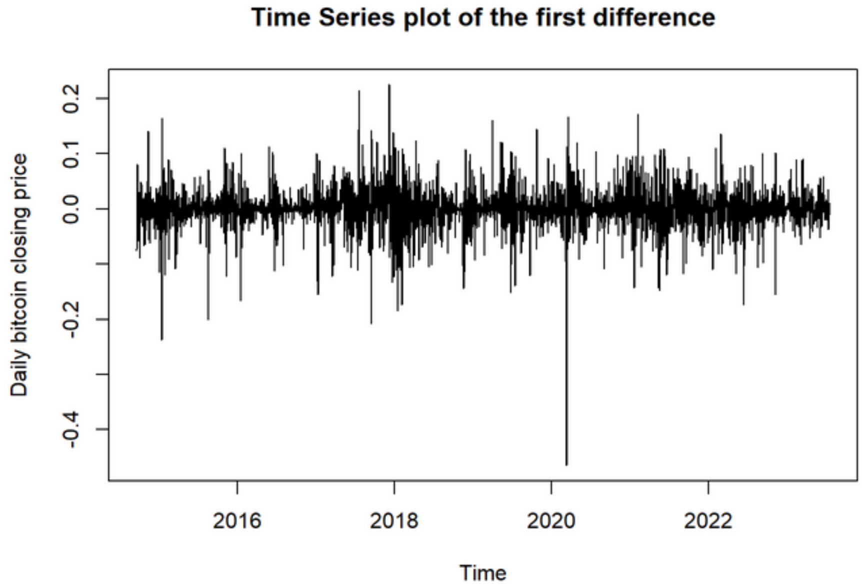

Figure 18 First Difference Time Series Plot

The graph in **Figure 18** shows that there is no obvious trend after first differencing, but random fluctuations throughout the period.

4.5.1 Stationarity Check after Differencing

Using the ADF test again to check for stationarity after differencing shows that the **p-value** is now **0.01** < 0.05. This means the null hypothesis, that the first difference of the log-transformed Bitcoin prices has a unit root, is rejected, and the differenced data is considered stationary.

```
adf.test(diff.bitcoin.t)
```

```
##
##  Augmented Dickey-Fuller Test
##
## data:  diff.bitcoin.t
## Dickey-Fuller = -13.988, Lag order = 14, p-value = 0.01
## alternative hypothesis: stationary
```

Figure 19 Augmented Dickey-Fuller Test on Differenced Bitcoin Data

4.6 ARIMA Model Selection

The process of selecting an appropriate ARIMA model involves first identifying the order of differencing (d), which was already established in previous sections. After that, the order of the

AR (p) and MA (q) parameters need to be determined. There are several methods that can help in the selection of the ARMA parameters.

4.6.1 Extended Autocorrelation Function (EACF)

The Extended Autocorrelation Function (EACF) helps identify the order of the model by producing a table of suggested autocorrelations that can be used to determine the AR and MA parameters. It helps identify the order of an ARMA model (p,d,q), where the AR order is determined by the row number and the MA order is determined by the column number at that point. An ‘o’ in the table indicates a significant autocorrelation at that particular combination of AR and MA orders, while an ‘x’ indicates a non-significant autocorrelation.

```
## AR/MA
##   0 1 2 3 4 5 6 7 8 9 10 11 12 13
## 0 o o o o o x o o o x o  o  o  o
## 1 x o o o o x x o o x o  o  o  o
## 2 x x o o o o o o o o o  o  o  o
## 3 x x o o o x o o o o o  o  o  o
## 4 x o x o o o o o o o o  o  o  o
## 5 x x x o x o o o o x o  o  o  o
## 6 x x x x o x o o o x o  o  o  o
## 7 x x x x o x x o o o o  o  o  o
```

Figure 20 Extended Autocorrelation Function Results for Model Selection

The possible p (AR order) and q (MA order) combinations proposed by the EACF as shown in **Figure 20** are: ARIMA(4,1,1), ARIMA(3,1,2), ARIMA(2,1,2), ARIMA(1,1,2), ARIMA(1,1,1), ARIMA(0,1,2), ARIMA(0,1,1), and ARIMA(0,1,0).

4.6.2 Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC)

The Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) also help by comparing different chosen models. Both methods try to balance model complexity and goodness of fit, so the model with the lower AIC or BIC score would be preferred.

Looking at the lowest AIC and BIC values of specific AR and MA orders in **Figure 21**, both criteria recommend the same set of parameters, which are ARIMA(2,1,2) and ARIMA(6,1,2).

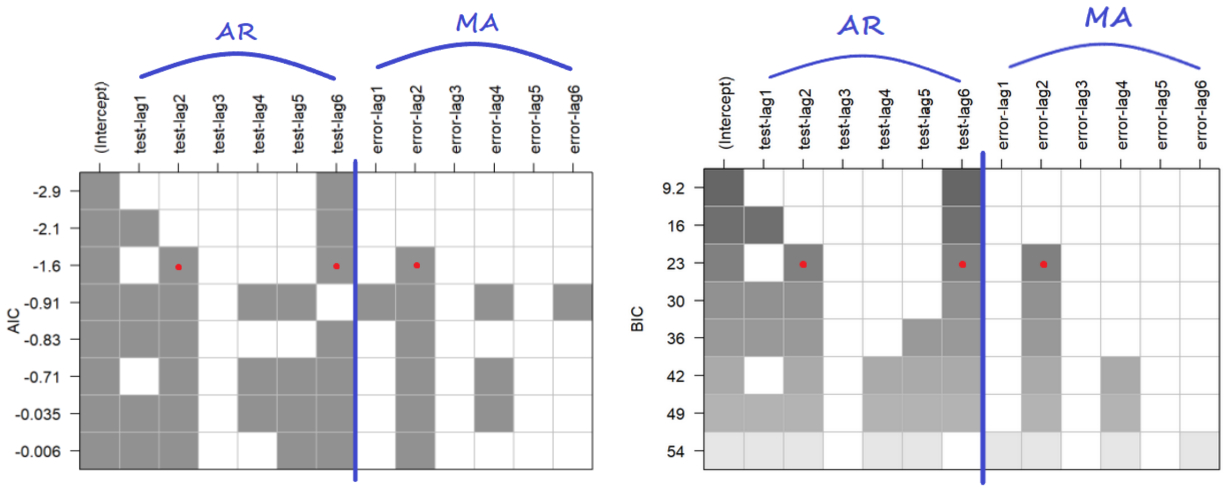

Figure 21 AIC and BIC Values for Model Selection

The BIC and AIC scores were also calculated for all recommended models, in which ARIMA(6,1,2) had the lowest AIC score and ARIMA(0,1,0) had the lowest BIC score, and ARIMA(0,1,1) coming in second for both, as shown in **Figure 22**.

```
aic_values <- sapply(list(model_010, model_011, model_012, model_111, model_112, model_212, model_312, model_411, model_612), AIC)
bic_values <- sapply(list(model_010, model_011, model_012, model_111, model_112, model_212, model_312, model_411, model_612), BIC)

aic_values
bic_values

[1] 52183.16 52183.13 52184.79 52185.11 52186.80 52188.81 52187.69 52185.80 52148.76
[1] 52189.24 52195.29 52203.02 52203.35 52211.12 52219.20 52224.17 52222.27 52203.47
```

Figure 22 AIC and BIC Scores for ARIMA Models

4.6.3 Conditional Sum of Squares (CSS) and Maximum Likelihood (ML)

The conditional sum of squares (CSS) and Maximum Likelihood (ML) methods are used to assess the parameters of a model. CSS minimizes the difference between the predicted value and the actual observed value by adjusting the model parameters in a way that makes the sum of squared differences as small as possible. While ML, on the other hand, maximizes the likelihood that a model's parameters will generate the actual observed value.

CSS and ML were applied on the nine ARIMA models suggested by EACF, and AIC and BIC. However, for the sake of simplicity, and considering that ML is often more accurate in assessing the parameters of ARIMA (Di Gangi, et al., 2022), only ML results are reported and compared in this thesis.

```r
# ARIMA (6,1,2)
model_612 <- arima(bitcoin, order = c(6,1,2), method = 'ML')
coeftest(model_612)
```

```
##
## z test of coefficients:
##
##       Estimate Std. Error  z value  Pr(>|z|)
## ar1 -0.1990905  0.0195793 -10.1684 < 2.2e-16 ***
## ar2 -0.9512583  0.0225076 -42.2638 < 2.2e-16 ***
## ar3 -0.0034887  0.0244952  -0.1424  0.886745
## ar4  0.0438404  0.0244910   1.7901  0.073444 .
## ar5  0.0432529  0.0182648   2.3681  0.017880 *
## ar6  0.0579368  0.0182811   3.1692  0.001528 **
## ma1  0.1735886  0.0088126  19.6977 < 2.2e-16 ***
## ma2  0.9678522  0.0140835  68.7226 < 2.2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```r
model_612_CSS <- arima(bitcoin, order = c(6,1,2), method = 'CSS')
coeftest(model_612_CSS)
```

```
##
## z test of coefficients:
##
##        Estimate  Std. Error  z value  Pr(>|z|)
## ar1 -0.46142900  0.02088904 -22.0895 < 2.2e-16 ***
## ar2 -0.94824746  0.02424563 -39.1100 < 2.2e-16 ***
## ar3 -0.00028294  0.02554822  -0.0111 0.9911637
## ar4  0.04898821  0.02554810   1.9175 0.0551758 .
## ar5  0.02440989  0.01985892   1.2292 0.2190100
## ar6  0.06195167  0.01816121   3.4112 0.0006468 ***
## ma1  0.43924995  0.01154587  38.0439 < 2.2e-16 ***
## ma2  0.95913736  0.01487297  64.4886 < 2.2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Figure 23 CSS and ML Application on One ARIMA Model

Based on the results of ML model testing, the top model with the most statistically significant coefficients is ARIMA(6,1,2).

4.8 Forecasting

4.8.1 Residual Analysis

It is natural for different criteria to suggest different models, as each of these methods has its strengths and weaknesses and may emphasize different aspects of the data. Therefore, the next step would be to check the residuals of each of the top models.

A good ARIMA model should have residuals that look like random/white noise. This indicates that the model has successfully captured all relevant predictors in the data, and that the remaining fluctuations or residuals are random and have no predictable pattern.

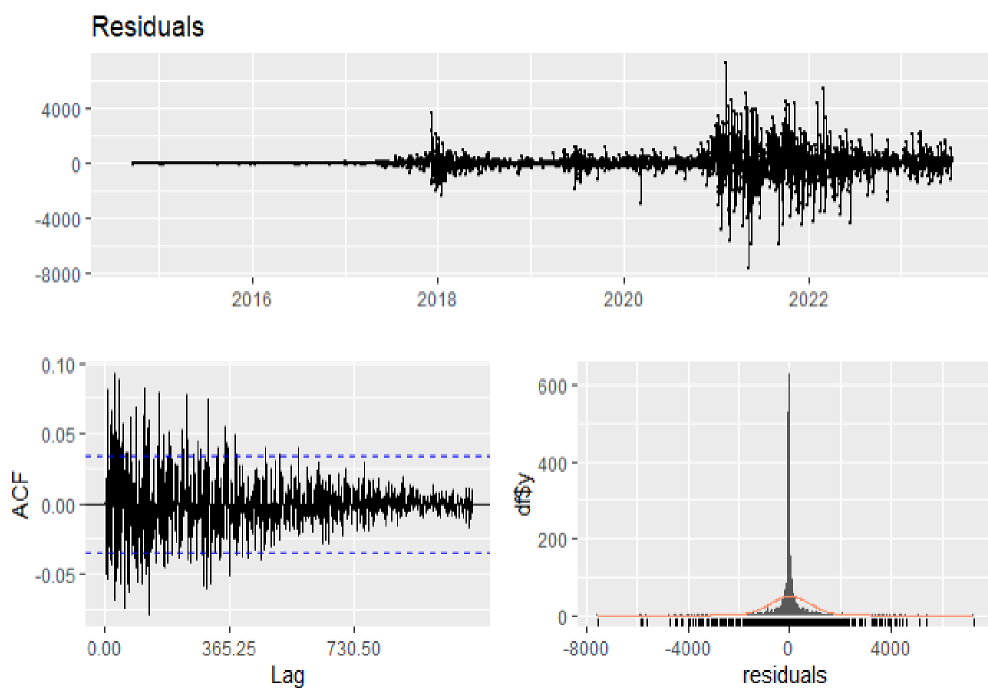

Figure 24 Residual Plots for One ARIMA Model

From the plots of residuals in **Figure 24**, it is evident that they are randomly distributed around zero and do not follow a specific pattern or trend. However, the ACF plot shows significant autocorrelations at various lags, and the histogram plot shows that it is not normally distributed.

4.8.2 Model Evaluation

All the models that were tested produced very similar plots. The next step is to evaluate the performance of the models using RMSE, MAE, and MAPE, where the combination with the lowest error values would be the most accurate and best for prediction.

	RMSE_Train	MAE_Train	MAPE_Train	RMSE_Test	MAE_Test	MAPE_Test
	<dbl>	<dbl>	<dbl>	<dbl>	<dbl>	<dbl>
ARIMA 6, 1, 2	0.03940042	0.02562542	Inf	0.03099448	0.02121623	123.7428
ARIMA 4, 1, 1	0.0394668	0.02560845	Inf	**0.03099013**	**0.0212102**	123.5946
ARIMA 3, 1, 2	0.03949352	0.02559431	Inf	0.03109535	0.02133982	139.8251
ARIMA 2, 1, 2	0.0394733	0.0256201	Inf	0.03099186	0.02121021	**123.5394**
ARIMA 1, 1, 2	0.03949498	0.02560444	Inf	0.0310946	0.02133874	139.7851
ARIMA 1, 1, 1	0.03949503	0.02559021	Inf	0.03110068	0.02134641	140.5897
ARIMA 0, 1, 2	0.03949517	0.02559036	Inf	0.03110073	0.02134648	140.5901
ARIMA 0, 1, 1	0.03948373	0.02565543	Inf	0.03099368	0.02121242	123.5551
ARIMA 0, 1, 0	0.05644119	0.03817685	Inf	0.05820128	0.05140668	1221.0308

Figure 25 ARIMA Models Evaluation

The dataset was split into a training set and a test set in which 80% of the observations were in the training set, while 20% were in the test set. Evidently from **Figure 25**, the training and test error rates are similar, indicating that the models performed well on both datasets and will likely perform well on new data as well. The model ARIMA(4,1,1) produced the lowest RMSE of 0.03099013, MAE of 0.0212102, and third lowest MAPE of 123.595. Model ARIMA(2,1,2) produced the lowest MAPE of 123.539, similar MAE to ARIMA(4,1,1), and second lowest RMSE of 0.03099186.

Chapter 5. Discussion

The main objective of this thesis was to review the literature on time series analysis and Bitcoin price prediction first and foremost, then develop a time series forecasting model for Bitcoin prices using only historical data, and lastly evaluate the accuracy of different models, highlighting the best one.

To reiterate, the primary research question for this thesis was whether time series analysis can be used to predict future prices of Bitcoin. The secondary research questions were: which is the best combination of p, d, q that produces the lowest error values? What level of prediction accuracy can be achieved using only historical price data? And, what are the challenges and limitations of using traditional time series analysis to predict cryptocurrency prices?

From the results, it is clear that time series analysis could in fact be used to on Bitcoin data even though it is characterized by high volatility and uncertainty. Secondly, according to the experiment, the best combination of ARIMA parameters with the lowest RMSE test set value is ARIMA(4,1,1). Thirdly, considering that cryptocurrencies' prices, specifically Bitcoin's, are influenced by various factors and not only their historical values, only a MAPE of 123% could be achieved using this type of data alone. A significant limit of traditional time series modeling techniques is that they only consider past values to predict future ones, disregarding other factors or features that could have an influence on future values. It is safe to assume that the past alone does not always predict the future, in the context of cryptocurrency.

Chapter 6. Conclusion

6.1 Conclusion

The research compared different ARIMA models highlighting the best model to forecast future Bitcoin prices. After analyzing the research results, we found that cryptocurrencies are challenging to predict due to their unique features. The ARIMA technique provided valuable insights regarding the characteristics of Bitcoin time series, which may not be easily interpretable by other prediction models. The literature suggests that ARIMA is accurate in forecasting time series data in the short-term, which was also determined by our experiment. The ARIMA(4,1,1) model yielded a satisfactory RMSE of 0.03099013, MAE of 0.02121, and MAPE of 123.595. Model ARIMA(2,1,2) also closely followed real Bitcoin prices with an RMSE of 0.03099186, same MAE value as ARIMA(4,1,1), and lowest MAPE of 123.539.

6.2 Recommendations and Future Work

While the objectives of this thesis were achieved, there are certainly some aspects that could be further studied in the future. One such aspect is exploring the correlation of Bitcoin with other cryptocurrencies, which could provide useful insights into this market's dynamics. Additionally, incorporating other related data such as regulatory changes, economic factors, number of transactions etc. could enrich the analysis and enhance the predictive capabilities of machine learning models. It could also be beneficial to explore implementing mixed approaches, combining ARIMA with CNN or LSTM models to predict cryptocurrency prices. This approach may enhance prediction accuracy by utilizing the strengths of both traditional time series and deep learning techniques.

References

1. Abu Bakar, N. & Rosbi, S., 2017. Autoregressive Integrated Moving Average (ARIMA) Model for Forecasting Cryptocurrency Exchange Rate in High Volatility Environment: A New Insight of Bitcoin Transaction. *International Journal of Advanced Engineering Research and Science,* 4(11), pp. 130-137.

2. Agrawal, R. K. & Adhikari, R., 2013. An Introductory Study on Time Series Modeling and Forecasting. *LAP Lambert Academic Publishing arXiv:1302.6613.*

3. Alahmari, S. A., 2019. Using Machine Learning ARIMA to Predict the Price of Cryptocurrencies..

4. Alessandretti, L., ElBahrawy, A., Aiello, L. M. & Baronchelli, A., 2018. Anticipating Cryptocurrency Prices Using Machine Learning. *Hindawi: Complexity,* Volume 2018.

5. Ambolis, D., 2023. *The Impact Of Ethereum On The Blockchain Industry: A Look At The Platform's History And Future.* [Online]
Available at: https://blockchainmagazine.net/the-impact-of-ethereum-on-the-blockchain-industry-a-look-at-the-platforms-history-and-future/

6. Benzekri, M. K. & Özütler, H. Ş., 2021. On the Predictability of Bitcoin Price Movements: A Short-term Price Prediction with ARIMA. *Journal of Economic Policy Researches,* 8(2), pp. 293-309.

7. Carmona, R., 2014. Time Series Models: AR, MA, ARMA, & ALL THAT. In: *Statistical Analysis of Financial Data in R.* New York, NY: Springer, pp. 345-421.

8. Chen, Z., Li, C. & Sun, W., 2020. Bitcoin price prediction using machine learning: An approach to sample dimension engineering. *Journal of Computational and Applied Mathematics,* Volume 365.

9. CoinMarketCap, 2023. *Bitcoin (BTC) Price, Charts, Market Cap, and other Metrics.* [Online]
Available at: https://coinmarketcap.com/currencies/bitcoin

10. Cortez, P., Rocha, M. & Neves, J., 2004. Evolving Time Series Forecasting ARMA Models. *Journal of Heuristics,* Volume 10, pp. 415-429.

11. Dama, F. & Sinoquet, C., 2021. Time Series Analysis and Modeling to Forecast: a Survey. *arXiv preprint arXiv:2104.00164.*

12. Darley, O. G., Yussuff, A. I. O. & Adenowo, A. A., 2021. Price Analysis and Forecasting for Bitcoin Using Auto Regressive Integrated Moving Average Model. *Annals of Science and Technology,* 6(2), pp. 47-56.

13. Demir, E., Gozgor, G., Lau, C. K. M. & Vigne, S. A., 2018. Does economic policy uncertainty predict the Bitcoin returns? An empirical investigation. *Finance Research Letters,* Volume 26, pp. 145-149.

14. Di Gangi, L., Lapucci, M., Schoen, F. & Sortino, A., 2022. Improved Maximum Likelihood Estimation of ARMA Models. *Lobachevskii Journal of Mathematics,* Volume 43, pp. 2433-2443.

15. Dickey, D. A. & Fuller, W. A., 1981. Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root. *Econometrica,* 49(4), pp. 1057-1072.

16. Federal Competitiveness and Statistics Centre (FCSC), 2021. *GDP - Production By Economic Sectors and Activities (Current Prices).* [Online]

Available at: https://uaestat.fcsc.gov.ae
[Accessed 2024].

17. Fernandes, M. et al., 2021. Bitcoin Price Prediction. *2021 7th IEEE International Conference on Advances in Computing, Communication and Control, ICAC3 2021.*

18. Fiaidhi, J., Sabah, A., Ansari, M. A. & Ayaz, Z., 2020. Bitcoin Price Prediction using ARIMA Model.

19. Flores, J. H. F., Engel, P. M. & Pinto, R. C., 2012. *Autocorrelation and partial autocorrelation functions to improve neural networks models on univariate time series forecasting.* Brisbane, IEEE.

20. Frankenfield, J., 2023. *Tether (USDT): Meaning and Uses for Tethering Crypto Explained.* [Online] Available at: https://www.investopedia.com/terms/t/tether-usdt.asp

21. Frankenfield, J., 2023. *What Was Mt. Gox? Definition, History, Collapse, and Future.* [Online] Available at: https://www.investopedia.com/terms/m/mt-gox.asp

22. Franz, J. & Spence, D., 2023. *Guide to recessions: 9 things you need to know.* [Online] Available at: https://www.capitalgroup.com/advisor/insights/articles/guide-to-recessions.html [Accessed 2024].

23. Giudici, G., Milne, A. & Vinogradov, D., 2020. Cryptocurrencies: market analysis and perspectives. *Journal of Industrial and Business Economics,* 47(1), pp. 1-18.

24. Hafid, A., Hafid, S. A. & Makrakis, D., 2022. Bitcoin Price Prediction using Machine Learning and Technical Indicators. *Preprints.org.*

25. Hodson, T. O., 2022. Root-mean-square error (RMSE) or mean absolute error (MAE): when to use them or not. *Geoscientific Model Development,* 15(14), p. 5481–5487.

26. Huang, C. & Petukhina, A., 2022. ARMA and ARIMA Modeling and Forecasting. In: *Applied Time Series Analysis and Forecasting with Python.* s.l.:Springer, Cham, pp. 107-142.

27. Hua, Y., 2020. *Bitcoin price prediction using ARIMA and LSTM.* s.l., EDP Sciences.

28. Investing.com, n.d. *Bitcoin Streaming Chart.* [Online] Available at: https://www.investing.com/crypto/bitcoin [Accessed 2024].

29. Jay, P. et al., 2020. Stochastic neural networks for cryptocurrency price prediction. *IEEE Access,* Volume 8, pp. 82804-82818.

30. Ji, S., Kim, J. & Im, H., 2019. A Comparative Study of Bitcoin Price Prediction Using Deep Learning. *Mathematics,* 7(10), p. 898.

31. Karakoyun, E. Ş. & Çıbıkdiken, A. O., 2018. *Comparison of ARIMA Time Series Model and LSTM Deep Learning Algorithm for Bitcoin Price Forecasting.* Prague, MAC Prague Consulting s.r.o.

32. Kayal, P. & Rohilla, P., 2021. Bitcoin in the economics and finance literature: a survey. *SN Business & Economics,* Volume 1, p. 88.

33. Lipovyanov, P., 2019. *Blockchain for Business 2019.* s.l.:Packt Publishing Limited.

34. Liu, M. et al., 2021. Forecasting the price of Bitcoin using deep learning. *Finance Research Letters,* Volume 40.

35. Mahmud, W., Fanani, A. Z., Santoso, H. A. & Budiman, F., 2023. *Review of Time Series Classification Techniques and Methods.* Semarang, IEEE.

36. McNally, S., Roche, J. & Caton, S., 2018. *Predicting the Price of Bitcoin Using Machine Learning.* Cambridge, UK, IEEE, pp. 339-343.

37. Meynkhard, A., 2019. Fair market value of bitcoin: Halving effect. *Investment Management and Financial Innovations,* 16(4), pp. 72-85.

38. Mittal, A., Dhiman, V., Singh, A. & Prakash, C., 2019. *Short-Term Bitcoin Price Fluctuation Prediction Using Social Media and Web Search Data.* Noida, India, IEEE, pp. 1-6.

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

40. Munim, Z. H., Shakil, M. H. & Alon, I., 2019. Next-Day Bitcoin Price Forecast. *Journal of Risk and Financial Management,* 12(2), p. 103.

41. Nagula, P. K. & Alexakis, C., 2022. A new hybrid machine learning model for predicting the bitcoin (BTC-USD) price. *Journal of Behavioral and Experimental Finance,* Volume 36.

42. Nakamoto, S., 2008. *Bitcoin: A Peer-to-Peer Electronic Cash System.* [Online]
Available at: www.bitcoin.org

43. NASA's Goddard Institute for Space Studies (GISS), 2022. *Vital Signs: Global Temperature.* [Online]
Available at: https://climate.nasa.gov/vital-signs/global-temperature/

44. Pappu, A. R., Kar, S. & Kadu, S., 2021. *ACF/PACF-Based Distance Measurement Techniques for Detection of Blockages in Impulse Lines of a Pressure Measurement Circuit for Nuclear Reactors.* s.l., Springer, Singapore.

45. PayPal, 2023. *Company Timeline.* [Online]
Available at: https://about.pypl.com/who-we-are/history-and-facts/default.aspx

46. Peng, Y. et al., 2018. The best of two worlds: Forecasting high frequency volatility for cryptocurrencies and traditional currencies with Support Vector Regression. *Elsevier,* Volume 97, pp. 177-192.

47. Petrică, A., Stancu, S. & Ghițulescu, V., 2017. Stationarity – The Central Concept in Time Series Analysis. *International Journal of Emerging Research in Management and Technology,* Volume 6, pp. 6-16.

48. Ponciano, J., 2021. *More Than $600 Million Stolen In Ethereum And Other Cryptocurrencies—Marking One Of Crypto's Biggest Hacks Ever.* [Online]
Available at: https://www.forbes.com/sites/jonathanponciano/2021/08/10/more-than-600-million-stolen-in-ethereum-and-other-cryptocurrencies-marking-one-of-cryptos-biggest-hacks-ever/?sh=d3410657f623

49. Poongodi, M., Vijayakumar, V. & Chilamkurti, N., 2020. Bitcoin price prediction using ARIMA model. *International Journal of Internet Technology and Secured Transactions,* 10(4).

50. Ranjan, S., Kayal, P. & Saraf, M., 2022. Bitcoin Price Prediction: A Machine Learning Sample Dimension Approach. *Computational Economics,* Volume 61, p. pages 1617–1636.

51. Reiff, N., 2023. *The Collapse of FTX: What Went Wrong with the Crypto Exchange?.* [Online] Available at: https://www.investopedia.com/what-went-wrong-with-ftx-6828447

52. Richter, C., Kraus, S. & Bouncken B., R., 2015. Virtual Currencies Like Bitcoin As A Paradigm Shift In The Field Of Transactions. *International Business & Economics Research Journal,* 14(4), pp. 575-586.

53. Roy, S., Nanjiba, S. & Chakrabarty, A., 2018. Bitcoin Price Forecasting Using Time Series Analysis. *2018 21st International Conference of Computer and Information Technology (ICCIT),* pp. 21-23.

54. Saad, M. & Mohaisen, A., 2018. *Towards characterizing blockchain-based cryptocurrencies for highly-accurate predictions.* Honolulu, HI, USA, IEEE, pp. 704-709.

55. Sin, E. & Wang, L., 2017. *Bitcoin price prediction using ensembles of neural networks.* Guilin, China, IEEE, pp. 666-671.

56. Singh, H. & Agarwal, 2018. *Empirical Analysis of Bitcoin Market Volatility Using Supervised Learning Approach.* Noida, India, IEEE, pp. 1-5.

57. Tredinnick, L., 2019. Cryptocurrencies and the blockchain. *Business Information Review,* 36(1), pp. 39-44.

58. U.S. Energy Information Administration, 2020. *U.S. natural gas consumption has both winter and summer peaks.* [Online]
Available at: https://www.eia.gov/todayinenergy/detail.php?id=42815
[Accessed 2024].

59. Urquhart, A. & Zhang, H., 2018. Is Bitcoin a Hedge or Safe-Haven for Currencies? An Intraday Analysis.

60. Verma, Y., 2022. *Quick way to find p, d and q values for ARIMA.* [Online]
Available at: https://analyticsindiamag.com/quick-way-to-find-p-d-and-q-values-for-arima/

61. Wirawan, I. M., Widiyaningtyas, T. & Hasan, M. M., 2019. *Short Term Prediction on Bitcoin Price Using ARIMA Method.* Semarang, Indonesia, IEEE, pp. 260-265.

62. World Bank – processed by Our World in Data, 2023. *17.13.1 - Annual GDP growth (%) - NY_GDP_MKTP_KD_ZG.* [Online]
Available at: https://ourworldindata.org
[Accessed 2024].

63. Yenidoğan, I. et al., 2018. *Bitcoin Forecasting Using ARIMA and PROPHET.* Sarajevo, Bosnia and Herzegovina, IEEE.