CHARLES UNIVERSITY
FACULTY OF SOCIAL SCIENCES
Institute of Economic Studies

The Profitability of Standard Trading Strategies in Cryptocurrency Markets

Bachelor thesis

Author: Miroslav Duda

Study program: Economic Theories

Supervisor: doc. PhDr. Ladislav Krištoufek Ph.D.

Year of defense: 2019

Declaration of Authorship

The author hereby declares that he compiled this thesis independently, using only the listed resources and literature, and the thesis has not been used to obtain any other academic title.

The author grants to Charles University permission to reproduce and to distribute copies of this thesis in whole or in part and agrees with the thesis being used for study and scientific purposes.

Prague, May 7, 2019

Miroslav Duda

Abstract

The thesis attempts to determine how strategies used for forecasting and trading on foreign exchange and stock markets perform when applied to cryptocurrency markets. The approaches explored are ARIMA, VAR, MA Crossover, and Granger Causality using gold prices and S&P 500. The currencies traded are Bitcoin, Ethereum, Binance Coin, and Basic Attention Token. The models are trained on logarithmically transformed and differenced time series composed of the currencies' daily and hourly closing prices. Applying these strategies mostly leads to ambiguous results, with MA Crossover generally performing better than VAR, which in turn performs better than ARIMA. However, every strategy was moderately successful for at least one of the currencies examined. Trading on the hourly dataset was negatively influenced by sudden price jumps. ARIMA and VAR perform better in the inter-bubble periods. No significant Granger causality was found.

Keywords	Cryptocurrency, Trading, Bitcoin, Ethereum, Binance Coin, Basic Attention Token, ARIMA, VAR, MA Crossover, Granger Causality

Title	The Profitability of Standard Trading Strategies in Cryptocurrency Markets

Author's e-mail	miroslav.duda11@gmail.com
Supervisor's e-mail	ladislav.kristoufek@fsv.cuni.cz

Abstrakt

Práce se pokouší určit, jakou úspěšnost mají strategie používané na forexových a akciových trzích při aplikaci na kryptoměnové trhy. Zkoumanými přístupy jsou ARIMA, VAR, MA Crossover (klouzavý průměr) a Grangerova kauzalita s využitím cen zlata a S&P 500. Obchodovanými kryptoměnami jsou Bitcoin, Ethereum, Binance Coin a Basic Attention Token. Modely jsou trénovány na logaritmicky transformovaných a diferencovaných časových řadách složených z konečných denních a hodinových cen jednotlivých měn. Aplikace strategií vede k nejednoznačným výsledkům. MA Crossover dosahuje obecně lepších výsledků než VAR, ARIMA pak vede k nejhorším. Přesto každá strategie funguje přijatelně pro alespoň jednu z měn. Obchodování na hodinových řadách bylo negativně ovlivněno náhlými cenovými skoky. ARIMA a VAR dosahují lepších výsledků v obdobích mezi cenovými bublinami. Signifikantní Grangerova kauzalita nebyla nalezena.

Klíčová slova	Kryptoměny, Obchodování, Bitcoin, Ethereum, Binance Coin, Basic Attention Token, ARIMA, VAR, Klouzavý průměr, Grangerova kauzalita
Název práce	Ziskovost standardních obchodních strategií na kryptoměnových trzích

E-mail autora	miroslav.duda11@gmail.com
E-mail vedoucího práce	ladislav.kristoufek@fsv.cuni.cz

Acknowledgments

The author would like to thank the supervisor, doc. PhDr. Ladislav Krištoufek Ph.D., for thesis topic suggestion, consultations of the thesis, and recommendations for its improvement, and RNDr. Michal Červinka Ph.D. for feedback on the thesis draft.

Typeset in FSV LaTeX template with great thanks to prof. Zuzana Havrankova and prof. Tomas Havranek of Institute of Economic Studies, Faculty of Social Sciences, Charles University.

Bibliographic Record

Duda, Miroslav: *The Profitability of Standard Trading Strategies in Cryptocurrency Markets*. Bachelor thesis. Charles University, Faculty of Social Sciences, Institute of Economic Studies, Prague. 2019, pages 54. Advisor: doc. PhDr. Ladislav Krištoufek Ph.D.

Contents

List of Tables

List of Figures

Acronyms

BTC Bitcoin

ETH Ethereum

BNB Binance Coin

BAT Basic Attention Token

USD United States Dollar

ARIMA Autoregressive Integrated Moving Average

VAR Vector Autoregression

MA Moving Average

Bachelor Thesis Proposal

Author	Miroslav Duda
Supervisor	doc. PhDr. Ladislav Krištoufek Ph.D.
Proposed topic	The Profitability of Standard Trading Strategies in Cryptocurrency Markets

Research question and motivation The main question is whether the trading strategies used in stock markets and forex markets can be applied to cryptocurrency trading with consistent success, and whether this approach leads to better results than simply holding a specific currency without trading. With cryptocurrencies being highly influenced by psychological factors, the reliability of strategies and models which work in other markets is questionable at best. The analysis may thus also help answer the question whether cryptocurrency trading as a whole can be profitable, or if it is only another layer of gambling in an already highly volatile market. Provided cryptocurrencies continue gaining in importance, answering this question may also provide insight into the workings of emerging markets, and may be used for comparison to these markets later when and if they stabilize, possibly allowing applications in volatile and "irrational" markets in general.

Contribution As cryptocurrency research is still relatively scarce with trading research mostly limited to trading based on public interest (ex. (3), (4), (6) in the academic literature section) confirming that sentiment-based trading can be effective, and explaining the currenciesâ€ volatility (ex. (5), (7)), both the empirical analysis of price data and the results may provide a foundation for further research, and offer insight into the relation between conventional markets and cryptocurrency markets. It can also be used to gain a rough estimate of the usability of some price and volume based strategies in cryptocurrency markets, which both researchers and traders may find helpful.

Methodology Data on cryptocurrency prices and volume will be gathered from CoinMarketCap.com or comparable sites. Trading strategies will then be simulated on a chosen period or periods of time. The end result of applying these strategies will be compared to the profitability of simply holding the chosen currencies in the same period. Results will be evaluated with an ultimate goal of discerning which strategies and models can be successfully used.

Outline After an introduction, the dataset will be presented. The dataset will then be analyzed. The thesis will conclude with comments on the analysis results.

- Introduction, Expected results
- Dataset, Analysis
- Comments, Conclusion

List of academic literature

1. Pärlstrand, E., & Rydén, O. (2015). Explaining the market price of Bitcoin and other Cryptocurrencies with Statistical Analysis
2. Cocco, L., Concas, G. & Marchesi, M. (2017). Using an artificial financial market for studying a cryptocurrency market
3. Georgoula, Ifigeneia and Pournarakis, Demitrios and Bilanakos, Christos and Sotiropoulos, Dionisios and Giaglis, George M., Using Time-Series and Sentiment Analysis to Detect the Determinants of Bitcoin Prices (May 17, 2015)
4. Garcia, David & Schweitzer, Frank. (2015). Social signals and algorithmic trading of Bitcoin. Royal Society Open Science.
5. Donier J, Bouchaud J-P (2015) Why Do Markets Crash? Bitcoin Data Offers Unprecedented Insights.
6. Kristoufek L (2015) What Are the Main Drivers of the Bitcoin Price? Evidence from Wavelet Coherence Analysis.
7. Cermak, Vavrinec, Can Bitcoin Become a Viable Alternative to Fiat Currencies? An Empirical Analysis of Bitcoin's Volatility Based on a GARCH Model (May 2, 2017).

Author Supervisor

Chapter 1

Introduction

Since the inconspicuous and shadow economy focused introduction of cryptocurrencies to the world through Bitcoin by (Nakamoto 2008), these decentralized alternatives to fiat currencies have gained mass public awareness, and grown into a relatively small but significant phenomenon distinguished especially by their enormous volatility with both unprecedented growth and devastating crashes, attracting many retail investors and technology enthusiasts, while discouraging most institutional investors as a result of cryptocurrencies' high degree of uncertainty. The main selling point of cryptocurrencies is their anonymity and decentralized nature, facilitated by the blockchain, a distributed ledger containing and permanently recording every transaction for which a specific cryptocurrency was used. Before the invention of the blockchain, without a trusted third party, digital currencies presented risks to their users, as it was impossible to ensure that the currency was only sent to one recipient at a time, ie. that the currency was only "spent" once. The structurally impartial and immutable transaction record solved this double spending problem, eventually leading to the emergence of a new market. While the cryptocurrency market is superficially similar to traditional stock and especially foreign exchange markets, an important difference exists in cryptocurrencies' large degree of uncertainty regarding their true value in comparison to a foreign exchange investor's value estimates of a currency based on their respective countries' economic performances and monetary policies, or a stock market investor's analysis of the underlying companies represented by the stocks.

However, a possibility still exists that the more technical approaches to valuation and trading used in the aforementioned traditional markets can be successfully applied to cryptocurrencies in spite of their idiosyncratic nature. Thus, the main purpose of the research is to investigate the profitability of these approaches in the cryptocurrency markets, compare them to long term "buy-and-hold" investment strategies without trading where cryptocurrency holdings are simply fixed to the initial investment level, and briefly consider potential influences of other price drivers, such as the prices of gold and stock market performance. In addition to the apparent benefit of deeper insight into which models and approaches can or cannot work for cryptocurrencies specifically, the research also aims to deepen comprehension of emerging and developing markets, as the conclusions viewed in relation to the results of these strategies in the traditional markets can be generalised, and used as a basis for understanding their interrelations. As the models are used for trading several different cryptocurrencies, analysis of the results also aims to discern the degree of intramarket cohesion, trying to answer whether the same trading strategies might be successfully applied to any cryptocurrency. The models and approaches used for forecasting and trading cryptocurrencies are ARIMA, VAR, MA Crossover, and Granger causality. In order to facilitate comprehension from an investor's point of view, the results are commented on both in terms of changes in cryptocurrency holdings, and in terms of changes in dollar value, as the extreme shifts in cryptocurrency prices can easily make trading without the USD feedback highly misleading.

The thesis has the following structure. The literature review, which immediately follows the introduction, summarizes the existing literature on the topic of cryptocurrency trading, and compares approaches directly using price and volume as opposed to external factors. In the dataset chapter, the price and volume time series for each cryptocurrency, gold, and S&P500 are presented along with the reasoning for the inclusion of the currencies, and the dataset treatment is explained. In the methodology section, the methods are presented along with expectations regarding their performance, and potential issues. Also included are the trading algorithms. The results and discussion section contains the descriptions of the outcomes of applying the algorithms to the datasets, and a summary debating why each approach succeeded or failed, with some comments on other interesting discoveries. The concluding section contains a concentrated results description and the main findings. Appendices contain an example of the algorithms, and tables showing detailed trading results.

Chapter 2

Literature Review

In existing research, trading strategies based on external factors, such as Wikipedia, Google or Twitter interest in cryptocurrency, measured by frequency of incidence of cryptocurrency-related terms on the aforementioned sites, appear to be slightly favoured over algorithmic trading derived from price action and trade volume. This is not surprising considering cryptocurrency's sensitivity to psychological influence.

Twitter sentiment analysis was examined by (Colianni *et al.* 2015), using Naive Bayes, logistic regression, and supporting vector machines, predicting the sign of Bitcoin price change. The Bernoulli Naive Bayes algorithm reached a day to day accuracy of 95%, and hour to hour accuracy of 76.23%, both of which are quite precise. However, Bitcoin volatility has risen since the publication date, possibly invalidating these results. Tweets were also used by (Kaminski & Gloor 2014), focusing on the context of the tweets, dividing them into groups based on positive or negative emotional signals as indicated by the word "Bitcoin" in conjunction with positively or negatively charged words. The Granger causality results were inconclusive, with the price and Twitter sentiment apparently moving jointly. It is of note that large trading volumes seemed to be correlated with the number of tweets expressing uncertainty rather than positivity or negativity.

Google Trends, Wikipedia and Bitcoin interactions were explored by (Krištoufek 2013). Vector autoregression (VAR(1)) was used for Google Trends, and vector error-correction model (VECM(7)) for Wikipedia. Bitcoin searches on Google and Wikipedia were found to positively influence Bitcoin price and vice versa, with the reasoning that the valuation of an asset with highly uncertain fundamentals will be driven mostly by speculation, but notifying of the difficulties with differentiating between interest generating price action, and price action causing a surge in interest. (Pärlstrand & Rydén 2015) have used log-log OLS to estimate the effects of prices of oil, exchange rates between the main fiat currencies, or the number of transactions on the blockchain on the prices of Bitcoin, Ripple and Litecoin. Similarly to (Krištoufek 2013), their main finding was the strong influence of Google searches on Bitcoin price, and a difference between these effects on the prices of Bitcoin and Litecoin, and on the prices of Ripple. Twitter sentiment analysis was also used by (Georgoula *et al.* 2015) through support vector machines analysis of Twitter posts, finding that Bitcoin price is positively related to sentiment ratio of Twitter users, and a negative relationship between S&P 500 and Bitcoin price in the long run. (Krištoufek 2015) utilised wavelet coherence to analyze the influence of economic drivers (demand for currency), transaction drivers (ex. traded volume), or the role of the Chinese market. The main conclusions were the presence of the effects of "traditional" market mechanisms, which do appear to have an influence on Bitcoin in the long run, and Bitcoin's cyclical boom-bust behaviour. Wavelets were employed also by (Phillips & Gorse 2018), to discover the effects of Wikipedia and Google search volume and uniquely also Reddit posts through activity on the subreddits (subforums) of several cryptocurrencies. Here, the relationship between interest and growth was again found to be significant, specifically a positive effect in the long term, and general strenghtening during cryptocurrencies' bubble periods.

Apart from the wavelet research, (Phillips & Gorse 2017) have also attempted a unique method. In addition to a fairly standard concept of social media activity usage, a model typically used by epidemiologists to search for possible outbreaks was used to determine the presence of a bubble, forecast its movements, and use this information to enter and exit the market accordingly. A strategy based on an unanimous agreement on the existence of an "outbreak" (bubble) by each of the five hidden Markov models (one for each input) led to 1380% returns over the examined period, outperforming a buy and hold strategy almost twice, while an alternative based on averaging probabilities performed worse at 775% returns. (Caporale & Plastun 2018) also tried an approach usually better suited to forex and stock markets, trying to discern the effects of the day of the week on cryptocurrency prices using t-tests on the returns for each day to test the hypothesis that the returns distributions are identical throughout the week. Of the four cryptocurrencies tested, only Bitcoin showed an abnormal distribution on Mondays. A strategy in which the trader opens a long position on Mondays, and closes it before the day ends, led to a 60% proportion of the trades being profitable.

Moving on to price and volume based trading, (Karakoyun & Cibikdiken 2018) compared the usefulness of an ARIMA model, classification algorithms, and an empirical conditional distribution model, concluding that the classification algorithm led to better returns than the empirical conditional distribution model, which in turn outperformed the ARIMA model. (McNally *et al.* 2018) explored machine learning methods, with ARIMA providing results much weaker than the recurrent neural network (RNN) or long short term memory (LSTM), although it is worth noting the ARIMA predictions were not made one step ahead. The RNN and LSTM models performed comparably well. However, the confusion matrix for both reveals fairly unsatisfactory results, with accuracy crossing the 50% level only by a narrow margin. Finally, (Makarov & Schoar 2018) researched inter-exchange arbitrage. Quite interestingly, exchanges operating in countries with strict capital control show greater arbitrage spreads, possibly as a result of the citizens of these countries valuing cryptocurrencies higher. While transaction costs, incurred both by the exchanges in the form of withdrawal and deposit fees, and by the currencies themselves, might initially seem as the main barrier to arbitrage, the paper implies these might in fact be overshadowed by the risk of entrusting one's cryptocurrencies to a potentially fraudulent exchange.

In conclusion, both direct and indirect methods do seem to have varying degrees of success. The prediction attempts based on psychological factors (Twitter, Wikipedia, Google) seem to provide slightly more impressive results. In spite of this, the price-based approaches still lead to reasonably relevant outcomes, justifying further exploration in this area.

Chapter 3

Data

The dataset for daily cryptocurrency price action and volume was sourced from CoinMarketCap[1], a site which aggregates price and volume data from numerous exchanges into a single volume-weighted average number, through package "rvest" by (Wickham 2016). The hourly dataset was sourced from CryptoDataDownload[2], specifically using data from the exchange Binance. Closing prices for each period ("tick") were chosen as representative. While unlike the stock and foreign exchange markets, cryptocurrency markets do not technically "close", the closing price is still registered at the end of each tick. Gold Prices were drawn from Quandl[3] using the "Quandl" package by (Raymond McTaggart *et al.* 2018), and the S&P 500 index from Yahoo Finance[4] using package "quantmod" by (Ryan & Ulrich 2018). Both gold and stock markets close during the weekends, which presented a problem for Granger causality testing due to missing values. For this reason, the price entries for each Friday were simply repeated for both days of the weekend.

The following cryptocurrencies were chosen for further examination.

Bitcoin was picked as the first currency. This choice was motivated by the decidedly dominant position of Bitcoin in the cryptocurrency space, representing roughly 50% of total market capitalization of the entire cryptocurrency market. It is relatively stable (compared to other cryptocurrencies), and correlates strongly with the smaller currencies.

---

[1] `https://coinmarketcap.com/`

[2] `https://www.cryptodatadownload.com/`

[3] `https://www.quandl.com/`

[4] `https://finance.yahoo.com/`

CoinMarketCap only offers price data for Bitcoin starting from April 28, 2013, meaning that the data for approximately two years are absent at the beginning. However, considering the extremely low liquidity and absolute size of the market in its initial period, these two years can be omitted without major detriment. The main points of interest on the Bitcoin price curve in chrono-

Figure 3.1: BTC closing prices

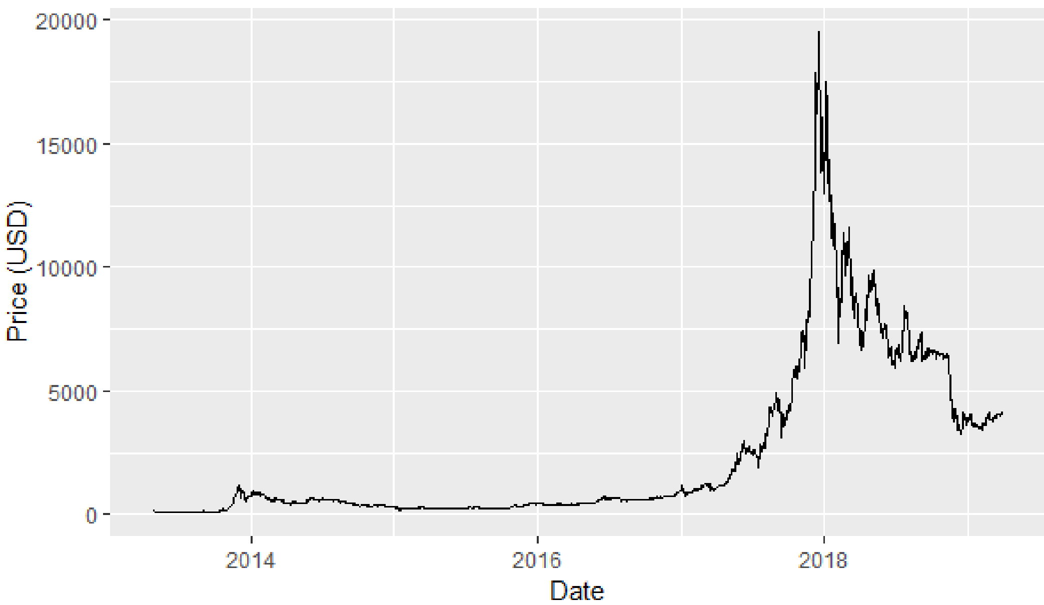

logical order are the first significant peak in late 2013, and, of course, the 2017 bubble and the subsequent crash. In December 2013, Bitcoin reached $1151 for 1 BTC, a notable success considering that the price had been at only $200 a month earlier. This first "bubble" burst once Mt. Gox, the exchange managing some 70% of all cryptocurrency trade volume at that point, announced that a significant amount of the BTC held in the exchange's wallets had been stolen. The second 2017 bubble's burst, where Bitcoin almost reached $20,000 likely cannot be attributed to a single factor. Another point of interest is the large price drop in late 2018.

The next cryptocurrency examined is Ethereum, the second largest currency in terms of market capitalization, occasionally losing this position to Ripple (XRP) and, unlike the purely currency oriented Bitcoin, a smart contract platform, providing its users with applications beyond the three main functions of money. Apart from simple contracts in which the Ethereum network serves as an independent arbiter, potential also lies in decentralized applications ("DApps"). Ethereum behaves similarly to Bitcoin, while at the same time showing sufficient divergence, making it valuable as a secondary indicator of the cryptocurrency market as a whole. While most of the main characteristics

Figure 3.2: ETH closing prices

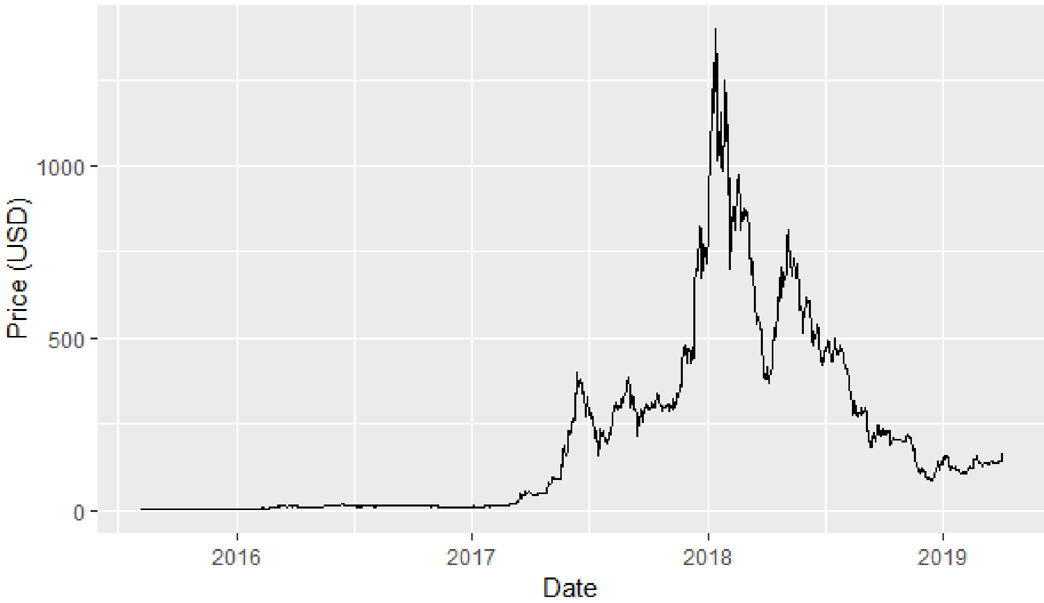

of the price development closely follow those of Bitcoin, it is worth noting that Ethereum reached its second peak in early 2018, slightly later than Bitcoin, and that within 2018 Ethereum price dropped below $100 per ETH, an over 90% decrease compared to the peak.

The third cryptocurrency chosen is Binance Coin. A currency intrinsically tied to its parent organization, the Chinese exchange Binance[5], its main function lies in lowering trading fees on this exchange, as long as BNB is used to cover them, with periodic reductions in supply, as Binance destroys ("burns") a portion of the currency received through the fees. While a currency as specialized as BNB could be expected to fill a highly specific niche at best, it instead rapidly rose to seventh position in market capitalization, mostly owing to Binance's impressive results. Binance Coin was included for its atypical behaviour, showing impressive growth in USD value even in the post-crash period, and noticeably differing from Bitcoin's price action. The main point of interest in the price development of BNB is the 2018 post-crash development.

[5] `https://www.binance.com/en`

Figure 3.3: BNB closing prices

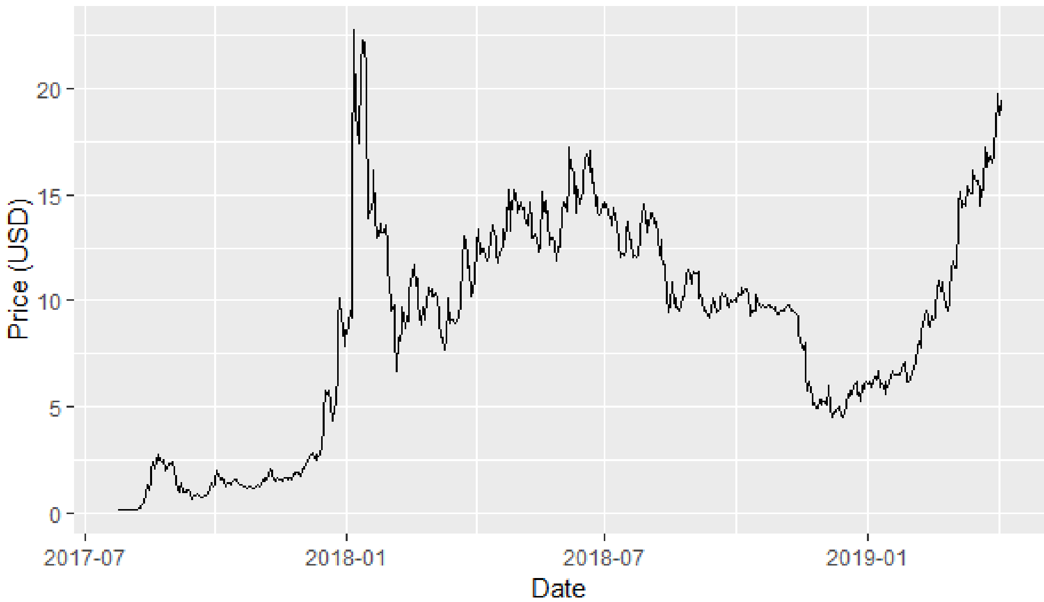

Where most cryptocurrencies, especially the smaller ones, lost over two thirds of their peak value, BNB swiftly recuperated, and even saw growth in the first half of 2018, which was, however, lost later in the same year. In early 2019, the currency's price surged, eventually crossing even the 2018 peak.

The last cryptocurrency included is Basic Attention Token as a representative of the smaller currencies. The token aims to serve as a unit of account on the BAT platform, a project attempting to facilitate interactions between advertisers, publishers, and users. Compared to other similar projects, it has a respectable track record, relatively active developers, and a clearly defined use. It is also one of the more predictable currencies, lacking the extremely persistent spikes and crashes typical of other smaller currencies, although its volatility is fairly high. The visible periodicity is BAT's most distinctive characteristic. Effects of the most critical events in the cryptocurrency market, such as the 2017/2018 peak and crash, are still apparent, but even in the post crash period, where many small currencies lost over 90% of their peak value without any later uptrends, the price manages to reach decent highs.

As the recent absolute size values of cryptocurrency price are orders of magnitude larger than the prices in early periods, a proper treatment can make the ensuing results more relevant. This can be achieved with a logarithmic

Figure 3.4: BAT closing prices

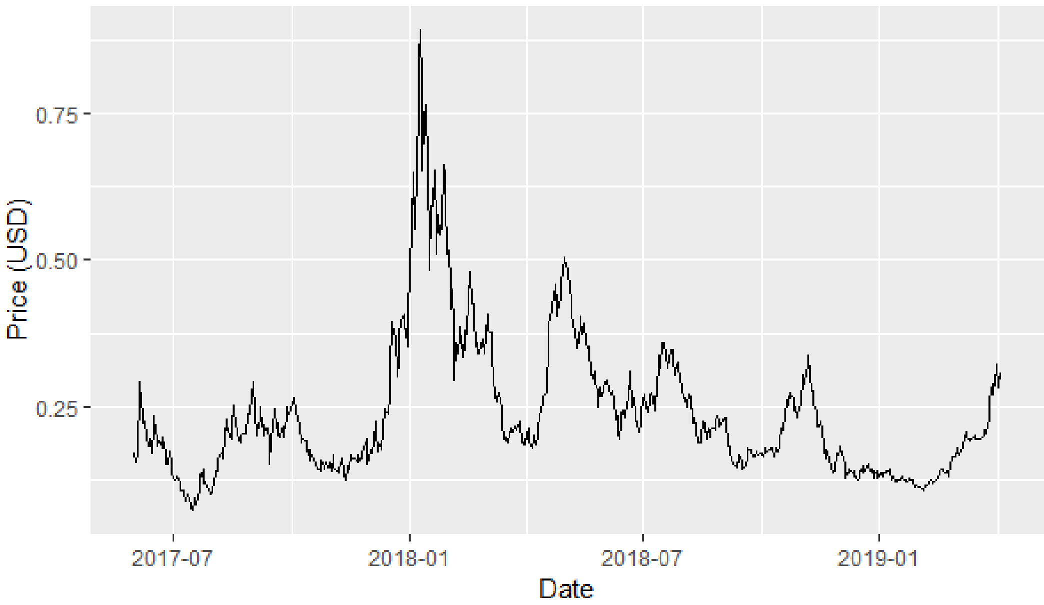

transformation of the time series, so that we obtain a series focused more on the relative rather than absolute changes in price, giving more weight to the earlier observations. It also leads to the series becoming closer to stationarity, both before and after differencing, working in favour of our assumptions.

The Box-Cox transformation, originally proposed by (Box & Cox 1964), from the package "forecast" by (Hyndman & Khandakar 2008) was picked as suitable. It applies the following function to the time series:

$$y(\lambda) = \begin{cases} \frac{y^{\lambda}-1}{\lambda} & \lambda \neq 0 \\ \log y & \lambda = 0 \end{cases}$$

with $\lambda \in [-5, 5]$. The optimal value of λ is determined through maximization of a log-likelihood function.

For the purposes of methods used, stationarity was a necessary requirement. For price action of most cryptocurrencies, it is obviously violated, as further supported by Augmented Dickey-Fuller (ADF) tests. After first differencing, the ADF tests report p-values below 0.01, allowing us to reject the null hypothesis of non-stationarity, and enabling the use of further analysis.

Chapter 4

Methodology

A 50:50 train-test split was used for all ARIMA and VAR models. MA Cross was applied to the entire dataset where possible, ie. after 10 periods when both moving averages could be defined.

4.1 Autoregressive Integrated Moving Average

One of the chosen tools for predicting cryptocurrency price action is the Autoregressive integrated moving average model, or the $ARIMA(p, d, q)$ model, the original fitting procedure for which was proposed by (Box & Jenkins 1970). It is composed of three main components:

1. The autoregressive part denoted by AR.

   The autoregressive model aims to quantify a linear dependence of the output values on previous lagged values, and attempts to make predictions based on this relationship.

2. The moving average model represented by MA

   The moving average model predicts the output values using the error terms in previous periods.

3. The degree of integration represented by I.

   The I informs that the model is integrated of order d, ie. the time series input is differenced d times to enable predictions which would otherwise be impossible due to absence of stationarity.

These three separate models are then generalized into the ARIMA model, which combines the functionalities of each of its constituent parts. The model is then defined as

$$y_t = c + \sum_{i=1}^{k} a_i y_{t-i} + \sum_{i=1}^{l} b_i \epsilon_{t-i}$$

where the c represents the optional intercept, a_i and b_i are the model's coefficients, y is the modelled time series, and ϵ is the error term.

After ensuring the modelled time series is stationary, the first challenge of ARIMA is correctly choosing the lag order for the autoregressive and moving average components, the p and q parameters. This can be done either using the autocorrelation and partial autocorrelation functions, which allow manually choosing the lag order beyond which the autocorrelations stop being significant, or using information criteria (Akaike information criterion as defined by (Akaike 1974) and the closely related Bayesian information criterion in this case) to determine the optimal level of lag to include. Tentative checking of the usability of ARIMA was done through the ACF and PACF. While no immediately obvious level beyond which further lags do not influence future values is present, and the (partial) autocorrelations may possibly be random, some levels of lag do have significant autocorrelation coefficients.

The profitability of using ARIMA predictions was tested by the following algorithm (see Appendix A for the ETH version):
An initial endowment of 1 BTC, 100 ETH, 1000 BNB, and 10,000 BAT was chosen arbitrarily. The algorithm then predicts the closing price for the following day. If this predicted value is lower than the most recent closing price, and the currency is owned, it is sold and converted into USD. If the predicted value is greater than the most recent closing price, and USD is owned, the currency is bought for owned USD. If neither of these conditions is satisfied, no action is taken. This process is then repeated for the following period, incorporating the closing price from the newly added day.

This strategy could realistically be hampered by the low depth of order books in cryptocurrency markets, causing the prices to move further before the desired quantity of a currency is bought or sold, and the exchange fees penalizing the trader for each trade. The liquidity problem was omitted due to the difficulty of its consideration in the algorithm, while the exchange fees could be included fairly easily, but were also omitted, as the fees vary greatly, ranging from zero percent of trade volume for offer makers on certain exchanges to up to tens of percent on others.

For the most significant currencies, two timescales are explored. One large, where the model is trained and used on daily price data, often encompassing information from the entirety of the currency's existence, and a smaller scale, where hourly data is used. The advantage of the larger scale is the smoother price curve, as individual investors cannot significantly alter the price in a short timespan, while the smaller timescale may show better results due to its lower usual volatility, with the added risk of unpredicted instant jumps. Due to the inferior availability of hourly and shorter price data for the smaller currencies, only Bitcoin and Ethereum are traded on the shorter timescale. Both Bitcoin and Ethereum were trained on the first three quarters of March, 2019, and tested on the last one, ie. the test set approximately corresponds to a week.

The main issue in using ARIMA for cryptocurrency analysis lies in its relative simplicity. It can be expected to struggle in an environment which is influenced by numerous external factors, as opposed to future values being determined solely by previous ones. The package "forecast" by (Hyndman & Khandakar 2008) was used for fitting the ARIMA models and creating forecasts.

4.2 Vector Autoregression

The second tool used for forecasting the price time series is vector autoregression, or VAR. A method similar to ARIMA, it incorporates the effect of lags of time series different than the one being forecast in addition to the autoregressive component described in the ARIMA section, thus making predictions based on a vector of time series. While any number of time series may be included, for the purposes of predicting cryptocurrency price action only a trading volume time series is appended.

For two series with one lag included, the model is defined as follows:

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix} = \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \end{pmatrix} + \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{pmatrix}$$

where y_1 and y_2 are the two time series, $a_{1,1}$ through $a_{2,2}$ are the model coefficients, and $\epsilon_{1,t}$ and $\epsilon_{2,t}$ are the error terms, ie. each time series has coefficients for the included lags of every other time series and itself.

Analogously to ARIMA, stationarity is again a necessary condition. For cryptocurrency data the log transformed and differenced series are used. The volume data behave very similarly to price data, with the obvious exception of volume remaining high even after price starts dropping. ADF tests confirm nonstationarity, so the volume series are transformed through the same process as the price series in ARIMA analysis. The second condition unique to VAR is that the order of integration has to be identical for all the time series used as inputs. For both prices and volumes, first differencing combined with the log transformation is sufficient for stationarity according to ADF tests. The order of integration is thus $d = 1$ in all cases.

The algorithm used for trading the currencies is analogous to the one used for ARIMA. The model is used to forecast the price for the next period one step ahead. If price is predicted to drop and the currency is owned, it is sold. In the inverse situation, it is bought. The initial endowments for each currency are identical to those for ARIMA.

The motivation for including VAR models in the analysis is the role of volume in determining future price development. While no patterns are absolutely definite, as a trend reaches its top or bottom, volume often decreases, signalling that no more buyers or sellers are willing to push the price further. This in turn leads to price stagnation followed by a correction in the direction opposite to the previous trend and, again, an increase in volume. However, in cryptocurrency markets this pattern is violated quite often, unfortunately also in situations critical for trading, such as the January 2018 peak and crash, or the first post-bubble drop below $7,000, where the trends reversed without the volume reduction. The main problem of using VAR on volume data is that volume reaches high levels during periods of both price increases and decreases, and can thus be interpreted as a measure of volatility rather than a price predictor. Nevertheless, it is expected to be a significant variable for price forecasts, and may improve on the univariate ARIMA modelling.

The package "tsDyn" by (Stigler 2010) was used for fitting the VAR models and using them for creating forecasts.

4.3 Moving Average Crossover

Trading strategies based solely on moving averages are closely related, but in principle different, to the autoregressive methods described above. These strategies are among the simpler approaches used both in forex and securities markets, using the average values of n last observations. The main trading signal is the "crossover", a situation in which the value of one moving average intersects another, or the price itself.

The value of the moving average at any point is defined as

$$MA_k = \frac{P_{n-k+1} + P_{n-k+2} + ... + P_{n-1} + P_n}{k}$$

where k is the desired number of periods included in the computation of the average, P are the price datapoints, and n is the total number of periods.

The main advantage of this approach is its ease of use. It can be applied to any price dataset at any timescale, as long as at least k previous periods are known. It is also readily available even to non-professional traders, with many exchanges offering the MA values as part of their trading interface. In cryptocurrency markets, the main motivation behind its usage is the strong persistence of trends. Unlike the more complex tools used for time series forecasting, the moving average strategy ensures that as long as a currency is increasing or decreasing in price, it will not be bought or sold too early, regardless of the volatility.

The main disadvantage is the moving average's lagged nature, As the MA is composed purely of past values. Furthermore, the MA does not predict future values of the time series, only gives a questionable signal about its future direction. It is also expected to struggle during low volatility periods, in which crossovers happen frequently without a strong trend following them, leading to purchases (sales) of currencies which subsequently drop (rise). The decision between different lag levels and crossovers is mostly arbitrary. Traditionally, MA_5, MA_{10}, and MA_{15} are used for short term trading, while the longer term moving averages are used to predict the general direction of the market as a whole.

As an alternative to the simple moving averages (SMA or only MA) described above, exponential moving averages (EMA) can be used if a more dynamic indicator is needed. The EMA assigns higher weight to more recent observations, improving its reactivity to sudden changes. This in turn leads to potentially more frequent false trade signals, making the decision between SMA and EMA context sensitive.

EMA is defined as

$$EMA_{k,t} = P_t \cdot \frac{l}{1+k} + EMA_{k,t-1} \cdot \left(1 - \frac{l}{1+k}\right)$$

where l stands for the smoothing parameter, which is set to 2 in the following algorithm.

The trading algorithm for moving averages works as follows:
The initial endowments for each currency are, again, identical to those in the ARIMA section. Since no train set is necessary for MA, data from the entire existence of the currency are used for trading. For the first approach, MA_5 and MA_{10} are chosen as important. If MA_5 falls below the MA_{10} curve, the cryptocurrency held is sold for USD. In the reverse situation, the cryptocurrency is bought for USD held. The second approach uses the more reactive EMA_5 in place of MA_5, while MA_{10} remains in its simple form.

4.4 Granger Causality

A method distinct from the models above, a Granger causality test was used to search for the effect of lagged S&P 500, representing the stock market, and gold prices, on the price of Bitcoin. The Granger test, searching for a causal relationship as first proposed by (Granger 1969), is fundamentally a Wald test, which compares two models, one of which includes a lag of the Granger-causal time series with a predetermined lag, and the other without the Granger-causal series, using a lag of the "caused" time series instead. The test then attempts to determine whether the model with the Granger-causal series is sufficiently better. If the test finds Granger-causality between gold or the S&P 500 and cryptocurrencies, this could be used for estimating future price movements.

For Granger testing, stationarity is, again, a necessary condition. To this end, gold and S&P 500 price data were differenced once, ensuring stationarity according to an ADF test. For Bitcoin, the Box-Cox transformed and differenced dataset was utilized. Granger causality was only tested between Bitcoin and gold, and Bitcoin and S&P 500. The underlying reasoning is that the cryptocurrency market's strong interdependence allows Bitcoin to serve as a proxy due to its relative importance. Alternatively, weighted price data for the entire market could be used based on market capitalization, with the downside of a much more complex dataset. Lags of up to seventh order were used, representing a week.

For Granger testing, the main issue lies in the fact that the causality test does not quantify the relationship between the series modelled, only showing a specific type of causality. This can still be used to forecast whether the price will rise or fall to some degree, although more precise estimates are unlikely to be produced. If the movements of Bitcoin are found to be explainable by Granger causality, with either gold or S&P 500 causing Bitcoin, a trading strategy can be created where Bitcoin is bought or sold based on the movements of the causing series with an appropriate lag.

Chapter 5

Results and Discussion

5.1 Autoregressive Integrated Moving Average

The algorithm described in the Methodology section led to the following results: In terms of Bitcoin, value held decreased considerably throughout the trading period, reaching a minimum at almost 85% BTC lost. During the first half of the trading period, value held shows a very strong downwards trend, then finds a bottom, and starts increasing in the last quarter, albeit at a slow pace. At the end of the period, value held in BTC is close to a third of one Bitcoin, a significant loss compared to the initial endowment of 1 BTC. Value held in terms of USD shows gradual growth throughout the entire period, with a significant spike corresponding to Bitcoin's twenty thousand dollar peak. However, even at this peak, only the equivalent of around four thousand dollars is held, as at this point BTC held is close to its minimal value. After the crash, the mild upwards trend resumes, and value held in USD remains in the 1500-2000 area. The results obtained are obviously unsatisfactory. The algorithm leads to significant loss in Bitcoin held, and while the value in dollars shows stable growth, this can mostly be ascribed to the growth of the value of cryptocurrencies in general. One possible advantage is the low volatility of USD value, but this comes at a massive cost in the form of approximately 80% of peak value lost. The algorithm only performs well in 2018 after the crash, where Bitcoin price oscillated between six and ten thousand USD, which apparently provided a suitable environment for ARIMA. The main obstacle is quite decidedly the cryptocurrency boom in the latter half of 2017. The predictions during this period commonly underestimate the speed of growth, leading to the Bitcoin losses mentioned above.

Figure 5.1: ARIMA trading BTC

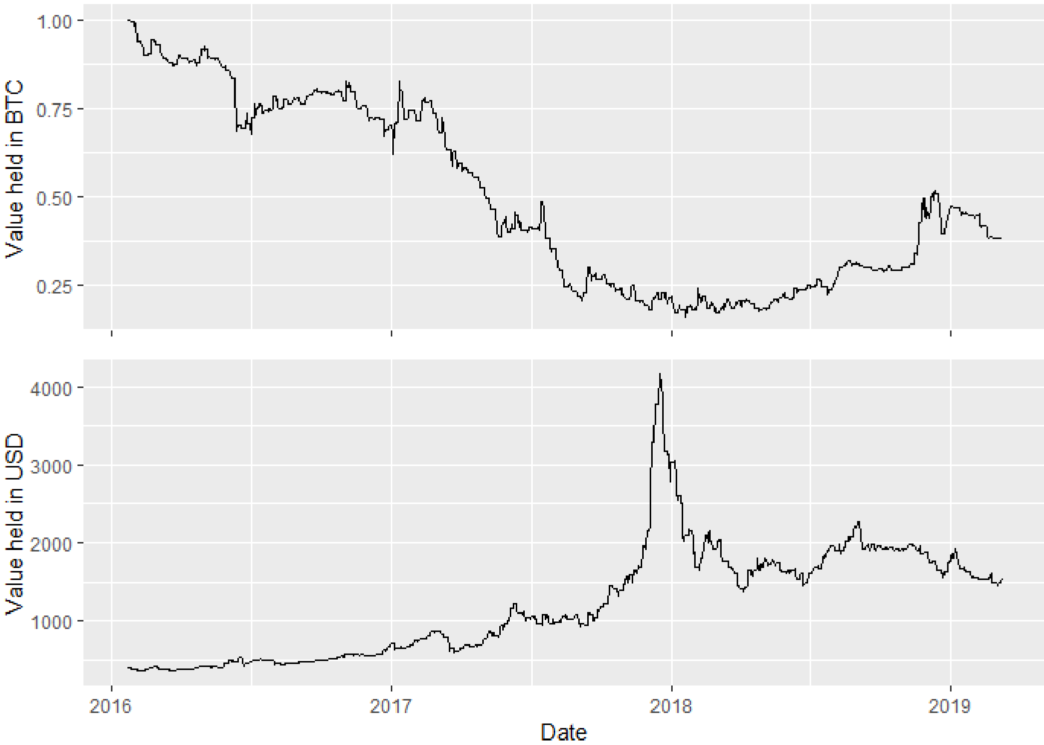

For Ethereum, the results are fairly similar. Large fraction of Ether held is lost during the 2017 bubble period (with a slight lag, as Ethereum peaked a few weeks after Bitcoin in January 2018), and is then slowly regained in the remainder of 2018. In contrast to ARIMA trading Bitcoin, Ether held does outperform a buy and hold strategy towards the end of the examined period, reaching almost 120 ETH, corresponding to an almost 20% increase over the initial endowment. Nevertheless, Ether held does drop back to the initial endowment level at the very end. Dollars held show a development much more worrying than that of Bitcoin. In spite of reasonable returns during the bubble, at approximately three times the initial investment, value in USD drops well below it in the 2018 crash. This can be attributed to Ethereum's massive drop in value. As Ether held stays roughly the same during this time, no buffer for the price crash is created.

Unlike the two large currencies above, the results for Binance Coin are quite positive. At one point, the algorithm reaches twice the initial investment, and while the BNB held does drop afterwards, the final returns still amount to

Figure 5.2: ARIMA trading ETH

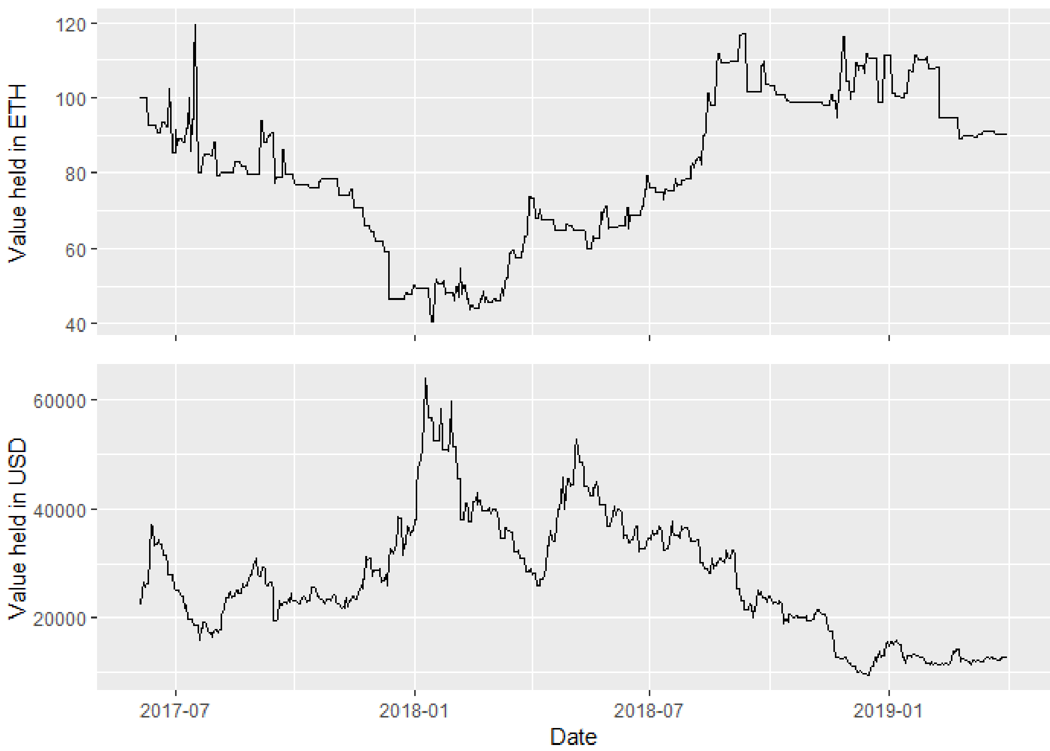

approximately 50%. Most of the increase in BNB held happens in the post-crash period in 2018. Unlike other major currencies, which rapidly dropped, and then oscillated around a decreasing mean, Binance Coin grew at a stable pace in the first half of 2018. This allowed the algorithm to demonstrate an impressive rate of growth in holdings, which was only interrupted when BNB started its strong upwards trend in early 2019. Due to the considerable increase in holdings, value in dollars remains stable throughout the 2018 crash and bear market, and then surges to twice the initial investment over the course of two months.

Finally, the behaviour of the algorithm when trading Basic Attention Token is somewhat erratic. The BAT held reaches close to 60% profit several times, but always falls back to the initial endowment level. The periods of growth in BAT held correspond to the periods where the price of BAT in USD terms falls, and vice versa. At the very end of the trading period, BAT held sharply drops to roughly 70% of the initial endowment. This is likely caused by the algorithm underestimating the large surge in BAT price. The development of USD held when trading BAT is probably the most underwhelming of all the traded

Figure 5.3: ARIMA trading BNB

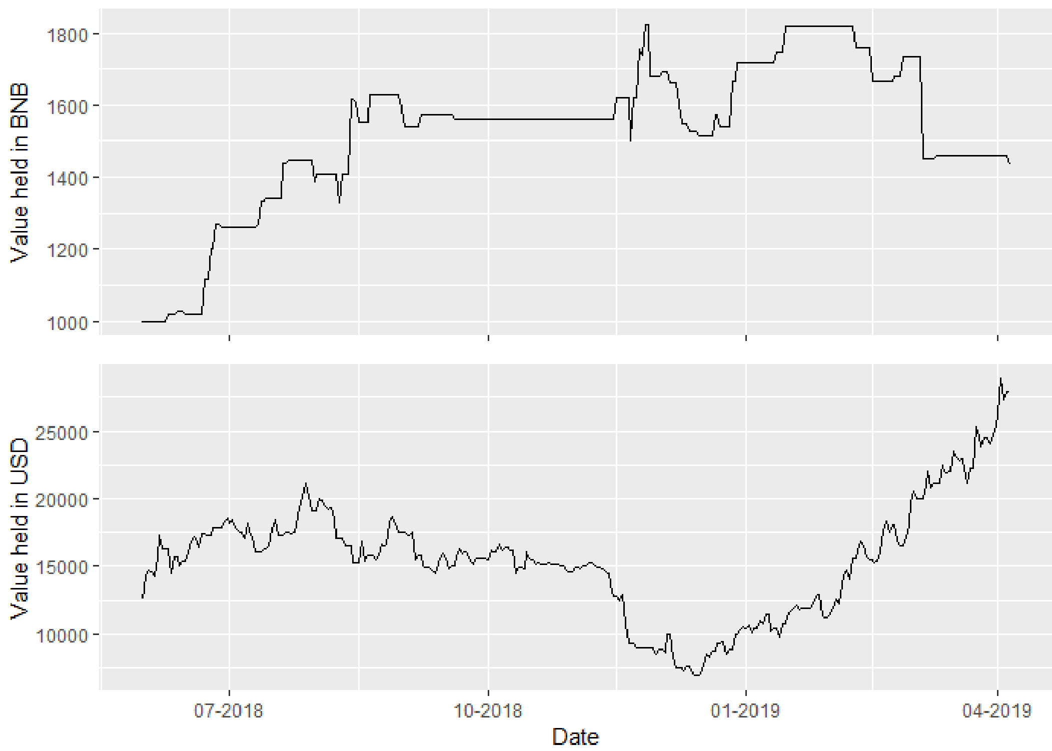

Figure 5.4: ARIMA trading BAT

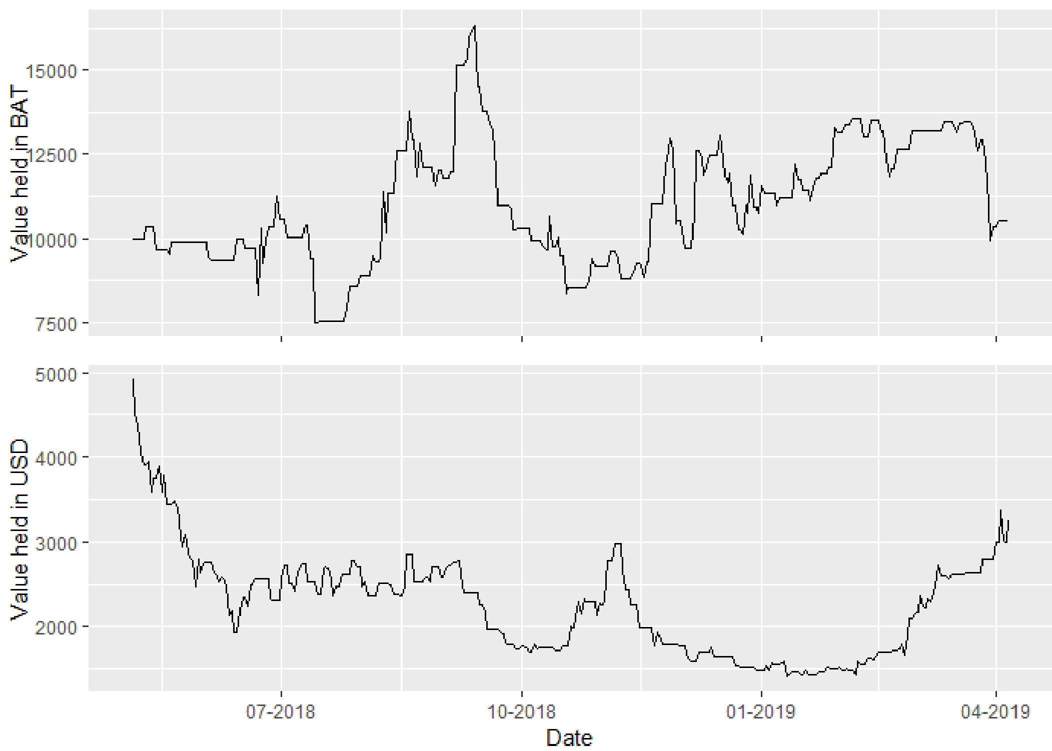

cryptocurrencies. Value held in USD drops below the initial investment at the very start, and then continues dropping, until it finally finds a bottom with an almost 60% loss. The smaller currencies were affected even more strongly than the large ones during the 2018 crash, and this is clearly visible on BAT. Drawing conclusions from the partial results above, ARIMA trading appears to be highly unreliable. The only currency where the ARIMA algorithm managed to produce consistent profits is Binance Coin, while in all three remaining cases, losses reach tens of percent, or peak value is significantly affected.

The main problem of ARIMA likely lies in its difficulties with the extreme price spikes symptomatic of cryptocurrencies. The predictions underestimate the speed of their growth, predicting negative price changes long before a crash actually happens, which leads to sales of the currencies held, and their subsequent repurchase once the strong upwards trends consolidate, losing significant portion of holdings in the process. The algorithm performs much better during periods where the currencies' price oscillates around a constant mean or a weak trend, such as the entirety of 2018, where the holdings increase sensibly. These results do imply potentially better outcomes for ARIMA-based intraday trading, as barring the nearly instantaneous and likely deliberate large jumps in 2018, the short-term volatility is much lower. In the following results, only the development of cryptocurrency holdings is commented on, as the prices, and thus the USD value held, remain relatively stable in these shorter timescales.

Unfortunately, the outcome for hourly trading is even worse than the one for the daily timescale. The volatility of returns is, of course, much lower. The algorithm reaches a high of 1% gain in the first two days, which is subsequently lost. After this, the Bitcoin held continues decreasing down to an approximately 1.5% loss at the end of the test period. For Ethereum, the hourly results are effectively identical to those of Bitcoin, with a slight increase in the first days, and a sharp drop afterwards.

When examined in context of the price action, these results are fairly easy to interpret. As per the expectations, the model cannot predict the sudden moves which frequently occurred in 2018. These drops and surges appear to be the results of either individuals or groups buying and selling massive amounts of the currencies, causing an avalanche effect which dissipates over several minutes to hours. Depending on the direction of the move, the price

Figure 5.5: ARIMA hourly trading

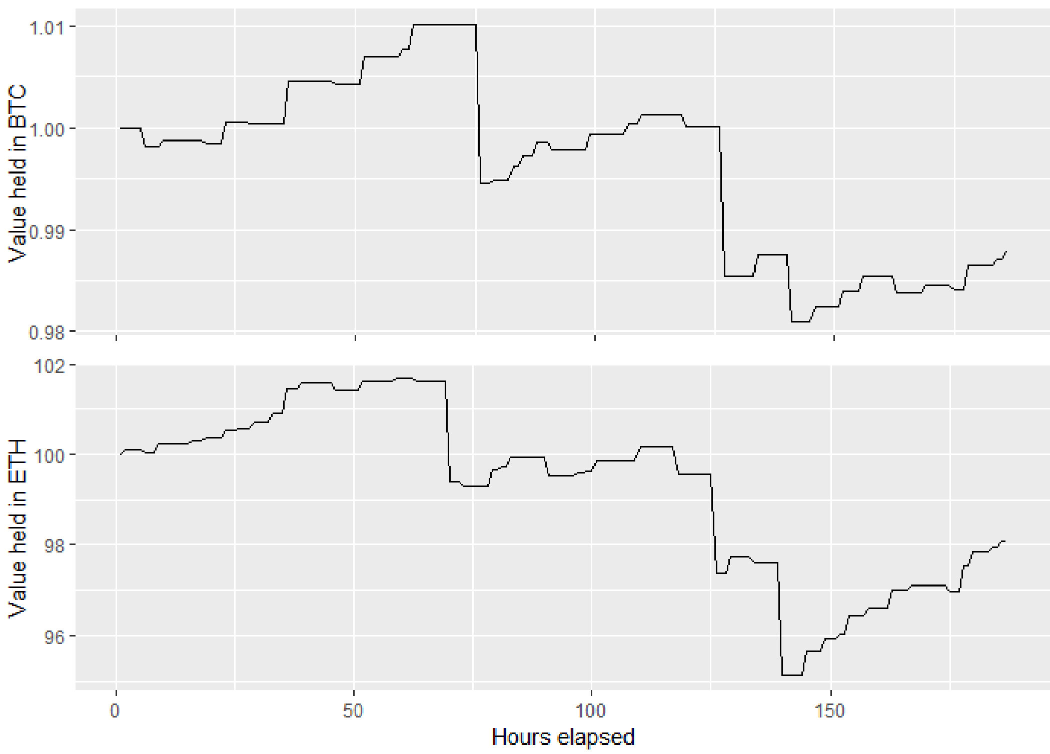

then stabilizes at a slightly higher or lower level compared to the original state. The majority of losses in holdings can be attributed to these jumps, which are extremely unlikely to be explainable through autoregression. Overall, the model and algorithm perform inadequately in almost every situation regardless of the timescale and the currency. The only exception is the larger timescale on Binance Coin, where the algorithm decisively outperforms a buy-and-hold strategy, although it is questionable to what degree this is a consequence of Binance Coin's favourable price development, and comparably low volatility. The algorithm consistently leads to losses in cryptocurrency holdings during periods of price growth and losses of USD value during periods of decline. As hypothesized above, this is likely caused by an underestimation of the trends' speed of growth and persistence in both directions. Nevertheless, cryptocurrency holdings do grow during downturns, which may cover and even exceed the losses incurred. The BNB trading results do imply that the model performs best when the price curve behaves linearly, ideally with a moderate slope.

Another promising approach is augmenting the ARIMA algorithm with an additional tool for estimating the length of a trend, so that the currency is not bought or sold prematurely. Trading volume can be incorporated as a variable, based on the assumption that as a trend nears its end, volume will decrease. A secondary condition for buying or selling could then be included in the algorithm.

5.2 Vector Autoregression

Trading Bitcoin using the VAR model proved to be highly inefficient. Bitcoin holdings drop throughout the 2017 bubble growth down to approximately 0.3 BTC, and no significant gain is made during the crash. The algorithm leads to a loss almost immediately. The initial losses can likely be attributed to the explosive growth before and during the bubble. While the results are very similar to ARIMA trading BTC, VAR does perform slightly better in this case, although the development is still almost strictly negative. This unimpressive

Figure 5.6: VAR trading BTC

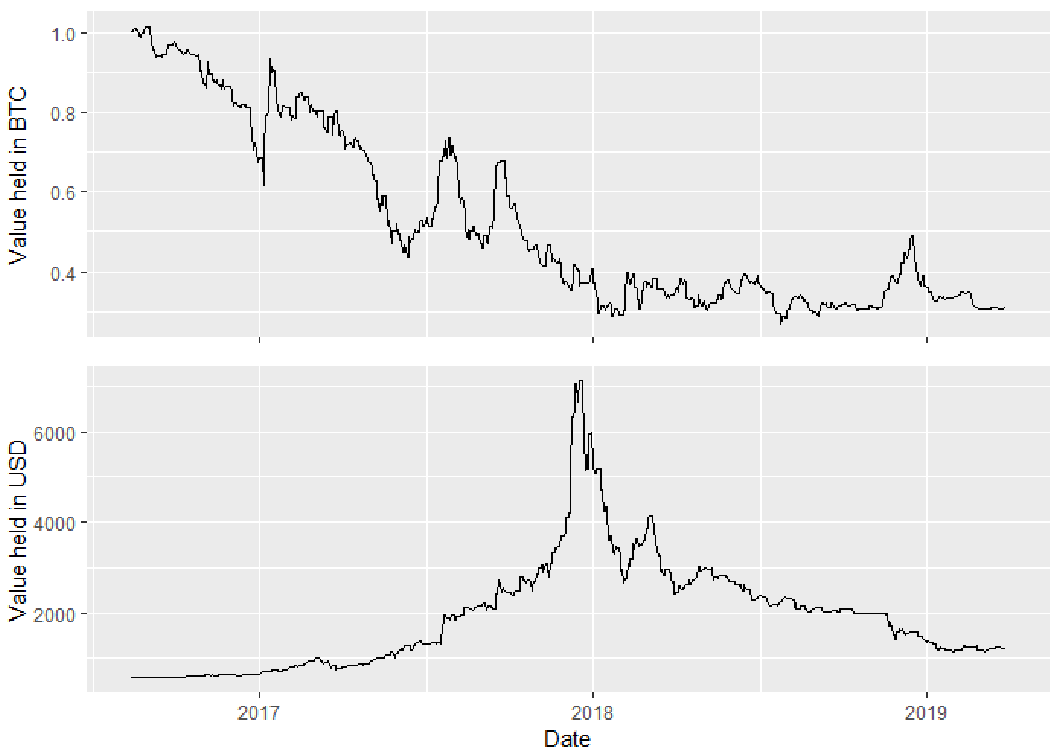

outcome can be clearly seen on the USD holdings. During the Bitcoin peak, $7000 is barely reached. As BTC held does not increase even in the crash, no decrease in USD volatility is achieved. There seems to be no reason for applying price and volume VAR to Bitcoin trading.

Ethereum VAR trading is dissatisfactory in absolute terms. However, the VAR results are again marginally better than the ones for ARIMA. While the holdings increase during the violent 2018 crash is initially slower using the VAR algorithm, the ETH held reaches a minimum of approximately 50 ETH, which is a slight improvement over the ARIMA bottom. This is counterbalanced by the performance at the trading period's end, where VAR reaches a 10% gain over initial investment compared to 20% with ARIMA. The USD value held

Figure 5.7: VAR trading ETH

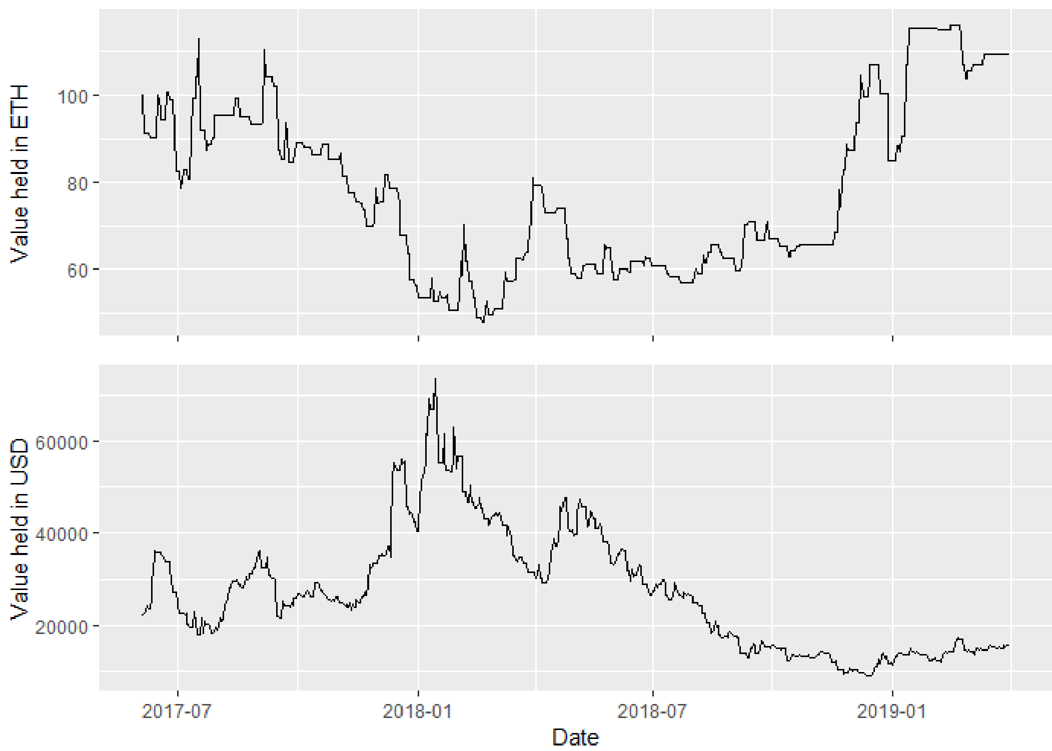

reaches a comparatively low peak due to the large loss in holdings, and as the algorithm outperforms a buy and hold strategy only marginally at the very end of the trading period, dollars held drop even below the pre-bubble levels. While the algorithm does lead to the 10% holdings gain at the very end, its subpar performance overall makes VAR an unsuitable choice for trading Ethereum.

The Binance Coin results for VAR are highly reminiscent of, but almost unequivocally worse than, those for ARIMA. Nevertheless, the algorithm reaches 1500 BNB representing a 50% gain, which is quite reasonable considering Binance Coin's absence of extraordinary crashes relative to other non-mainstream cryptocurrencies. Unfortunately, BNB's unexpected growth in the end of the trading period returns the holdings to the initial level. As BNB holdings re-

Figure 5.8: VAR trading BNB

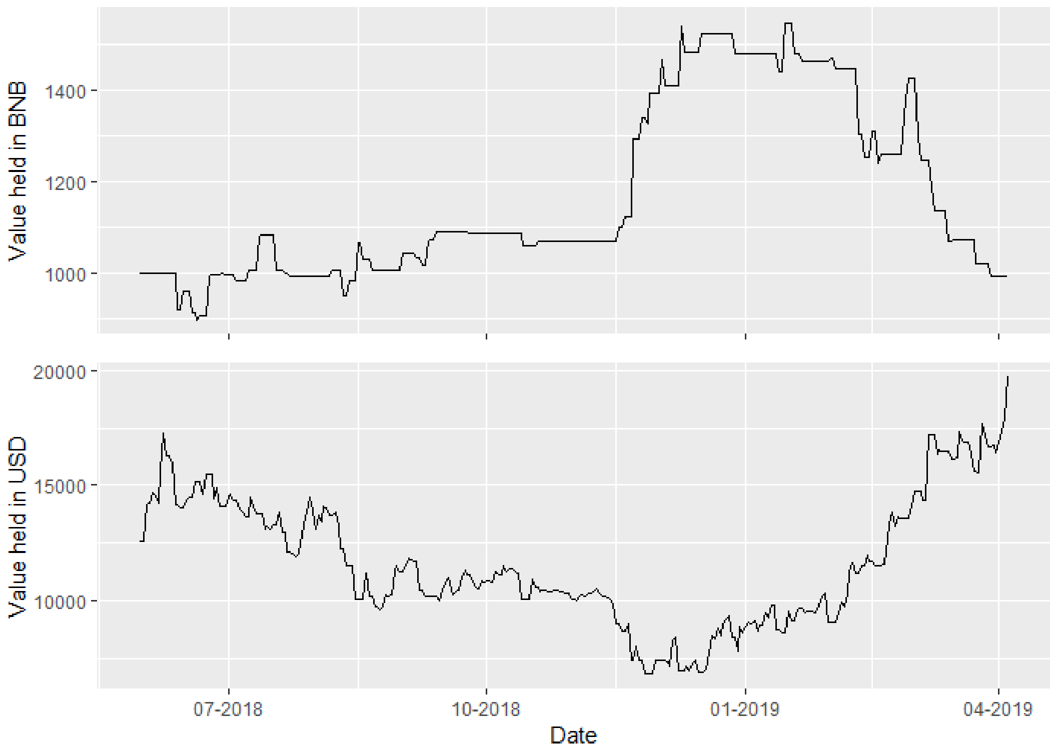

main mostly constant in the first half of the trading period, USD value drops by almost 50%. Again, it is important to note that this ranks BNB among the better performing currencies in the 2018 crash. The growth in both price and holdings in the second half cover these losses, ultimately reaching almost \$20,000. While the outcome is acceptable in absolute terms, the exceptional ARIMA results still indicate it as the better option.

Finally, rather unexpectedly, the results for trading BAT using VAR are very good in cryptocurrency holdings terms. Two peaks are reached, both crossing 20,000 BAT, corresponding to an over 100% gain, and while the first peak is eventually followed by a bottom at the period's midpoint with a minimum of 8,500 BAT, the algorithm returns to profitable levels quite swiftly. These results ironically do not imply a comparably positive outcome in USD holdings, as the vast majority of the period is spent in a loss. This can be slightly misleading, as the period's beginning coincides with a crash following one of BAT's frequent spikes. Overlooking the absolute profit, and assuming a more

Figure 5.9: VAR trading BAT

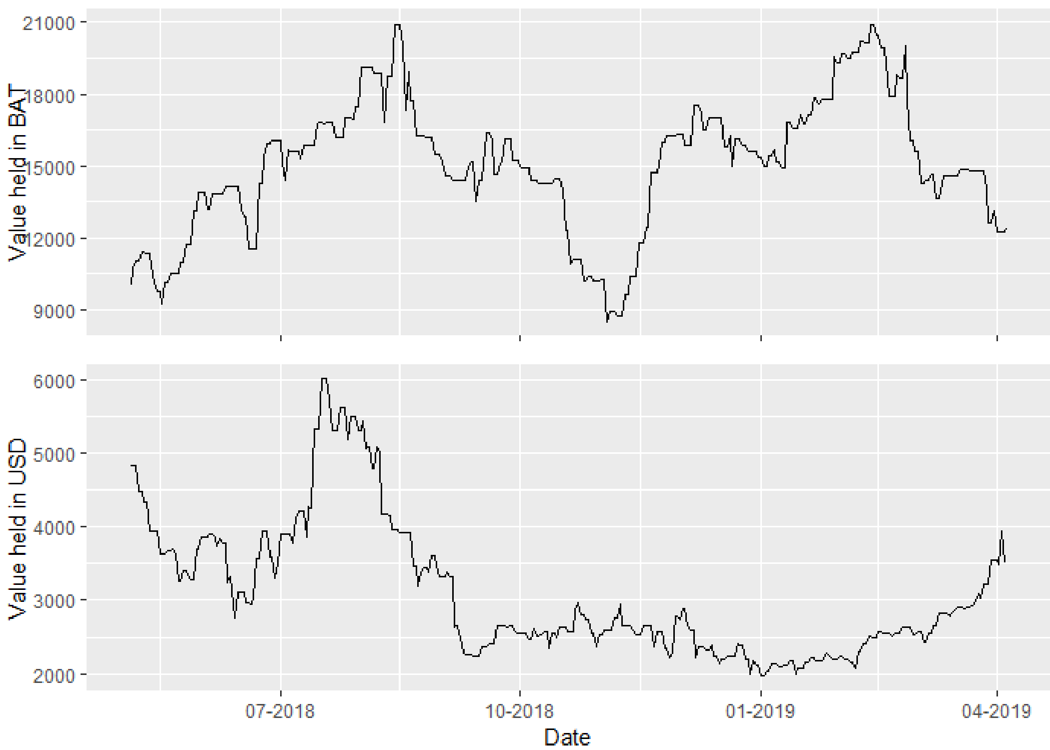

opportune introduction, the plot shows a maximum of over $6,000, a 20% gain even in comparison to the high entry point. In conclusion, vector autoregression performs quite well with Basic Attention Token, where it outperforms ARIMA decisively. For all the other currencies, the results are very similar to ARIMA, but effectively slightly better, or, as in the case of Ethereum, comparable and less volatile. As the addition of volume proved to be beneficial, further additions could perhaps lead to even better models.

5.3 Moving Average Crossover

Taking into consideration that the main issue with ARIMA and VAR forecasting was likely difficulties with cryptocurrency's volatility, the moving average crossover strategy is expected to perform slightly better, as it directly addresses this problem, specifically the issue of "irrationally" strong and persistent trends.

For Bitcoin, omitting a short-lived peak at 3.8 BTC, an unprecedented high of 3.4 BTC is reached, representing an over 240% gain, at the midpoint of the period, corresponding to January 2016, the approximate bottom of the bear market following the drop after the 2014 Mt. Gox scandal. Past this point, the holdings start decreasing down to a minimum of 0.64 BTC at the 2018 bubble's pinnacle. The trend is then reversed, and the algorithm slowly accumulates Bitcoin throughout 2018. In the USD value development, the BTC losses are

Figure 5.10: MA Cross trading BTC

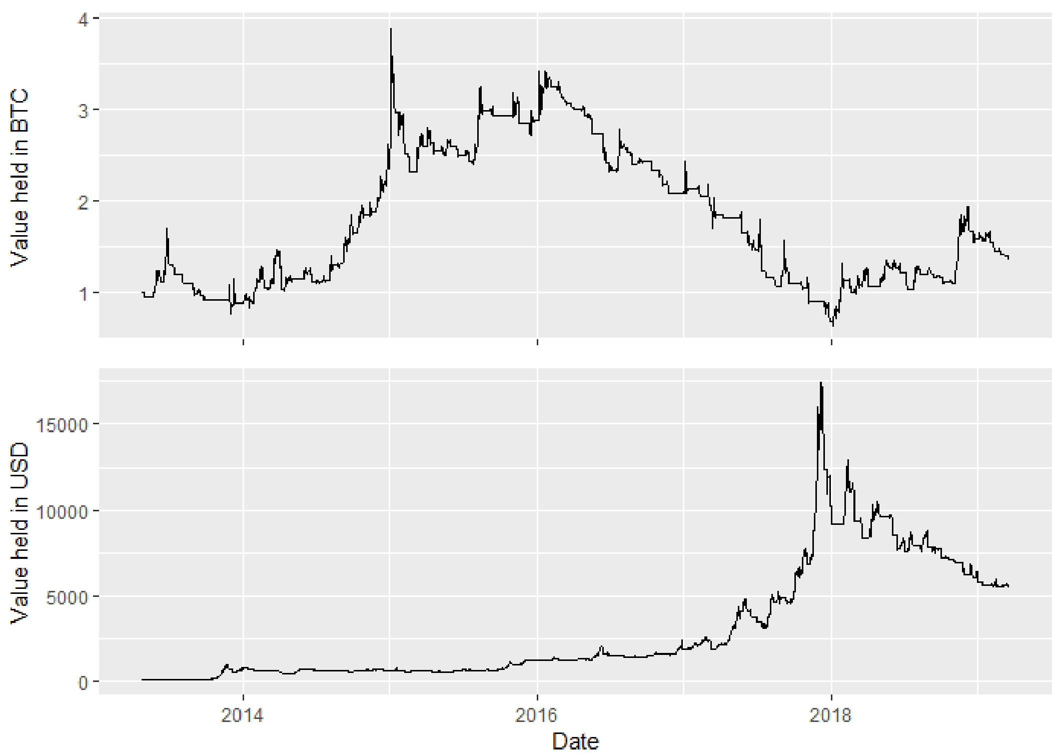

clearly seen, as the highest amount of USD held reached is only $17400. However, the algorithm performs well during and after the crash, causing the USD value to stabilize around $5500. These results imply that for trading Bitcoin, the MA crossover strategy could be efficiently utilised as a volatility and thus risk mitigating instrument. It is worth noting that the losses in cryptocurrency holdings would have been much higher without the extremely profitable period in 2015, which covered the losses in late 2017. Without it, the algorithm would perform much worse.

Surprisingly, while Ethereum generally behaves similarly to Bitcoin, the ETH moving average results are significantly different. From August 2015 to the $1400 peak in January 2018, the holdings fluctuate in the 100-200 ETH range. The minimum of 77.1 ETH corresponds to the period immediately following the peak. After that, holdings start growing rapidly, reaching almost 500 ETH during 2018. This impressive result can be traced to Ethereum's 93% drop in

Figure 5.11: MA Cross trading ETH

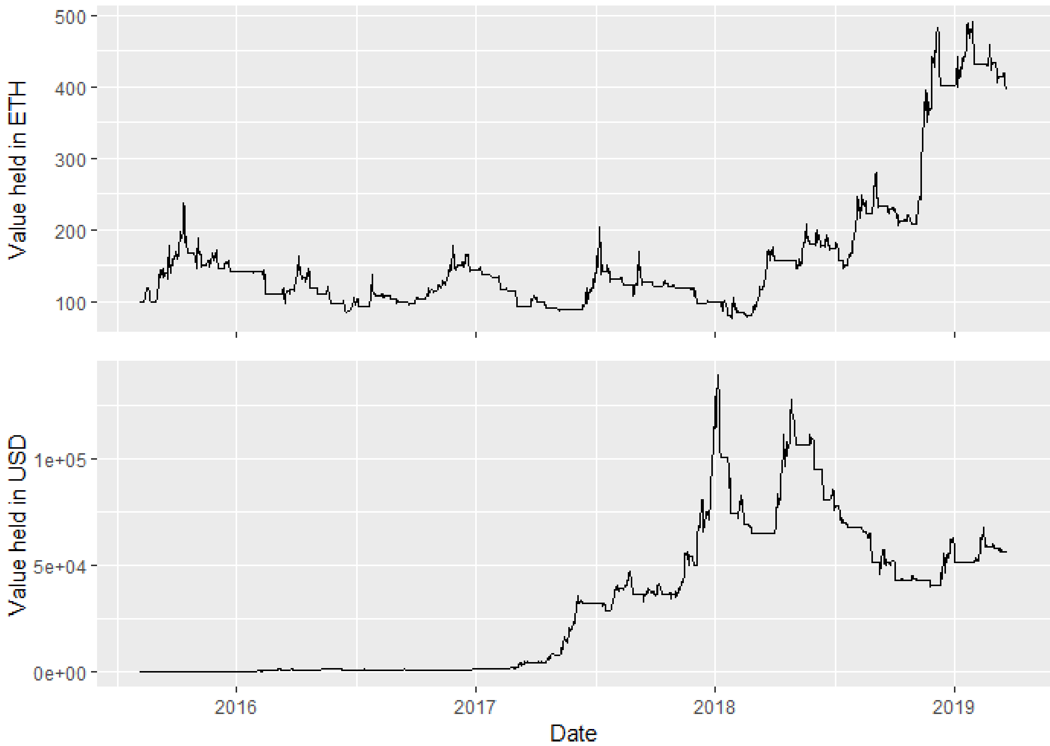

value in 2018. In spite of this, the USD value results show that the value held in dollars never drops below the levels immediately preceding the peak, and in fact surpasses them at the very end of the period. Ethereum apparently works soundly with the MA strategy. Compared to the acceptable but mediocre results for Bitcoin, Ethereum's volatility and the suddenness of its movements are well attuned to the principles of MA. The algorithm managed to maintain holdings in the growth period, and led to considerable gains in the crash.

In contrast with these promising results, the Binance Coin performance is much weaker. The holdings curve only shows gains in the very beginning, reaching an almost 100% increase over initial endowment, although these should be interpreted with a degree of scepticism, as at this point, the currency was still a very recent addition to the market. Past this point, the holdings shrink below 500 BNB, and never even reach the initial level, ultimately dropping below 250 BNB, representing a 75% loss. This disappointing result may be a consequence of the same attributes that allowed Binance Coin to synergize with ARIMA - the relatively low volatility, gradual rises and drops, and no significant spikes.

Figure 5.12: MA Cross trading BNB

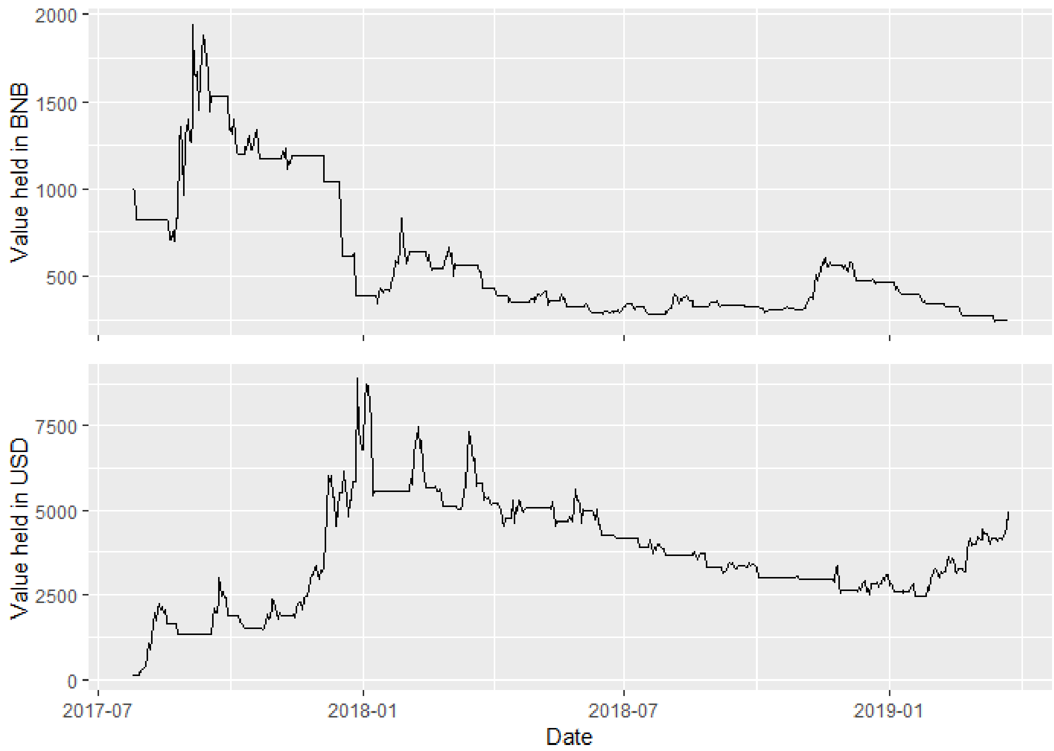

However, in spite of the large losses in holdings, the USD curve shows a tolerable development. Regardless of this, there seems to be little reason for using MA on Binance Coin in light of its phenomenal results with ARIMA.

The outcome for Basic Attention Token is rather inconclusive, and demonstrates a moderate degree of volatility. In a manner similar to trading BNB, the holdings reach over 100% gain, drop shortly after, reach a bottom, and start increasing at a slow pace. Over the course of the trading period, holdings mostly remain in the 7500-15000 range, ultimately reaching a slightly profitable level of 10,442, with no readily visible trend. As a consequence of the stable trading results, the USD value curve unsurprisingly closely follows the closing price curve. The currency shows a fairly consistent pattern of reaching a momentary peak and rapidly dropping in value. Unlike many other cryptocurrencies, especially in the smaller currency group, these cycles happen quite frequently, and the drops in value are not absolutely catastrophic. With EMA_5 used instead of SMA_5 as the more dynamic MA indicator, the results for Bitcoin are almost strictly better. Although the plots understandably look similar, and a lower bottom is reached in the first 500 days, the more reactive

Figure 5.13: MA Cross trading BAT

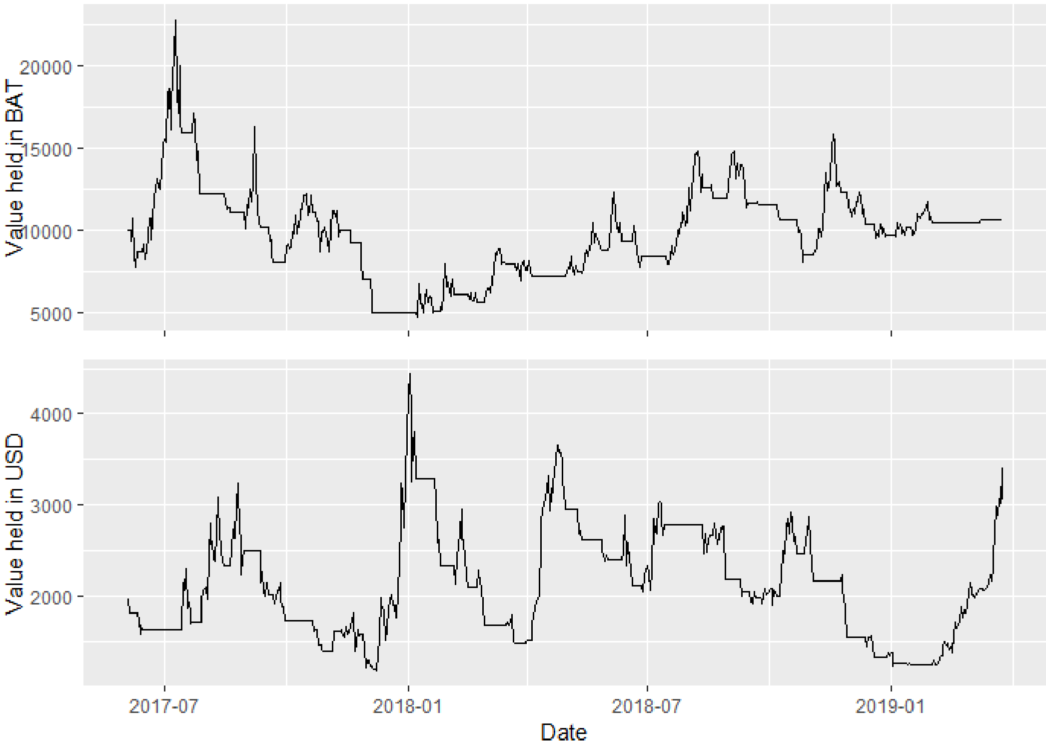

algorithm considerably improves on the first approach in the remainder of the trading period, reaching a higher peak, a higher post-crash bottom, and leads to a faster growth in holdings in 2018. The same distinction can be seen in Ethereum trading. The algorithm reaches over 600 ETH from the initial 100 ETH endowment, a considerable improvement when compared to the 491 ETH peak using SMA_5, and performs slightly better in general. Beyond this, the results are again comparable. Somewhat surprisingly, while the outcome is less impressive than the one for Bitcoin and Ethereum, compared to the disastrous SMA_5 approach, the exponential moving average leads to decent gains for Binance Coin. After the first peak at 2500 BNB, two more periods of holdings growth lead to approximately 50% gains. In spite of this, most of the period is spent in a loss. When compared to the superior ARIMA, there seems to be little reason for using MA in either variant for BNB. A marked difference between SMA and EMA is visible when trading BAT. After the almost identical first third, holdings rise past the 20,000 level, corresponding to a 150% gain, and stay in the 12,000 to 23,000 range, never dropping below initial investment. Overall, the results are again ambiguous. Of the four currencies traded, Binance Coin is obviously unsuitable for the MA strategy, Bitcoin and

Figure 5.14: EMA Cross trading BTC

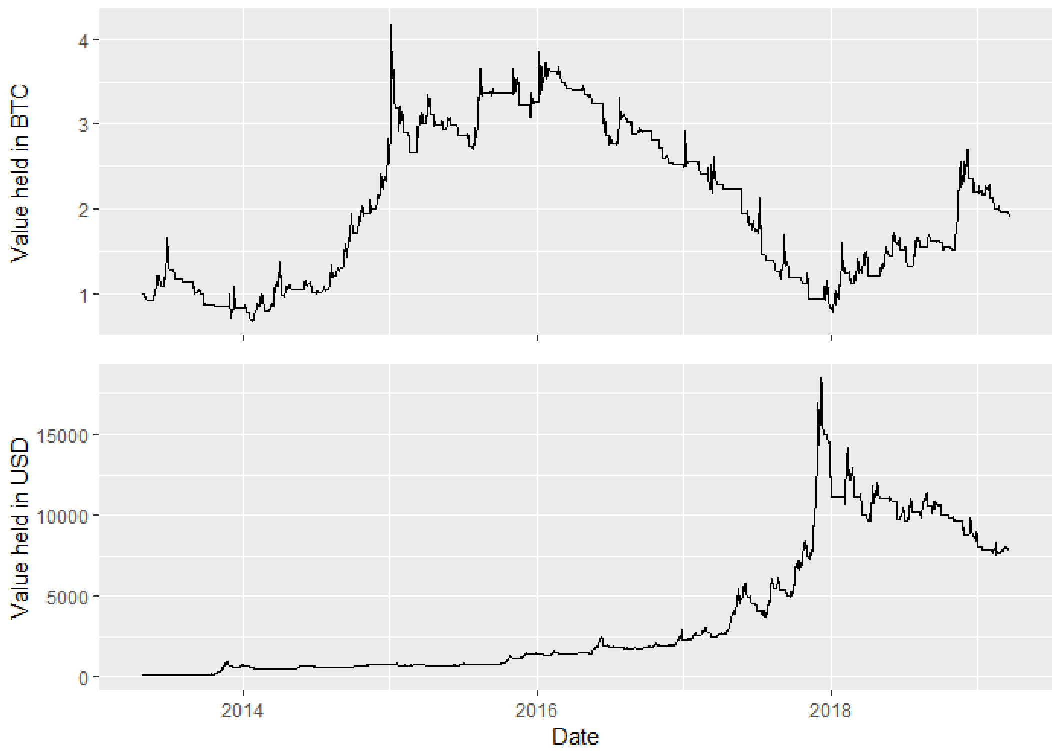

Figure 5.15: EMA Cross trading ETH

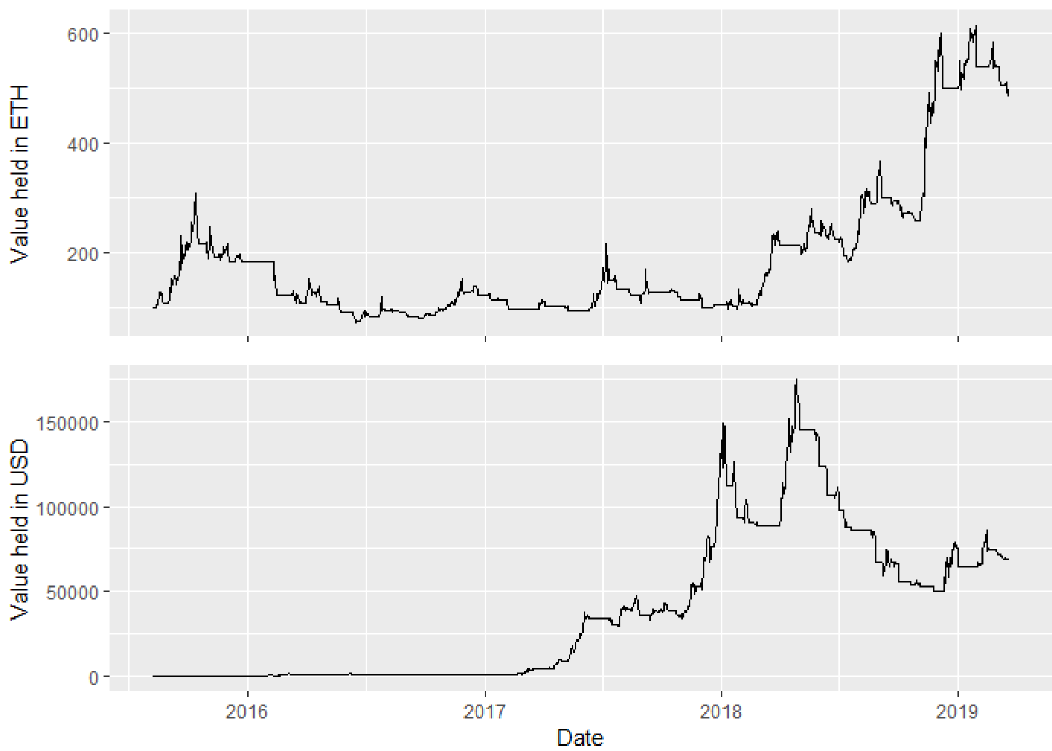

Figure 5.16: EMA Cross trading BNB

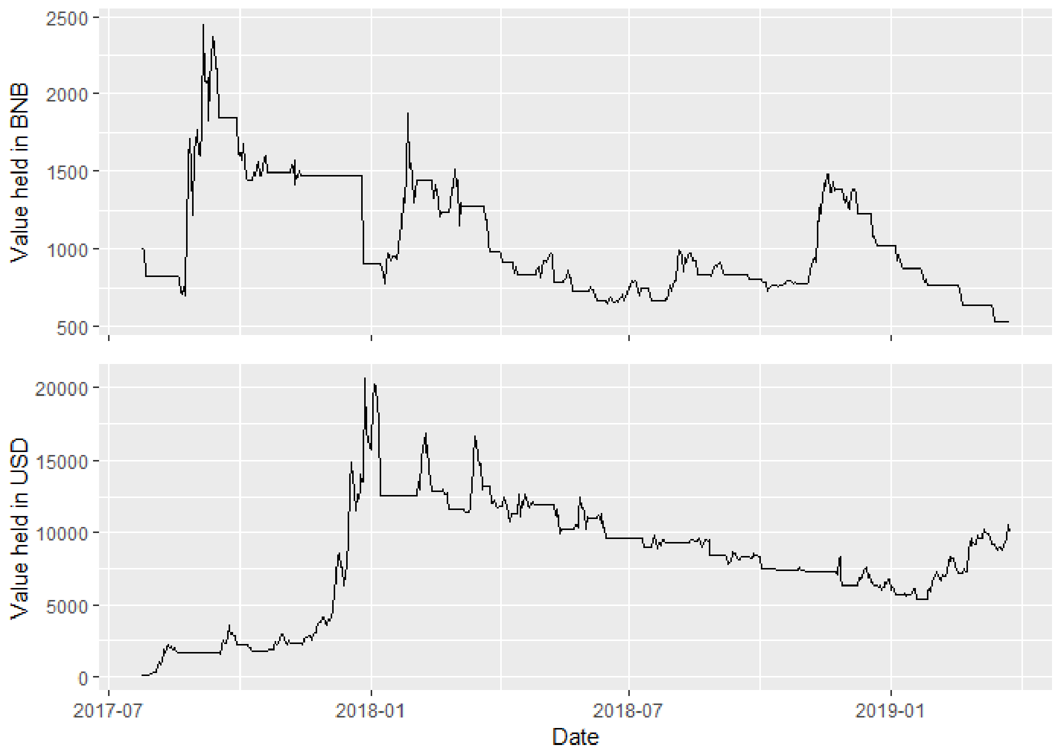

Figure 5.17: EMA Cross trading BAT

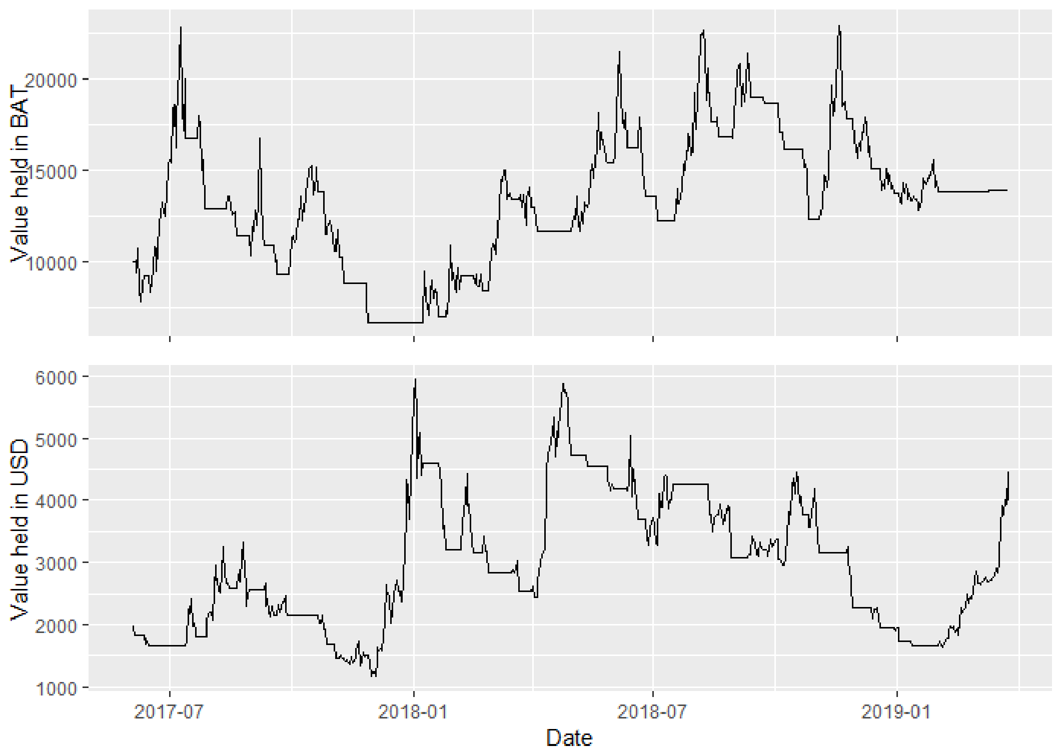

Basic Attention Token work moderately well, and the results for Ethereum are, rather unexpectedly, remarkable, especially considering the solid performance both during the period of rising price and the crash. When compared to the performance of the ARIMA and VAR model, each are seemingly suitable for different currencies. Between the SMA_5, SMA_{10} and EMA_5, SMA_{10} strategies, the exponential moving average approach led to higher gains and smaller losses in virtually every situation and for every traded currency. Extrapolating this to conclude that MA approaches tending towards more dynamic components are necessarily better might be too ambitious in a market with this level of heterogeneity, but it does imply that cryptocurrency price trends are rather short-lived compared to foreign exchange and the stock market.

5.4 Granger Causality

In accordance with expectations, no strong Granger causality was found either between gold and Bitcoin, or S&P 500 and Bitcoin, in either direction. Surprisingly, the only level of p-value significant at the 95% significance level was found in the model testing Bitcoin Granger causing gold. This can likely be attributed to spurious results, as it is improbable that movements in Bitcoin price would cause changes in the price of a market as large and mature as gold, especially specifically with a three day lag. The p-values of Granger tests with S&P 500 and gold causing Bitcoin never drop below 0.2 in the tested lag range, eliminating the need for further examinations of the profitability of a strategy based on Granger causality, as no significant causality was found. However, the possibility remains that there is a causal relationship between the aforementioned series, with the effects materializing in a shorter than daily timespan, for example on an hourly basis.

5.5 Comments

To summarize the findings above, no attempted strategy led to consistent gains on all four examined currencies. Quite unexpectedly, each of the strategies performed well for at least one of the currencies. While Bitcoin is traditionally assumed to be the main determinant of the entire market, these findings suggest that beyond the superficial, at least some currencies do have their idiosyncrasies, necessitating care when generalizing. While it may initially appear that algorithmic cryptocurrency trading is simply a matter of choosing appropriate strategies for each currency, the ARIMA on BNB results show that this might be a dangerous approach. Here, the algorithm did lead to decent and, more importantly, consistent gains as long as the curve more or less behaved linearly. However, the strong price uptrend in early 2019 led to a significant loss in holdings. Binance Coin is admittedly a young currency, and the full breadth of its uses, such as its usage for margin trading fees on its parent exchange, is still not unambiguously defined. It does however show that historical data are not necessarily a useful predictor of future behaviour for cryptocurrencies, not only in the sense of autoregressive price determination, but also for choosing which models are most likely to produce desirable results based on a currency's general behaviour.

An important finding is that most strategy-currency pairs finished at either a loss or very close to the initial investment. Still, many of these pairs also led to respectable gains at some points in the trading periods. This inconsistent performance essentially confirms that all the forms of cryptocurrency trading explored, buy and hold included, are better suited for risk-loving investors, possibly serving as a gambling substitute. The successes and failures of the trading algorithms are obviously closely tied to the current market behaviour. Unlike foreign exchange and stock markets, where the price movement follows a relatively stable curve, and crashes or spikes are exceptions to the norm, the cryptocurrency price action is composed almost entirely of these "uncommon" situations. Where a model applied to the forex market may successfully use short term fluctuations around a general trend to increase holdings of a currency even as said currency rises in value on a larger scale, cryptocurrency prices often move with such speed and suddenness, that any sale during an uptrend necessarily causes a loss in cryptocurrency holdings, and any purchase during a downtrend a loss in USD value held. Assuming that the cryptocurrency market does behave periodically, which is currently uncertain, as only two full boom-bust cycles have emerged to date, training the models and trading exclusively in the inter-bubble periods may be a functional approach, although it necessitates a secondary tool for forecasting the existence of a bubble, a problem in its own right. Trading within the bubble periods appears inadvisable, as a single incorrect forecast may lead to extreme losses due to the volatility.

In light of these findings, the unsatisfactory results for ARIMA trading on the hourly timescale are fairly peculiar. A closer examination of the holdings plots reveals that almost all of the losses incurred are consequences of nearly instantaneous jumps. As was previously discussed, these jumps may potentially be caused either by an avalanche effect, where a minor decrease or increase in price leads to several waves of panic sales or purchases, which is further reinforced for future instances as additional traders are conditioned to expect these movements and react accordingly. Another possibility is that they are simply caused by manipulation. Of course, these hypothetical causes are not mutually exclusive. Without these occurrences, the algorithms perform reasonably well, leading to consistent increases. In spite of this, it still seems imprudent to use the hourly model and algorithm during the bubble periods, as training the model on both the bubble and inter-bubble data would lead to worse performance in both cases, and training exclusively in the bubble periods requires the usage of data from previous bubbles, which may behave differently.

Chapter 6

Conclusion

While no clear winner arose within the models and algorithms used, some patterns did emerge. It is interesting to note that while no approach performed well for all the examined currencies, every approach led to positive outcomes for at least one currency, showing that while Bitcoin price strongly influences or at least correlates with the prices of most other currencies, simply identifying the price development of every currency with Bitcoin when testing cryptocurrency trading models. Overlooking the unusable results for Granger causality, the MA crossover generally led to the best results. ARIMA led to impressive gains in Binance Coin trading, and acceptable gains for Basic Attention token, ARIMA applied to Ethereum trading only outperformed a buy and hold strategy by a narrow margin while also spending a significant portion of the trading period in a loss, and Bitcoin trading caused major losses in holdings without any redeeming features. The reasoning behind ARIMA's disappointing performance is the significant difference in volatility between cryptocurrency bubbles and inter-bubble periods, which causes the model to forecast negative price development during the most profitable periods, where a premature sale and a later repurchase lead to large losses, and vice versa for the crashes. Trading on the hourly timescale is hampered by the existence of fast price jumps. Without these, the results are quite favourable. The results for VAR were in general slightly better than those for ARIMA, the reasoning for its questionable success being essentially identical to ARIMA. VAR performed better in Bitcoin and Basic Attention Token trading, while Binance Coin trading led to a worse outcome. The main finding here is that volume does improve the predictions in most cases, in spite of its debatable effect on the direction of price development, which suggests potential model improvements by including further variables.

Between using SMA_5 and EMA_5 as the more dynamic moving average, with SMA_{10} serving as the stabler one in both cases, EMA_5 performed slightly better, although the results are comparable. The success of the MA crossover approach can be traced to its ability to process the extreme volatility and strong trends, which proved to be the main obstacle for both ARIMA and VAR. Ethereum was the most MAC-compatible currency, while MAC Binance Coin trading was outperformed by both ARIMA and VAR. Granger causality unfortunately failed, as causal relationship was found neither between gold and Bitcoin prices, nor S&P 500 and Bitcoin prices.

Another important conclusion is the qualitative difference between the bubble and inter-bubble periods. If a secondary tool is used for predicting the existence of a bubble, both ARIMA and VAR can be used somewhat successfully for trading during periods of lower volatility. Intra-bubble trading appears to be better suited to the MA crossover approach, or simply a cessation of trading. As an alternative not addressed in this work, tools used to predict bubble bursts in forex and stock markets are another trading possibility for high volatility periods.

As for further questions regarding the use of autoregression and moving averages for cryptocurrency trading, the performance of VAR models may improve if more explanatory variables are included in them. These could range from fairly obvious ones, such as the price developments of other currencies, to more obscure ones, such as Bitcoin's share of total cryptocurrency market capitalization or the volumes of a currency's blockchain transactions. If predicting cryptocurrency bubbles can be feasibly done as was discussed previously, ARIMA and VAR can be exclusively trained and applied to the inter-bubble periods. Considering the favourable results of MA Crossover, combinations of different SMA and EMA levels may prove to be profitable, especially in long term trading. Some potential may also be hidden in different Granger causality approaches, either examining causality between other time series, or changing the timescale.

Bibliography

Akaike, H. (1974): "A new look at the statistical model identification." In "Selected Papers of Hirotugu Akaike," pp. 215–222. Springer.

Box, G. & M. Jenkins (1970): "Time series forecasting analysis and control." *San Franscisco: Holden Day* .

Box, G. E. & D. R. Cox (1964): "An analysis of transformations." *Journal of the Royal Statistical Society: Series B (Methodological)* **26(2)**: pp. 211–243.

Caporale, G. M. & A. Plastun (2018): "The day of the week effect in the cryptocurrency market." *Finance Research Letters* .

Colianni, S. G., S. M. Rosales, & M. Signorotti (2015): "Algorithmic trading of cryptocurrency based on twitter sentiment analysis."

Georgoula, I., D. Pournarakis, C. Bilanakos, D. Sotiropoulos, & G. M. Giaglis (2015): "Using time-series and sentiment analysis to detect the determinants of bitcoin prices." *Available at SSRN 2607167* .

Granger, C. W. (1969): "Investigating causal relations by econometric models and cross-spectral methods." *Econometrica: Journal of the Econometric Society* pp. 424–438.

Hyndman, R. J. & Y. Khandakar (2008): "Automatic time series forecasting: the forecast package for R." *Journal of Statistical Software* **26(3)**: pp. 1–22.

Kaminski, J. & P. Gloor (2014): "Nowcasting the bitcoin market with twitter signals." .

Karakoyun, E. & A. Cibikdiken (2018): "Comparison of arima time series model and lstm deep learning algorithm for bitcoin price forecasting." In

"The 13th Multidisciplinary Academic Conference in Prague 2018 (The 13th MAC 2018)," pp. 171–180.

Krištoufek, L. (2013): "Bitcoin meets google trends and wikipedia: Quantifying the relationship between phenomena of the internet era." *Scientific reports* **3**: p. 3415.

Krištoufek, L. (2015): "What are the main drivers of the bitcoin price? evidence from wavelet coherence analysis." *PloS one* **10(4)**: p. e0123923.

Makarov, I. & A. Schoar (2018): "Trading and arbitrage in cryptocurrency markets." .

McNally, S., J. Roche, & S. Caton (2018): "Predicting the price of bitcoin using machine learning." In "2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)," pp. 339–343. IEEE.

Nakamoto, S. (2008): "Bitcoin: A peer-to-peer electronic cash system." .

Pärlstrand, E. & O. Rydén (2015): "Explaining the market price of bitcoin and other cryptocurrencies with statistical analysis."

Phillips, R. C. & D. Gorse (2017): "Predicting cryptocurrency price bubbles using social media data and epidemic modelling." In "2017 IEEE Symposium Series on Computational Intelligence (SSCI)," pp. 1–7. IEEE.

Phillips, R. C. & D. Gorse (2018): "Cryptocurrency price drivers: Wavelet coherence analysis revisited." *PloS one* **13(4)**: p. e0195200.

Raymond McTaggart, Gergely Daroczi, & Clement Leung (2018): *Quandl: API Wrapper for Quandl.com.* R package version 2.9.1.

Ryan, J. A. & J. M. Ulrich (2018): *quantmod: Quantitative Financial Modelling Framework.* R package version 0.4-13.

Stigler, M. (2010): *Threshold cointegration: overview and implementation in R.* R package version 0.7-2.

Wickham, H. (2016): *rvest: Easily Harvest (Scrape) Web Pages.* R package version 0.3.2.

Appendix A

ARIMA ETH Algorithm

```
ethereum = 100  # Starting with one hundred ETH
dollars = 0  # and zero USD
for (i in 668:length(ethClose)) {
  prediction = forecast(ethClose[1:i], h = 1, model = ethARIMA)
  if ((prediction$mean[[1]] < ethClose[i]) &&
  (ethereum != 0)) {
    dollars = ethereum*ethClose[i]
    ethereum = 0
  } else if ((prediction$mean[[1]] > ethClose[i]) &&
  (ethereum == 0)) {
    ethereum = dollars/ethClose[i]
    dollars = 0
  }
  arimaTradeEth$eth[i-667] = ethereum
  arimaTradeEth$usd[i-667] = dollars
}
for (i in 1:length(arimaTradeEth$eth)) {
  if (arimaTradeEth$eth[i]==0) {
    arimaTradeEth$eth[i]=arimaTradeEth$usd[i]/ethClose[i+667]
  }
}
for (i in 1:length(arimaTradeEth$usd)) {
  if (arimaTradeEth$usd[i]==0) {
    arimaTradeEth$usd[i]=arimaTradeEth$eth[i]*ethClose[i+667]
  }
}
```

Appendix B

Detailed Result Tables

Note: Each pair of unseparated rows represents the results for one currency. The initial endowments are 1 BTC, 100 ETH, 1000 BNB, and 10,000 BAT.

Table B.1: ARIMA Results by Currency

Statistic	Mean	Min	Pctl(25)	Median	Pctl(75)	Max
BTC	0.49	0.186	0.2	0.4	0.8	1.055
USD	1,205.03	350.50	552.82	1,066.60	1,729.94	4,169.83
ETH	81.08	41	65.9	80.0	98.9	120
USD	27,304.67	9,418	19,522.4	25,834.5	34,676.9	63,968
BNB	1,525.53	1,000	1,450.4	1,559.5	1,667.1	1,826
USD	15,466.59	6,911	12,748.3	15,628.6	17,509.8	28,866
BAT	11,002.55	7,512	9,715.7	10,741.7	12,469.9	16,339
USD	2,259.76	1,409.00	1,726.21	2,311.89	2,620.27	4,915.80

Table B.2: VAR Results by Currency

Statistic	Mean	Min	Pctl(25)	Median	Pctl(75)	Max
BTC	0.52	0.270	0.3	0.4	0.7	1.016
USD	1,917.35	572.21	916.71	1,930.30	2,560.65	7,125.07
ETH	78.27	48	61.7	74.2	93.5	116
USD	27,172.15	8,859	14,904.3	25,972.7	35,004.2	73,616
BNB	1,168.45	900	1,006.2	1,074.6	1,380.1	1,547
USD	11,625.45	6,836	9,760.0	11,139.9	13,819.8	19,692
BAT	15,075.55	8,517	13,856.4	15,193.1	16,774.4	20,927
USD	3,090.27	1,969.46	2,418.21	2,661.15	3,628.16	6,007.54

Table B.3: MA Cross Results by Currency

Statistic	Mean	Min	Pctl(25)	Median	Pctl(75)	Max
BTC	1.78	0.639	1.1	1.6	2.5	3.884
USD	2,976.96	106.08	664.99	1,333.08	5,083.85	17,401.89
ETH	162.85	77	105.5	131.3	167.6	492
USD	31,090.87	103	1,215.8	23,528.5	54,984.7	139,259
BNB	585.52	243	329.2	396.1	638.6	1,946
USD	3,676.64	106.64	2,642.94	3,399.79	4,982.28	8,905.26
BAT	9,843.99	4,779	8,078.3	10,072.1	11,332.8	22,827
USD	2,124.28	1,176.82	1,638.52	2,068.39	2,571.90	4,437.49

Table B.4: EMA Cross Results by Currency

Statistic	Mean	Min	Pctl(25)	Median	Pctl(75)	Max
BTC	2.03	0.678	1.2	1.9	2.9	4.184
USD	3,585.61	106.08	686.18	1,512.18	5,839.02	18,455.76
ETH	187.60	73	104.6	127.6	216.9	616
USD	37,740.66	109	1,045.7	24,873.9	67,183.9	174,491
BNB	1,053.07	534	773.3	906.8	1,389.1	2,450
USD	8,039.51	106.64	5,703.01	8,387.98	11,112.48	20,639.80
BAT	13,309.10	6,664	11,128.5	13,548.0	16,005.7	22,990
USD	2,951.76	1,158.86	2,087.32	2,844.32	3,760.50	5,943.76