ERASMUS UNIVERSITY ROTTERDAM
ERASMUS SCHOOL OF ECONOMICS
MSc Economics & Business
Master Specialisation Financial Economics

Efficiency on the Cryptocurrency Market

The Predictability of Returns

Author: P. L. Huyshest
Student number: 356479
Thesis supervisor: Dr. J. J. G. Lemmen
Finish date: December 2018

PREFACE AND ACKNOWLEDGEMENTS

The foundation of this research topic originates from my interest in finance and coding. The creation of cryptocurrency has drawn the attention of a sizable amount of new type of investors. As of 2018, the majority of people has heard of the term cryptocurrency, yet are unaware of what it exactly constitutes. As a Finance student, I wonder how well traditional finance theories perform in such a dynamic market. I hope this research provides new insights into cryptocurrency and its technological implications within financial settings.

I first would like to thank my thesis supervisor, professor J. J. G. Lemmen, for narrowing down the topic, and for guiding me through the process of writing this thesis.

On a special note, I would like to take this opportunity to express my gratitude towards my parents for the love and support they have given me throughout life. I would not have been able to accomplish this without them.

Pascale Huyshest

ABSTRACT

This research analyzes weak form market efficiency on five cryptocurrencies. The analysis is defined in two ways. First, the research uses historical return observations to examine whether the return series are formed by a random walk process. The following tests are executed: the Ljung-Box test; the Augmented Dickey-Fuller test; the Phillips-Perron test; the runs test; the Brock-Dechert-Scheinkman test; and finally, an autoregressive process of order 5 is established. Ethereum, Ripple, Litecoin, and Monero are not in line with the random walk theory. The outcome for the Bitcoin market is somewhat debatable, as some tests show consistency with the random walk theory, whereas other tests do not. The cryptocurrency market does not reside in a state of efficiency. Second, trading volume data is taken into account. The objective is to find out if volume data is useful for the prediction of future returns. The results claim that volume data does not Granger-cause returns. Historical volume observations do not contain additional information beyond that what is implicit in historical return observations. This research shows that technical analysis is profitable when investors trade on historical price observations. Yet, trading on volume is not a worthwhile activity.

Keywords: Cryptocurrency, Market Efficiency, Return Predictability, Trading Volume

JEL Classification: G12, G14, G17, G19

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1. Introduction

This research analyzes the degree of efficiency on the cryptocurrency market, by looking at five cryptocurrencies that are among the top ten cryptocurrencies with the largest market capitalization. The objectives of this research are twofold. On the one hand, it focuses on return predictability by looking at historical values. On the other hand, the research tries to determine whether or not there exists a clear relationship between trading volume and returns.

In 2008 Satoshi Nakamoto introduces Bitcoin, world's first fully-operating cryptocurrency. In his whitepaper Nakamoto (2008) illustrates the disadvantages of electronic payment systems as we know it. The whitepaper presents an answer on how to eliminate interference from banks and governments – Nakamoto achieves this by implementing blockchain technology in an innovative way. In this regard, Bitcoin has paved the way for other cryptocurrencies. Now, a decade after the creation of Bitcoin, there are over 1,900 cryptocurrencies as listed on the Coinmarketcap website.

The cryptocurrency market is still very much in its infancy, and it has not been widely researched yet. There is much room for improvement into a better understanding of its underlying dynamics. This research sets out to examine the relationship between information and returns, which is defined in two ways. First, the objective is to find out whether historical return observations can help predict future returns. If investors are capable of forecasting returns with use of past prices, the market is said to be weak form inefficient according to the Efficient Market Hypothesis (Fama, Efficient capital markets: a review of theory and empirical work, 1970). This inefficiency arises due to the predictability of returns, thus creating the opportunity to beat the market. According to the EMH, in a state of efficiency no investor is capable of outperforming the market by using informational trading strategies. Second, the research looks at the relationship between historical trading volume and returns. Although exchanges display trading volume alongside the price data, it is a rather overlooked measure in analyses. Even though a market may be efficient in terms of past prices, investors still seem to be making profits by applying technical analysis. The price statistic may not provide a full explanation. Researchers have found volume to be positively correlated with price changes, but they hardly ever captured the reason behind this relationship. Blume, Easley, and O'Hara (1994) and Antoniou, Ergul, Holmes, and Priestley (1997) uncover volume data to play a crucial role in the flow of information to the market. For some investors, the price is fully revealing. Technical analysis is not required, as these investors can deduce the correct value just by looking at the price sequence. However, not all investors are capable of making a distinction between news and noise solely from the price. For these less skilled investors, trading volume contains information on the value placed by skilled investors.

Finance literature on cryptocurrencies is mostly limited to Bitcoin. Urquhart (2016) is the first to examine weak form efficiency on the Bitcoin market by executing a variety of tests. The test results conclude that the Bitcoin market is significantly inefficient. Nevertheless, it might transition into becoming more efficient over time. Contrarily, Bartos (2015) concludes that the price of Bitcoin is consistent with efficiency, i.e., the current price of Bitcoin is instantly adjusted to the underlying information set. This means no investor can beat the market using trading strategies based on this information. Furthermore, Balcilar, Bouri, Gupta, and Roubaud (2017) are among the few to analyze the relationship between trading volume and returns on Bitcoin. The test results suggest that volume can indeed help predict future returns. However, this paper concludes that predictability only exists when the market is in its regular, normal mode. In the event of a bearish or bullish market, volume does not seem to hold any predictive power. Therefore, in these markets, volume-based trading strategies are meaningless.

As mentioned previously, the majority of studies examine predictability with regard to Bitcoin. This research adds to prior literature by looking into five different cryptocurrencies. Furthermore, most studies merely look into the relationship between the current price and its past values. This research also takes trading volume into account. The main research question is as follows:

Is the cryptocurrency market efficient according to the Efficient Market Hypothesis?

First, the research looks at general weak form efficiency, with the following null hypothesis:

H_0: The process generating the return sequence is formed by a random walk.

Weak form market efficiency is derived from the random walk theory, in which the process that generates prices is said to be random. In case this null hypothesis holds, the future price is unforeseeable. Investors therefore cannot derive future prices from past values. This null hypothesis is examined with a selection of the following tests: the Ljung-Box test for serial correlation (autocorrelation); the Augmented Dickey-Fuller test and the Phillips-Perron test for the detection of unit root; the runs test and Brock-Dechert-Scheinkman test for serial independence; and finally, an autoregressive model is established based on the 5-day lagged historical return. Next, this research looks at the relationship between trading volume and returns.

H_0: Trading volume data is a useful predictor of future returns.

This null hypothesis is concerned with whether the movement of trading volume and price changes are somehow related to one another. This null hypothesis is tested by executing the Granger-causality test. The Granger-causality test analyzes whether or not the trading volume sequence is useful for the prediction of future returns.

The data set consists of the daily prices and daily trading volume of the following cryptocurrencies: Bitcoin (BTC), Ethereum (ETH), Ripple (XRP), Litecoin (LTC), and Monero (XMR). The data is retrieved from cryptocompare.com. The sample period starts at September 01st 2015 and ends at August 31st 2018. Additionally, the full sample is split in two subsamples, in which each subsample contains 18 months' worth of data.

#	Name	Symbol	Market Cap	Price	Circulating Supply	Volume (24h)
1	Bitcoin	BTC	$108.910.055.374	$6.330,18	17.204.887	$4.034.916.364
2	Ethereum	ETH	$32.824.033.748	$324,11	101.273.047	$1.781.767.588
3	XRP	XRP	$12.024.547.414	$0,305969	39.299.874.590 *	$247.620.263
4	Bitcoin Cash	BCH	$9.901.141.894	$572,71	17.288.203	$342.393.661
5	EOS	EOS	$4.622.012.669	$5,10	906.245.118 *	$666.661.321
6	Stellar	XLM	$4.196.097.052	$0,223537	18.771.403.505 *	$85.404.228
7	Litecoin	LTC	$3.477.327.697	$60,15	57.810.763	$261.377.072
8	Cardano	ADA	$2.944.339.443	$0,113562	25.927.070.538 *	$59.475.350
9	Tether	USDT	$2.413.415.146	$1,00	2.407.140.346 *	$2.556.249.493
10	Monero	XMR	$1.540.798.405	$94,72	16.266.706	$29.124.414

Fig. 1. Snapshot of the top 10 largest cryptocurrencies, dated August 12th 2018 (*www.coinmarketcap.com*)

The results of the first null hypothesis show that the cryptocurrency market is relatively inefficient. Ethereum, Ripple, Litecoin, and Monero reject (nearly) all of the diagnostic tests of randomness. Bitcoin barely passes half of the tests. This indicates that Bitcoin is neither strictly efficient, nor strictly inefficient. With regard to the second null hypothesis, the Granger-causality test cannot be rejected. This statement implies that trading volume data is not a useful indicator for the prediction of future returns. In general, it seems likely that investors are capable of outperforming the market by trading on the information implicit in past prices. However, the information implicit in trading volume seems to be incorporated in the price. Trading on volume does not lead to any excess return.

The layout of this research is the following. Section 2 explains some of the basics of blockchain technology and cryptocurrency. Section 3 covers the theoretical framework with respect to the EMH and the relationship between trading volume and returns. Section 4 provides a review of finance-related cryptocurrency literature. Section 5 contains an in-depth view of the data and methodology. Section 6 presents the results of this research. Last but not least, section 7 includes a final conclusion, discusses some drawbacks, and presents some suggestions for future research.

2. Cryptocurrency and Blockchain Technology

2.1. The Essence of Cryptocurrencies

Cryptocurrencies are virtual currencies which are used for electronic payments. The term *crypto* is derived from the fact that the currencies are encrypted for security purposes. Cryptocurrencies are built upon blockchain technology. The blockchain resembles a digital ledger, or database, in which pieces of information (*blocks*) are recorded. The database is distributed to all participants, and is continuously updated whenever a block is added to it. The blockchain is therefore a decentralized distributed system. The database is synchronized throughout the entire network, which means that every participant possesses the exact same version. The primary objective of blockchain technology is to enhance transparency, without putting a central organization in place (Böhme, Christin, Edelman, & Moore, 2015).

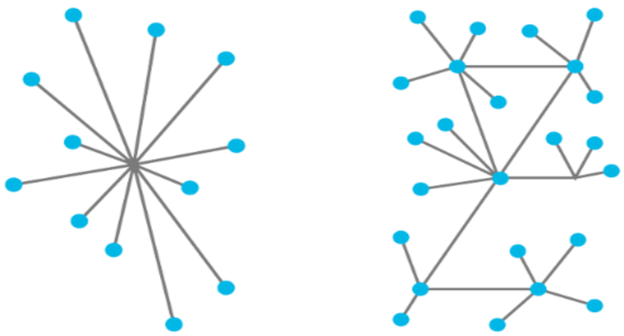

Fig. 2.a. A centralized system (left) vs. a decentralized system (right) (*www.blockgeeks.com*). In a centralized system, there is one single authority that has control over the decision-making process. This authority keeps track of the database. In a decentralized system, the decision-making process is executed at multiple levels.

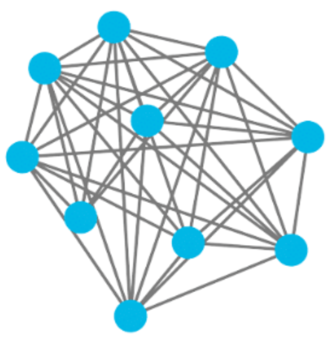

Fig. 2.b. A decentralized distributed ledger (*www.blockgeeks.com*). In a distributed system, all participants own the exact same database. Whenever a new block is added to the blockchain, all participants receive an updated version of the ledger.

2.2. Bitcoin: World's First Operating Cryptocurrency

The popularity of cryptocurrency has risen substantially over the past years. Nonetheless, there still exists quite some confusion on what cryptocurrency is. In *Bitcoin: A Peer-to-peer Electronic Cash System*, Nakamoto (2008) describes some issues the current digital world faces when making online payments. These issues arise due to interference of banks and governments. Prior to the implementation of Bitcoin, all electronic payments were regulated by trusted third parties, such as banks. The main tasks of these institutions consist of facilitating the payment process, acting as mediators, and minimizing the amount of fraud. In order for the electronic payment system to run smoothly, these third parties ought to be aware of all payments that go through the system. Whenever a bank receives an online payment request, the bank ensures that the money is deducted from the bank account of the sender, and is added to the bank account of the recipient. Intermediaries have to prevent double-spending from happening[1]. Whereas giving full authorization to third parties works well, it also poses some problems. The necessity of installing an intermediary to prevent fraud by double-spending leads to the following disadvantages. Intermediaries cannot avoid mediating in conflicts between the sender and recipient. These conflicts indicate it is not possible to make irreversible transactions online. If a transaction were to be permanent and if a conflict were to arise, there would be no need for an intermediary in the first place, as there would be nothing the intermediary could do to resolve the conflict. The expenses associated with mediation and the prevention of double-spending are reflected within the transaction costs, deeming it unfavorable to make small payments. Relying on these intermediaries turns the electronic payment system into an expensive and slow process (Nakamoto, 2008).

2.3. The Transaction Process

Nakamoto (2008) continues that Bitcoin offers the solution to the issues that come with double-spending and interference. Cryptographic proof replaces the trust previously put in third parties. The blockchain eliminates the reversibility of transactions. Irreversibility is a necessity in reinforcing the level of consensus within the network. A transaction is only successful if the majority of participants, or nodes, agree on the validity of the transaction. The only way a decentralized system is able to prevent double-spending, is for the majority of nodes to agree on the transactions taking place. First, the network receives a payment request.

[1] Double-spending is the act of spending the exact same currency more than once. This problem does not occur with physical currency: once a person hands over a coin or bill, he does not own it any longer, and he is not capable of spending it again.

With use of a timestamp server, the network stamps the exact time of occurrence on the request. The earliest timestamp is the one that counts, and each single time frame has its own unique stamp[2]. This ensures that all participants agree on one single timeline. After the request is stamped, the network broadcasts the payment request to all nodes. Next, the network implements a proof-of-work system. A proof-of-work is an algorithm, which requires nodes to put in a certain amount of effort, i.e., processing power (CPU), to solve a complex puzzle (the process of creating a block is also called *mining*). The complexity of each new proof-of-work increases exponentially (Böhme et al., 2015). This demonstrates how difficult it is to corrupt the network. If a hostile node were to attack the system – to steal back its recent payment – it would have to redo not only the proof-of-work of the payment it wishes to change, but also redo all proof-of-work that came after that transaction, and on top of that surpass the work of the nodes that are working on the current block. Whenever a node creates a block, it sends out the block to all other nodes. The nodes only accept the block if the entire transaction history embedded in it is valid, and not already spent. Afterwards, the block is added to the blockchain, and the network continues working on the next block (Nakamoto, 2008).

2.4. Incentives

The Bitcoin system rewards miners for expending CPU and electricity in two separate ways. The first one is by solving the proof-of-work to create entirely new Bitcoin. The creation of a new coin consists of transferring the coin from the 'mine' to the creator (a *coinbase transaction*) (Antonopoulos, 2015). Other than the creation of new coins, miners can earn rewards by facilitating the transaction process. During a transaction, senders pay a fee on top of the amount they are transferring. The miner to first complete the proof-of-work receives the transaction fee as an incentive. As opposed to ordinary currencies, the total Bitcoin supply is predetermined and fixed[3]. At a certain point in time, all coins will have entered circulation. Once this occurs, the incentives solely consist of transaction fees. Other than just rewarding participants, these incentives will also keep out hostile nodes. The only two options an attacker has, are to either steal back his recent payments, or using his CPU to cooperate with

[2] It is not possible to have the same stamp for different time frames. The timestamp for a specific date and time is converted into a number. For example, the UNIX timestamp server is based on the seconds that have passed since 12 a.m. January 01st 1970.

[3] This is not the case for all cryptocurrencies. There are currencies that have an infinite supply of coins. However, for the coins with infinite supply there is often a predetermined limit on how many coins there will be produced each year.

the network. As mentioned previously, stealing back his payments would not benefit the attacker at all, as he has to expend an enormous amount of CPU into redoing many proof-of-work puzzles. Therefore, the hostile node is better off being honest (Nakamoto, 2008).

This section summarizes the idea behind blockchain technology and the Bitcoin network. The implementation of Bitcoin has led to the invention of many other cryptocurrencies. The vast majority of cryptocurrencies apply roughly the same principles within their own blockchain. Generally, these other blockchain work in a similar way to that of Bitcoin.

3. Market Efficiency

The Efficient Market Hypothesis (EMH) considers a market to be efficient, when all available information is fully reflected in the price at any given point in time (Fama, 1970). The market shows the correct price according to the relevant available information. Investors are not capable of systematically earning any excess return by employing investment strategies that are based on this information.

3.1. The Study on Efficient Capital Markets

3.1.1. The Three Forms of the Efficient Market Hypothesis

The EMH consists of three separate forms of efficiency. Early adoptions of the model simply speak of market efficiency in terms of the full incorporation of information. Fama (1970) narrows down the general theory into three forms of market efficiency: weak form efficiency, semi-strong form efficiency, and strong form efficiency. There exists a level of dependency between the three forms, i.e., a market has to be weak form efficient and semi-strong form efficient in order to be strong form efficient.

In weak form efficient markets, the information set solely consists of historical prices. Generally, no other historical information is considered, just the lagged values of the price. In weak form efficient markets, investors cannot systematically realize any excess return with use of technical analysis. Weak form market efficiency implies that the process that generates prices is random. The sequence of past prices does not hold any predictive power. The majority of efficiency studies concern weak form efficiency. This research deals with weak form efficiency as well.

The second stage of market efficiency is semi-strong form efficiency. In semi-strong form efficient markets, the information set consists of all, past and current, publicly available information (e.g., news articles, annual reports, balance sheets). The main question is whether information is immediately incorporated into the price. Whenever new information arrives, the market reacts instantly, causing the price to adjust to its efficient level. Semi-strong form efficiency indicates that investors are not able to outperform the market by either technical or fundamental analysis. This type of efficiency implies that the only way an investor is able to realize excess return, is by engaging in insider trading.

Last but not least, the third form of market efficiency is the strong form. In this final form, privately held information is added to the information set. Strong form efficient markets assume all relevant information to be incorporated into the price. The act of trading on inside information does not lead to any systematically realized excess return.

The general outcome of the entire EMH is that efficiency indicates it is not possible to beat the market by expert stock-picking. The EMH, however, does not necessarily imply expected return to be zero. The EMH only suggests that it is not possible to systematically realize any extra return on top of the market return by employing trading strategies that are based on the underlying information set.

3.1.2. The Fair Game Model

Fama (1970) starts by clarifying what is meant when prices fully reflect all available information. The EMH statement about efficiency is too broad, and needs to be narrowed down in order to make it empirically testable. The formation of prices needs to be defined. Fama suggests to employ the Capital Asset Pricing Model (CAPM)[4] to compute equilibrium prices. Fama also explicitly mentions that there are several different theories to compute this, each of which assesses different measures of risk. However, whatever method chosen, they can all be described by the following equation:

$$E(\tilde{p}_{j,t+1}|\ \Phi_t) = [1+E(\tilde{r}_{j,t+1}|\ \Phi_t)]p_{j,t} \qquad (1)$$

where E denotes the expected value of a factor; $p_{j,t}$ and $p_{j,t+1}$ denote the price of asset j at time t and $t+1$, respectively; Φ_t is the information set and is assumed to be fully reflected in the price; finally, the future return, r_{t+1}, is calculated by $(p_{t+1} - p_t)/p_t$. The tildes atop the p and r indicate that p and r are random at time t. So far, the model assumes two things. First of all, the market equilibrium is specified in terms of expected return. Second of all, the expected return fully incorporates the relevant information set. These assumptions lead to the statement that capital allocation is a "fair game". It is not possible to beat the market and earn excess return based on trading on information. This relationship is illustrated by the following equation:

[4] The CAPM describes the relationship between market risk and return. Generally, an increase in risk is accompanied by an increase in compensation for taking this risk, i.e., riskier assets typically exhibit higher returns. The model distinguishes between systematic and idiosyncratic risk. Systematic risk affects the entire market, and cannot be eliminated. Idiosyncratic risk is stock-specific, and can be reduced. By holding a well-diversified portfolio, these stock-specific risks cancel each other out. In the event of a truly well-diversified portfolio, the investor is mostly exposed to systematic risk.

$$x_{j,t+1} = p_{j,t+1} - E(p_{j,t+1} | \Phi_t) \qquad (2)$$

This equation shows that $x_{j,t+1}$, which stands for the excess market value at time $t+1$, is equal to the observed future price minus the expected future price. In a fair game, the actual future price and the expected future price are exactly the same:

$$E(\tilde{x}_{j,t+1} | \Phi_t) = p_{j,t+1} - E(p_{j,t+1} | \Phi_t) = 0 \qquad (3)$$

Equation (3) explains that, for any investor, the expected future sequence of x equals zero (let r replace p, and we see that the excess future return should be equal to zero as well). Any type of trading strategy that is based on information set Φ_t results in an excess return equal to zero. On an aggregate level, the total excess market value, V, also adds up to zero.

$$V_{t+1} = \sum_{j=1}^{n} \alpha_j(\Phi_t)[r_{j,t+1} - E(\tilde{r}_{j,t+1}|\Phi_t)] = 0 \qquad (4)$$

where $\alpha_j(\Phi_t)$ denotes how much an investor should invest in strategy j, and in which j depends on information set Φ_t. The total future market value is simply the summation of each investor's gains for employing informational trading strategies. Therefore, on an aggregate level the total gain of the entire market is equal to zero.

$$E(\tilde{V}_{t+1} | \Phi_t) = \sum_{j=1}^{n} \alpha_j(\Phi_t)E(\tilde{z}_{j,t+1}|\Phi_t) = 0 \qquad (5)$$

Based on the information set, it is a fair game to all market participants. However, the fair game model in itself is insufficient to explain the EMH, as the model only states that the equilibrium price should somehow be established in terms of expected return. The model hardly says anything about the stochastic process that generates prices. Nevertheless, the idea behind the fair game model serves as a basis for the random walk theory. The random walk model enables us to make statistical comments about the stochastic process.

3.1.3. The Random Walk Model

As mentioned previously, weak form efficiency implies that prices are formed by a random walk process. Prior to Fama (1970) little was known about this stochastic process. Early versions of market efficiency merely observed a randomness of some sort within the formation of prices, but did not provide an answer to possible causes of this randomness. Past literature, even well before Fama's paper on the EMH, concluded that the subsequent changes in price are independent, and that the probabilities of these price changes are identically distributed (i.i.d.). These two assumptions form the basis of the general random walk model. Independence implies that the occurrence of an observation within the sequence does not affect the occurrence of other observations. In the presence of independence, the serial covariance between lagged values are equal to zero, thus indicating randomness. An identical distribution within the sequence means that the probability of occurrence is equal for each observation. This is captured within the following equation:

$$f(r_{j,t+1} | \Phi_t) = f(r_{j,t+1}) \qquad (6)$$

where f is the probability function. Equation (6) has more implications than the fair game model under (1), as it denotes that probability function f is the same for any random variable at any point in time. Despite these assumptions, Fama quickly observes that independence is not necessarily required in order for a market to be efficient. There are cases in which nonlinear dependence is present (suggesting it is possible to employ profitable strategies), and in which at the same time the serial covariance equals zero. This statement is crucial into understanding efficiency. It should be stressed that, by saying this, Fama implicitly points out that the assumption of linear independence (of which serial correlation, the linear relationship between a variable and its lagged values, is the indicator) should be fulfilled[5]. If there is linear dependence, both the random walk and the EMH are rejected. However, a market can still be efficient in the absence of nonlinear independence, as long as there is no serial correlation. In statistical terms, a general simplified random walk model is expressed as follows:

$$y_t = y_{t-1} + \varepsilon_t \quad (7)$$[6]

where ε_t is the error term, assumed to follow $\varepsilon_t \sim (0, \sigma^2)$. The randomness comes from the error term, which represents the part that cannot be predicted. The equation states that the best forecast of the current observation is equal to the preceding value, accounted for an estimation error. For every increase in time t, all that is left at the end is simply the increasing error term. Each next step within the sequence is a random one, thus suggesting it is futile to make any predictions.

---

[5] The random walk hypothesis has three different specifications. The most stringent form is the RW1, in which error term, ε_t, is i.i.d. with a mean of zero and variance σ^2. In the RW2 specification, the distribution need not be identical. The error term still needs to satisfy the assumptions of nonlinear independence and zero serial correlation (linear independence). The specification mentioned by Fama is the RW3, in which the assumption of nonlinear independence does not necessarily have to be satisfied. As long as there is no linear correlation between the lags, weak form efficiency is not violated.

[6] Equation (7) is a basic simplification of the random walk model. Under certain circumstances the model may include an intercept and a trend.

$$y_t = \alpha + \theta t + y_{t-1} + \varepsilon_t$$

where α is the intercept, and θ equals the time trend. For the sake of understanding the basics of the random walk theory, both α and θ are set to zero in equation (7).

Although the random walk hypothesis sounds straightforward, it might cause some confusion. A common misconception exists with regard to the distribution of future returns. Oftentimes, historical information is believed to be of no importance when it comes to determining the distribution of future returns. This is not what the random walk hypothesis implies. The random walk states that the order in which the historical returns came does not determine the distribution of future returns.

3.1.4. Weak Form Efficiency Tests

Fama (1970) starts off by clarifying that efficiency is most likely to occur in instances where there are no transaction costs, where all information is freely available to all market participants, and where all participants agree on the implications of the underlying information set. However, Fama also stresses that, while these assumptions are sufficient to indicate efficiency, they do not necessarily lead to efficiency. Moreover, violation of any of these three assumptions might be a source of inefficiency. Tests of weak form efficiency are typically categorized in two divisions.

The first division includes testing various trading strategies directly. Alexander (1961) provides an example by constructing a filter-based trading strategy. In this trading strategy, the market movement has to equal or exceed a certain percentage size for it to be considered for examination. All market movements below this percentage are filtered out. The filter strategy is set on the upward and downward trend of price movements, i.e., momentum. Let any arbitrary investor set the filter at 5%. If the price increases at least 5% from a previous low, the investor should go long by buying the stock. Whenever the price then decreases at least 5%, the investor should sell the stock and go short. If the filter is set at a stricter level, there are fewer losses to be made, but at the same time, profits are smaller as well. Alexander (1964) concludes that the filter strategy, in absence of transaction costs, outperforms the buy-and-hold strategy. This does not hold, however, when analysis takes transaction costs into account. These results are partially supported by Fama and Blume (1966). Their test results suggest that the filter strategy is inferior to the buy-and-hold strategy even in absence of transaction costs. Fama and Blume attribute this difference to a bias within Alexander's calculations, thus overestimating the actual profitability of the filter strategy.

The second division consists of statistical tests for randomness. This includes the typical tests examining serial correlation, dependence, and unit root within the return sequence. The amount of literature related to these tests is quite extensive. This research also concerns the statistical approach to finding a stochastic data generating process. Section 5 includes an elaborate description of the tests used in this research.

3.2. The Relationship Between Trading Volume and Efficiency

In spite of the observation of weak form efficiency, investors still seem to benefit from technical analysis. There appears to be a paradox between the general message of the EMH and reality.

3.2.1. On the Importance of Trading Volume Data

In terms of weak form market efficiency, researchers often only take historical prices into account. As mentioned by Fama (1970), differences in beliefs or in the interpretation of information might be sources of inefficiency. These discrepancies lead to inefficiency when some investors are capable of consistently making better estimations based on the information inherent to the price statistic. This indicates that the price statistic only captures a fragment of the underlying information set as assumed under weak form efficiency. If this holds, there is reason to believe that there are other measures that add value to technical analyses. One of these measures is trading volume data. Karpoff (1987) provides a synopsis of papers that deal with the relationship between trading volume and price changes. The majority of these papers find volume to be positively related to the change in price. According to Karpoff, research on the relationship between price and volume data is useful. It may shed a light on the structure of financial markets, as it tells something about the way investors process information. Under the assumption of rationality, investors make the best predictions given the circumstances. However, this does not imply that investors do not make errors, investors simply try to do the best they can with whatever information they have. Differences in the interpretation of information leads to the possibility of some investors outperforming others. Investors who are not as skilled, or are not as much in luck, need additional information to uncover the correct, efficient value.

3.2.2. The Evidence of Trading Volume Tests

Blume et al. (1994) find that the price statistic is not fully revealing. Their research shows that volume data contains information and tells something about the way market participants interpret this information. By including trading volume data into technical analysis, investors are capable of distinguishing between the information that is implicit in volume data to that implicit in the price statistic. Volume data increases the precision to which informational signals reach market participants. Antoniou et al. (1997) find similar results. For some investors it can be difficult to draw accurate inferences from the price statistic. In such events, volume data tells these uninformed, or less skilled, investors something about the value given by fully informed investors. The findings show that volume indeed contains valuable information. First, Antoniou et al. try to find weak form efficiency. Next, from their subset of weak form efficient markets, they take historical trading volume into account. Over half of the weak form efficient markets now display inefficiency. These results suggest that technical analysis, with the inclusion of volume data, is a worthwhile activity, as is often seen in reality.

3.3. Criticism on the Nature of the Efficient Market Hypothesis

The general statement of the EMH seems rather strict in its conclusion, especially looking at the strong form. The conclusion that society resides in an efficient state of optimal allocation at any point in time is not realistic. Furthermore, the specifications can be hard to verify (e.g., what piece of information is relevant, or how do investors interpret information).

One of the main criticisms on the EMH is with regard to the amount of time it takes for prices to fully reflect information. The implications of the EMH show that prices immediately adjust to their correct value as soon as new information is released. Critics argue that it is not realistic for prices to instantly adjust to their optimal state, and that it takes time for the market to incorporate new information. However, oftentimes, these critics assume that information comes in a definite, sure form. This is in itself not a realistic assumption either. In reality, information mostly comes gradually[7]. The EMH assumes every smaller fragment of information to be incorporated accordingly, whereas critics often think that the EMH assumes the entire piece of information to be incorporated instantly.

In addition, some critics do not support the main result of the EMH, i.e., that prices are correct. They conclude that this statement is rather restrictive, and that prices in real markets

[7] Consider the following example: airplane manufacturer Airbus plans on developing a new type of plane. This plane takes you from one place to the other within half the amount of time it would normally take. According to critics, the EMH states that the market instantly implements this information as a whole. On the contrary, the EMH does not treat this news as one entire piece of information. Rather, Airbus first announces their plans. Next, they continue developing the plane, and based on these developments, the project either succeeds or fails. The market incorporates these pieces of information – the announcement, development, and result – separately.

possibly are under- or overvalued. There are two ways of dealing with this view. On the one hand, the EMH simply states that prices are correct based on the available information set – the price is neither under- nor overvalued based on the relevant information. On the other hand, the less strict version of the EMH allows prices to deviate from their correct value. The cause for this deviation does not come from information, but possibly lies in behavioral factors, e.g., a bias towards certain companies, over-optimism, or other discrepancies (Malkiel, 2003).

Despite the complexity and criticism on the EMH, economists are still committed to the hypothesis. It has shown its usefulness throughout the past decades.

4. Literature Review

This section highlights some of the finance literature with respect to the cryptocurrency market.

4.1. General Weak Form Efficiency and Cryptocurrencies

One of the leading articles with respect to this research, is focused specifically on Bitcoin return predictability. Urquhart (2016) is the first to examine weak form efficiency on the Bitcoin market. In this paper the full sample consists of the daily closing price of Bitcoin, starting from August 01st 2010 to July 31st 2016. Other than just using the full sample, Urquhart creates two additional subsamples to test if the degree of Bitcoin efficiency varies throughout time. The first subsample covers the first three years of the full sample (August 01st 2010 to July 31st 2013), and the second subsample covers the final three years of the full sample (August 01st 2013 to July 31st 2016). In order to compute the returns, Urquhart takes the continuously compounded values of the price sequence:

$$R_t = Ln(P_t / P_{t-1}) * 100 \qquad (8)$$

where *t* and *t-1* denote *P*'s current value at time *t*, and *P*'s previous value at time *t-1*, respectively. Urquhart performs five tests for randomness to analyze weak form market efficiency. Firstly, Urquhart carries out the Ljung-Box test. The null hypothesis of this test implies that there is no autocorrelation. The Ljung-Box null hypothesis is rejected in the full sample and the first subsample, indicating that the returns are serial correlated. This means that the value of a specific return is affected by its previous values. Secondly, Urquhart executes the runs test and the Bartels test. The null hypothesis of both tests indicate independence. The runs test and the Bartels test are rejected in each of the samples, suggesting that there exists no randomness between the return observations. Thirdly, the automatic variance ratio (AVR) test checks if the return data follows a random walk. Again, for the full sample and the first subsample, the p-values are well-below the 5% significance level, thus pointing out that the successive values of the returns are not random. Furthermore, the Brock-Dechert-Scheinkman (BDS) test examines whether the data sequence is independent and identically distributed (i.i.d.). This test also rejects its null hypothesis, indicating that the return observations are correlated to their past values. Last but not least, the paper includes a rescaled range R/S Hurst exponent. The Hurst exponent tests for long memory, i.e., correlation within the return series. The test results show a pattern of negative correlation, or anti-persistence. The general conclusion of this research is that empirical results reject weak form efficiency. In the full sample, as well as in the first subsample, all randomness tests are rejected. Yet, for the second subsample, the Ljung-Box test and AVR test are not rejected. This may be an indication of Bitcoin transitioning into a more efficient market over time.

Chu and Nadarajah (2017) examine the exact same data set that is used in Urquhart's (2016) efficiency research. Similarly, there are three sample periods under consideration: the full sample, and two subsamples (each containing three years of data). They transform the data by multiplying the return series to the power of *m* (*m* is 17, but could be any odd integer). Chu and Nadarajah use the following tests of predictability: the Ljung-Box test; the runs test; the Bartels test; the wild-bootstrapped automatic variance ratio test; and the BDS test. Furthermore, they add three extra tests. The first of three is the spectral shape test, which searches for random walk. The second one is a portmanteau test for no serial correlation. The final test is a generalized spectral test on the martingale difference hypothesis[8]. None of the tests, with the exception of the runs test and Bartels test, show signs against their null hypothesis, i.e., the p-values are well-above the 5% significance level. According to these results, the return series are not serial correlated, are independent and identically distributed, and follow a random walk in all samples. This research strongly contradicts the main findings as reported by Urquhart (2016).

Bartos (2015) detects efficiency on the Bitcoin market by analyzing the effect of announcements involving Bitcoin. Bartos collects data from news articles and events, and creates a variable conveying positive signals and a variable conveying negative signals. These variables are regressed on the log transformation of the aggregate Bitcoin price index. The outcome reveals that the market immediately reacts on announcements, whether they are positive or negative. The Bitcoin price is significantly higher (lower) on days with positive (negative) announcements compared to days without any announcements. These results indicate that Bitcoin is in a state of efficiency.

As in the case of Urquhart (2016), Chu and Nadarajah (2017), and Bartos (2015), others also report contradictory findings on the predictability of Bitcoin returns. Bariviera (2017) searches for long memory within the return series. The R/S Hurst exponent exhibits long

[8] The martingale difference hypothesis states that the forecast of an observation is zero relative to its past value.

memory for the entire sample period. Use of the Detrended Fluctuation Analysis (DFA), however, exhibits long memory only in the period prior to 2014. After 2014, the return series tends towards following a memoryless, non-serial correlated process. Jiang, Nie, and Ruan (2018) discover persistence in Bitcoin returns. According to their tests, future returns are highly predictable, and the rate of predictability does not decrease over time. Whereas Brauneis and Mestel (2018) declare Bitcoin (and other cryptocurrencies) to follow the EMH, Caporale, Gil-Alana, and Plastun (2018) find Bitcoin to exhibit inefficiency. Although quite extensively examined, academics have yet to find a conclusive answer on weak form efficiency on the cryptocurrency market.

4.2. The Volume-Price Relationship on the Cryptocurrency Market

Little research, however, has been done on the volume-price relationship of cryptocurrencies. Balcilar et al. (2017) test whether there exists a causal relationship between trading volume and returns of Bitcoin. Papers on trading volume typically focus on finding a volume-price relationship within the conditional mean of the return distribution. Balcilar et al. aim to find if such a relationship exists, taking the entire distribution into account. They test this by proportioning the distribution into quantiles. This enables them to search for predictability within the mean, as well as in the tails of the distribution. The paper first looks for standard linear Granger causality, in which the null hypothesis states that volume does not Granger cause returns. The test statistic indicates that the null hypothesis cannot be rejected. Volume data does not predict returns. Nonetheless, it may be that the Granger condition of linearity is specified erroneously. Balcilar et al. use the BDS test to examine whether the relationship between volume and price is in fact linear. The null hypothesis of i.i.d. residuals is rejected, indicating there exists nonlinearity within the volume-price relationship. The use of linear Granger causality is incorrect, as the relationship displays nonlinearity. In order to deal with nonlinearity and the fat tails of the data, the paper searches for Granger causality within the quantiles of the distribution. The null hypothesis is rejected within the range of 0.24 to 0.66, which means that volume can be used to predict returns. In the quantiles below and above this range, however, the null hypothesis cannot be rejected. The overall conclusion of this study shows that volume holds predictive power only when the market is in its normal, median mode. In extreme market regimes – a bearish (lower quantiles) or bullish (upper quantiles) market – trading volume data holds no predictive power at all. In these instances, volume-based trading rules are not profitable.

4.3. Anomalies on the Cryptocurrency Market

A subject closely linked to efficiency, or rather deviations from efficiency, is the observance of anomalies. An anomaly indicates that there is a pattern in the return sequence. This contradicts the efficiency theory, as the EMH assumes the return sequence to be free of any observable patterns. The detection of market inefficiency oftentimes serves as a basis to look into the possible causes of the deviations, i.e., if any type of observable anomaly exists. Although finding anomalies is beyond the scope of this analysis, some comments on the existence of anomalies may be of use for future research.

Bartos argues (2015) that most of the well-known anomalies are non-existent within the cryptocurrency market due to the market's dynamics. Bartos claims that periodically reoccurring anomalies, such as the Weekend Effect, are not supposed to occur, since cryptocurrencies are traded seven days a week, 24 hours a day. Moreover, many anomalies relate to the fundamentals of companies, which are of no concern in the case of cryptocurrencies[9]. This statement, however, is deduced by Bartos' own assumptions rather than research. Despite his arguments, there are papers that report findings in favor of the existence of anomalies within the cryptocurrency market.

Kurihara and Fukushima (2017) find evidence in favor of the Day-of-the-Week Effect. They perform an Ordinary Least Squares regression, as well as a Robust Least Squares regression on the daily Bitcoin return sequence, with dummy variable *D* as a proxy for the days of the week. In addition, they perform the Ljung-Box test. The Ljung-Box test is rejected, and the regressions show that Bitcoin is significantly inefficient. The test results reveal that the anomaly seems to occur over the weekend, as the effect is smaller than during weekdays.

Similar to Kurihara and Fukushima (2017), Caporale and Plastun (2017) present results in line with the Day-of-the-Week effect. The working paper executes several tests for the detection of anomalies across four cryptocurrencies. The research proposes to apply a trading simulation, to see whether any profits can be made from the anomaly. Out of the four currencies, only Bitcoin showcases an anomaly. Bitcoin returns are higher on Monday, and Caporale and Plastun use the simulator to go long on Monday and close this position at the

[9] Cryptocurrencies are not backed by any performance measures (e.g., financial statements, revenues, and the likes), as they generally are not the result of a company's effort to generate profits.

end of the day. Overall, the trading strategy seems to be profitable. Nonetheless, when the full sample is divided into individual years, this strategy does not seem to have returns very different to that of applying a random trading strategy.

The existence of anomalies is oftentimes hard to pinpoint, and the reason for the occurrence of anomalies is even more difficult to examine. Yet, more research into this topic could possibly be beneficial into understanding the dynamics of cryptocurrencies.

5. Data and Methodology

5.1. Data Description

This research examines the degree of return predictability of the following cryptocurrencies: Bitcoin (BTC), Ethereum (ETH), Ripple (XRP), Litecoin (LTC), and Monero (XMR). The data set consists of price and trading volume observations at a daily level, as of September 01st 2015 to August 31st 2018. The data is retrieved from cryptocompare.com. The price data is the volume-weighted average price, taken from several leading cryptocurrency exchanges. The volume data takes on the aggregate value from these exchanges.

In addition to the full sample period, this research generates two subsample periods to check whether the degree of efficiency varies over time. Each subsample period contains 18 months' worth of observations: the first subsample ranges from September 01st 2015 to February 28th 2017, and the second subsample ranges from March 01st 2017 to August 31st 2018. The choice of two additional similar-sized subsamples is based on the assurance that there would be enough observations left in each of the subsamples.

The raw price data carries properties that are common within price series (e.g., trends, heavily correlated residuals, and the inability to make comparisons), making it a less favorable series to work with statistically. Hence, a continuously compounded return series is created by taking the natural logarithm, which is characterized by the following equation:

$$R_t = Ln(P_t / P_{t-1}) * 100 \qquad (9)$$

5.2. Methodology

5.2.1. Statistical Random Walk Tests

Serial Correlation

The very first statistical test searches for possible serial correlation (*autocorrelation*) within the series. In the event of serial correlation, observations within the series are affected by their historical values. Serial correlation indicates linear dependence and is a violation of the random walk model. The Ljung-Box test, which is an improved version of the Box-Pierce test, detects whether or not serial correlation is present (Box & Ljung, 1978). Instead of testing serial correlation at individual lags, the Ljung-Box test looks for joint serial correlation up to any m lags[10]. The Ljung-Box test statistic (Q) is computed as follows:

$$Q = n(n+2) \sum_{k=1}^{m} \frac{r^2(k)}{n-k} \quad (10)$$

n equals the sample size, m is the number of lags tested, and r is the autocorrelation at lag k. The Ljung-Box test follows a χ^2 distribution, and the null hypothesis declares the correlation up to a set of lag m (here 12) to be equal to zero. Failure to reject the null hypothesis is in line with random walk. It means that the degree of serial correlation does not significantly differ from zero and that the observations are linearly independent of one another. If this is the case, investors cannot predict returns by looking at historical observations.

Unit Root

Generally, hypothesis testing (e.g., model forecasting) requires the data generating process to be at least weakly stationary, i.e., the mean, variance, and covariance of the distribution does not change throughout time. This means it is not difficult to predict future values, which is inconsistent with weak form efficiency. This research first applies the Augmented Dickey-Fuller (ADF) test to check whether or not the return series are stationary (Dickey & Said, 1984). The ADF test is a revised and improved test that allows for higher order integration. The null hypothesis states that the series contain unit root[11], or equivalently, is non-stationary. The data generating processes are consistent with random walk if the test results fail to reject the null hypothesis. The graphs in Appendix B show that none of the log return series contain a trend. However, it may be an appropriate measure to include an intercept, as the mean of the currencies do not exactly evolve around zero. The ADF test fits the following first-difference regression to test for unit root:

$$\Delta y_t = \alpha + \delta y_{t-1} + \varepsilon_t \quad (11)$$

where intercept α may or may not be included, depending on the structure of the series, and where $\delta = 0$ is equivalent to the null hypothesis of unit root. Instead of the regular p-value, the ADF test results rely on the MacKinnon p-value. Prior to MacKinnon's (1996) paper, calculating the significance level was a complicated and tedious task when it involved unit root tests. A possible flaw of the ADF test is that the right-hand side of the test depends on the number of lags included. An increase in the number of lags diminishes the number of

[10] Instead of looking at the null hypotheses of H_0: $\rho_1 = 0$, H_0: $\rho_2 = 0$, up to H_0: $\rho_m = 0$, individually, the Ljung-Box test looks at H_0: $\rho_1 = \rho_2 = \ldots = \rho_m = 0$.

[11] See Appendix A for the mathematics behind the unit root assumption of the random walk model.

observations that are left. In this research, the optimal number of lags is chosen under the Schwarz Information Criterion.

Unfortunately, the ADF test is not very powerful, especially if the root is close to unity. For example: for $\varphi = 0.95$, the ADF test may not be able to perfectly distinguish between stationarity and non-stationarity (especially when the size of the sample is small). It is biased towards non-stationarity, while the series in fact is stationary. To account for possible biases, this research performs a second unit root test, the Phillips-Perron test (Perron & Phillips, 1988). The PP null hypothesis is similar to that of the ADF test: the time series contains unit root. One of the advantages of testing the PP test over the ADF test is that the PP test is nonparametric, which means that the test does not necessarily depend on a certain type of distribution. A second advantage is that the test is not only robust against serial correlation – like the ADF test, but that it also appropriately takes care of heteroskedasticity. A final advantage is the fact that the test does not depend on choosing any lags. The ADF and PP test often have similar outcomes, however, the test results will slightly differ due to the difference in sample size (the PP test does not include any lags in its calculations).

Runs Test

The next test in this analysis is the runs test for randomness (Wald & Wolfowitz, 1940). The runs test is a nonparametric test. However, with a large enough sample size, the test statistic more or less follows a normal distribution, with the following mean and variance:

$$E(R) = \frac{2N_+N_-}{N_+ + N_-} + 1 \qquad (12)$$

$$\sigma^2(R) = \frac{2N_+N_-(2N_+N_- - N_+ - N_-)}{(N_+ + N_-)^2(N_+N_- - 1)} \qquad (13)$$

The runs test computes the amount of runs[12] in the entire sample and checks whether the total amount of runs is in line with the expected amount of runs under randomness. The null hypothesis deals with whether the order of the observations is random. If the total amount of runs is far below or beyond what would be expected, then the underlying process is assumed to not be random. In this analysis, an event within the sample is given the symbol '+' ('–') if the return is above (below) the mean. In case an observation equals the threshold (here: mean), it will be discarded. A run consists of a sequence of the same sign up to the point

[12] A run is a sequence of successive observations which are assigned the same sign or symbol.

where it changes to the other sign[13]. Rejection of the null hypothesis is an indication that the sample does not display randomness. This subsequently comes down to a rejection of the random walk model.

Brock-Dechert-Scheinkman Test

Section 3 of this research briefly highlighted the existence of multiple types of dependence (linear serial correlation versus nonlinear dependence). The paper written by Brock, Dechert, LeBaron, and Scheinkman (1996) shows a nonparametric method on how to detect nonlinearity. The BDS null hypothesis declares the sequence of observations to be drawn from an i.i.d. data generating process. The BDS test is a portmanteau test, as its alternative hypothesis is unspecified[14]. The test statistic depends on the correlation integral ($C_{\varepsilon,m}$). The correlation integral is a measure of the probability that any pair of observations are within a certain distance, ε, of each other.

$$C_{\varepsilon,m} = \frac{1}{Nm(Nm-1)} \sum_{i \neq j} I_{i,j;\varepsilon} \quad (14)$$

where $I_{i,j:\varepsilon}$ equals one if the distance between object i and j is smaller than or equal to ε (i.e., they are near one another). Variable I takes on the value of zero in all other situations. The power of the test depends on the choice of distance ε, and embedding dimension m. There are no exact guidelines to follow when choosing these two variables. Nonetheless, Brock, Hsieh, and LeBaron (1991) show that the power of the test statistic is greatest when setting distance ε as a function of the standard deviation, in between the range of 0.50σ (0.25σ for extremely large sample sizes) and 2.00σ. Moreover, one should carefully consider the level of embedding dimension m, as an increase in m could negatively affect the test results[15]. In this analysis distance ε is tested at multiple levels of σ (0.50σ; 0.75σ; 1.00σ; 1.25σ; and 1.50σ). Embedding dimension m ranges from two to five. Under the assumption of i.i.d., the probability of a distance being smaller than or equal to ε is constant. The joint probability ($C_m(\varepsilon)$) of the pairs being within distance ε of each other equals the product of all individual probabilities ($C_1^m(\varepsilon)$).

---

[13] A sequence of '+ + – + – – + – + + – – –' is made up of 13 events, and 8 runs. Run 1 consists of + +; run 2 is simply one – symbol; run 3 exists of a single +, and so on.

[14] The BDS test results do not identify what type of nonlinear dependence exists. It only detects whether or not the data sequence has been formed by some type of nonlinearity. This makes the BDS test useful in multiple settings. Additional diagnostic tests are required in order to specifically tell what type of nonlinearity there is. However, this is not of much importance within this research. The only thing that bears relevance is whether or not the series are independent.

[15] If one were to choose m too large relative to the finite size of the sample, it would not produce the most reliable results. The finite sample properties (e.g., unbiasedness) worsen as m increases.

$$C_m(\varepsilon) = C_1^m(\varepsilon) \quad (15)$$

Rejection of the BDS test implies the existence of dependence, thus rejecting the random walk model. It should be noted that this test is a complementary test, as nonlinearity does not necessarily imply inefficiency.

Autoregressive Process of Order 5

Last but not least, the research tests whether or not lagged returns affect current returns based on the execution of simple Ordinary Least Squares regressions. Autoregressive (AR) models of order 5 are constructed to see whether the 1- to 5-day lagged values have any significant impact on current returns. The choice of order five is made based on the following principles. The foremost reason for including more than one lag is purely because it may be interesting to see whether current returns do not only depend on yesterday's value, but also whether they depend on earlier observations. However, with the inclusion of additional lags as potential significant regressors, the model loses observations. This can negatively affect the power of the test. Therefore, the limit is set at five lags as to not include too many independent variables. The AR(5) model is constructed as follows:

$$R_{j,t} = \beta_1 c + \beta_2 R_{j,t-1} + \beta_3 R_{j,t-2} + \beta_4 R_{j,t-3} + \beta_5 R_{j,t-4} + \beta_6 R_{j,t-5} + \varepsilon_t \quad (16)$$

As autoregressive models of order 2 and up by definition include more than one regressor, other than intercept *c*, it is appropriate to conduct an F-test. In this analysis, instead of looking at each lag individually (t-test), the F-test checks whether coefficients β_2 to β_6 jointly differ from zero. The null hypothesis is stated as follows:

$$H_0: \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = 0$$

The null hypothesis tests whether the coefficients of the lagged returns are jointly significantly different from zero[16]. If the test results fail to reject the null hypothesis, the coefficients are said to be insignificant, i.e., the lagged return has no effect on current return. Failure to reject the null hypothesis is in line with the random walk.

5.2.2. Testing the Relationship between Volume and Price

The second null hypothesis of this research focuses on the relationship between trading volume and cryptocurrency returns. This part of the analysis introduces trading volume data

[16] The alternative hypothesis involves 'or'-statements rather than the 'and'-statements as under the null hypothesis. The alternative hypothesis is concerned whether at least one of the coefficients is not equal to zero.

$$H_a: \beta_2 \neq 0, \text{ or } \beta_3 \neq 0, \text{ or } \beta_4 \neq 0, \text{ or } \beta_5 \neq 0, \text{ or } \beta_6 \neq 0$$

and tests whether or not it contains information beyond that what is implicit in the price statistic. The analysis carries out Granger-causality tests. Granger (1969) developed a statistical test in which the lagged values of an independent variable are included as a regressor. One has to be cautious when testing Granger-causality as the name of the test might be misleading. The test does not capture the true causality between variables, rather, it describes the degree of predictability of one variable on another. The Granger-causality test is based on testing if the lagged observations of some independent variable is significant in predicting some other dependent variable beyond the predictability of the dependent variable's own lagged values. For simplicity, let X_t and Y_t be the variables under consideration. Variable X_t is said to Granger-cause Y_t if X_{t-i} improves the prediction of Y_t after the inclusion of lagged values of Y_t. In this research, the main objective is to find out if the inclusion of trading volume data (X_t) affects the predictability of cryptocurrency returns (Y_t). The null hypothesis states that some variable X_t does not Granger-cause Y_t. In case the results cannot reject this null hypothesis, there is reason to assume that X_t holds no predictive power beyond Y_t's own lagged values. In a bivariate model, Granger-causality is either unidirectional (X_t Granger-causes Y_t, or Y_t Granger-causes X_t), bidirectional (X_t Granger-causes Y_t and Y_t Granger-causes X_t simultaneously), or non-existent (independence between X_t and Y_t).

In order to test the possible existence of Granger-causality, the volume series ought to be stationary. Non-stationarity could lead to incorrect results. Therefore, first of all, the ADF test is conducted to test whether or not the series contain unit root. In case the original series is non-stationary, it must be differenced until it is stationary. The next step is to find the corresponding amount of lags to include with respect to the return and trading volume data. The relationship of interest is the following:

$$R_{j,t} = \alpha_0 + \beta_1 R_{j,t-1} + \ldots + \beta_5 R_{j,t-5} + \chi_1 V_{j,t-1} + \ldots + \chi_5 V_{j,t-5} + \varepsilon_t \quad (17)$$

The test results report an F-statistic as to conclude whether or not the combined effect of R_{lag} and V_{lag} is significant for the prediction of returns.

6. Results

6.1. Descriptive Statistics

Tables 1.a through 1.c report the descriptive statistics of the log return series over the full sample period, and the subsample periods, respectively. The tables show that all cryptocurrencies have a high standard deviation from their mean values. The amount of the minimum and maximum returns is quite substantial, especially for Ripple. Bitcoin displays the smallest minimum and maximum returns. Intuitively, this is reasonable, considering the fact that Bitcoin has been the most popular among all cryptocurrencies. At the start of the sample period, Bitcoin had already matured more in terms of price levels compared to the other cryptocurrencies. Other than that, the subsamples show that there are some fluctuations in the degree of skewness, either changing from approximate symmetry to moderate or heavy skewness, or vice versa. All five currencies share the property of excess kurtosis in each of the samples, indicating high probabilities of extreme values. Overall, the combined result of skewness and kurtosis indicate that none of the currencies are normally distributed. Despite these statistics, nothing seems extremely out of the ordinary compared to return series in other finance literature.

Table 1.a. Descriptive statistics of the ln return over the full sample period (09/01/2015 – 08/31/2018)

Currency	N	Mean	Median	SD	Min	Max	Skewness	Kurtosis
BTC	1,095	0.313	0.324	4.044	-18.917	22.762	-0.198	7.131
ETH	1,095	0.487	0.000	7.210	-31.010	38.299	0.158	6.555
XRP	1,095	0.325	-0.191	8.933	-65.299	102.800	2.099	27.089
LTC	1,095	0.283	0.000	5.838	-31.248	55.165	1.788	17.973
XMR	1,095	0.499	0.000	8.500	-51.083	64.087	0.786	11.469

Table 1.b. Descriptive statistics of the ln return over the first subsample period (09/01/2015 – 02/28/2017)

Currency	N	Mean	Median	SD	Min	Max	Skewness	Kurtosis
BTC	546	0.303	0.245	2.900	-16.742	11.928	-0.965	11.245
ETH	546	0.452	-0.080	7.715	-31.010	38.299	0.018	7.170
XRP	546	-0.098	-0.194	7.024	-45.179	40.689	0.792	13.302
LTC	546	0.055	0.000	3.095	-22.896	21.595	-0.081	16.834
XMR	546	0.590	0.000	9.459	-51.083	64.087	0.945	12.526

Table 1.c. Descriptive statistics of the ln return over the second subsample period (03/01/2017 – 08/31/2018)

Currency	N	Mean	Median	SD	Min	Max	Skewness	Kurtosis
BTC	549	0.324	0.424	4.929	-18.917	22.762	-0.027	5.119
ETH	549	0.522	0.137	6.678	-26.868	26.024	0.374	5.094
XRP	549	0.746	-0.189	10.483	-65.299	102.800	2.291	25.534
LTC	549	0.510	-0.139	7.642	-31.248	55.165	1.523	11.590
XMR	549	0.409	-0.122	7.432	-27.646	42.764	0.387	6.449

6.2. Random Walk Tests

Serial Correlation Test

Tables 2.a through 2.e reveal the Ljung-Box test results per cryptocurrency, over all sample periods. The tables report the autocorrelation levels up to 12 lags[17]. In the case of Bitcoin, none of the samples rejects the null hypothesis (with the minor exception of the first subsample at lags 8 and 9). The Q-statistics are smaller than the χ^2 critical value, and the p-values are well above the 5% significance level. Returns on the Bitcoin market seem to be in line with the random walk hypothesis. As evident from the subsamples, the inability to reject the null hypothesis grows stronger over time (the results for some of the lags are barely over the 5% significance level in the first subsample, but are larger in the second subsample), which might imply that the accuracy of not rejecting the 'no autocorrelation' statement increases. Ethereum and Litecoin on the other hand, both seem to reject the Ljung-Box null hypothesis over the full sample and the first subsample. Looking at the second subsample, both currencies fail to reject the null hypothesis. Ethereum and Litecoin seem to transfer into becoming random over time. Moreover, Ripple and Monero undoubtedly reject the null of no serial correlation over all sample periods – nearly all p-values are below 5% or even 1%, regardless of how great or small the set of lags is. According to the serial correlation test, investors generally are capable of generating excess return by applying technical analysis. This statement does not include Bitcoin, however, as the Bitcoin market seems to support the random walk hypothesis.

[17] Increasing the amount of lags within the set mostly generated the same results – the pattern of (no) significant autocorrelation remained with the inclusion of more than 12 lags.

Table. 2.a. Autocorrelation (AC) coefficient and Ljung-Box Q-statistic (*Bitcoin*).

	Full Sample			First Subsample			Second Subsample		
Lag	AC	Q-Stat	P-value	AC	Q-Stat	P-value	AC	Q-Stat	P-value
1	-0.002	0.005	0.947	-0.006	0.017	0.895	-0.001	0.000	0.986
2	0.012	0.158	0.924	-0.090	4.450	0.108	0.047	1.202	0.548
3	0.005	0.184	0.980	-0.027	4.862	0.182	0.016	1.336	0.721
4	-0.035	1.529	0.821	0.088	9.134	0.058	-0.077	4.649	0.325
5	0.044	3.674	0.597	0.020	9.350	0.096	0.053	6.204	0.287
6	0.031	4.738	0.578	0.046	10.520	0.104	0.026	6.573	0.362
7	-0.020	5.165	0.640	-0.032	11.102	0.134	-0.016	6.725	0.458
8	0.003	5.174	0.739	-0.104	17.160	0.028	0.039	7.565	0.477
9	-0.004	5.191	0.817	-0.028	17.595	0.040	0.005	7.580	0.577
10	0.042	7.132	0.713	-0.023	17.887	0.057	0.063	9.804	0.458
11	0.011	7.272	0.777	0.021	18.126	0.079	0.007	9.831	0.546
12	0.002	7.276	0.839	-0.027	18.520	0.101	0.014	9.948	0.621

Table. 2.b. Autocorrelation (AC) coefficient and Ljung-Box Q-statistic (*Ethereum*).

	Full Sample			First Subsample			Second Subsample		
Lag	AC	Q-Stat	P-value	AC	Q-Stat	P-value	AC	Q-Stat	P-value
1	-0.055	3.285	0.070	-0.120	7.906	0.005	0.031	0.531	0.466
2	0.010	3.399	0.183	-0.004	7.915	0.019	0.026	0.905	0.636
3	0.088	11.905	0.008	0.116	15.312	0.002	0.046	2.100	0.552
4	-0.000	11.905	0.018	-0.021	15.548	0.004	0.024	2.412	0.661
5	-0.002	11.911	0.036	-0.061	17.630	0.003	0.076	5.603	0.347
6	0.007	11.965	0.063	-0.045	18.749	0.005	0.075	8.767	0.187
7	0.051	14.806	0.039	0.096	23.832	0.001	-0.010	8.819	0.266
8	-0.044	16.976	0.030	-0.073	26.789	0.001	-0.005	8.834	0.357
9	0.016	17.244	0.045	0.019	26.985	0.001	0.016	8.980	0.439
10	0.013	17.440	0.065	0.009	27.030	0.003	0.022	9.246	0.509
11	0.057	21.054	0.033	0.060	29.019	0.002	0.051	10.704	0.468
12	0.023	21.656	0.042	0.023	29.314	0.004	0.018	10.881	0.539

Table. 2.c. Autocorrelation (AC) coefficient and Ljung-Box Q-statistic (*Ripple*).

	Full Sample			First Subsample			Second Subsample		
Lag	AC	Q-Stat	P-value	AC	Q-Stat	P-value	AC	Q-Stat	P-value
1	-0.164	29.554	0.000	-0.383	80.586	0.000	-0.070	2.7051	0.100
2	0.094	39.327	0.000	-0.020	80.813	0.000	0.142	13.925	0.001
3	0.068	44.370	0.000	0.023	81.115	0.000	0.085	17.941	0.000
4	-0.015	44.619	0.000	0.031	81.662	0.000	-0.039	18.777	0.001
5	0.051	47.533	0.000	-0.007	81.687	0.000	0.074	21.811	0.001
6	0.013	47.728	0.000	-0.050	83.085	0.000	0.039	22.677	0.001
7	0.028	48.576	0.000	0.018	83.260	0.000	0.030	23.173	0.002
8	0.057	52.143	0.000	0.028	83.684	0.000	0.067	25.692	0.001
9	0.039	53.803	0.000	-0.031	84.225	0.000	0.067	28.212	0.001
10	0.056	57.220	0.000	0.037	84.996	0.000	0.061	30.316	0.001
11	-0.036	58.676	0.000	-0.042	85.978	0.000	-0.038	31.130	0.001
12	-0.001	58.677	0.000	0.000	85.978	0.000	-0.004	31.138	0.002

Table. 2.d. Autocorrelation (AC) coefficient and Ljung-Box Q-statistic (*Litecoin*).

	Full Sample			First Subsample			Second Subsample		
Lag	AC	Q-Stat	P-value	AC	Q-Stat	P-value	AC	Q-Stat	P-value
1	0.011	0.133	0.715	0.028	0.423	0.516	0.007	0.024	0.877
2	-0.015	0.371	0.831	-0.118	8.120	0.017	0.001	0.024	0.988
3	0.022	0.883	0.829	-0.116	15.562	0.001	0.042	1.025	0.795
4	0.034	2.150	0.708	0.034	16.194	0.003	0.032	1.605	0.808
5	0.010	2.266	0.811	0.019	16.390	0.006	0.007	1.634	0.897
6	0.112	16.092	0.013	-0.013	16.485	0.011	0.131	11.153	0.084
7	-0.032	17.205	0.016	-0.032	17.039	0.017	-0.034	11.791	0.108
8	-0.052	20.249	0.009	-0.083	20.838	0.008	-0.050	13.167	0.106
9	-0.002	20.255	0.016	-0.019	21.044	0.012	-0.001	13.168	0.155
10	-0.007	20.304	0.027	-0.031	21.597	0.017	-0.005	13.181	0.214
11	0.055	23.632	0.014	0.016	21.742	0.026	0.060	15.196	0.174
12	-0.032	24.762	0.016	-0.037	22.507	0.032	-0.032	15.789	0.201

Table. 2.e. Autocorrelation (AC) coefficient and Ljung-Box Q-statistic (*Monero*).

	Full Sample			First Subsample			Second Subsample		
Lag	AC	Q-Stat	P-value	AC	Q-Stat	P-value	AC	Q-Stat	P-value
1	-0.167	30.775	0.000	-0.207	23.479	0.000	-0.104	6.009	0.014
2	0.000	30.775	0.000	-0.012	23.551	0.000	0.018	6.190	0.045
3	0.015	31.038	0.000	-0.006	23.574	0.000	0.050	7.561	0.056
4	-0.018	31.382	0.000	-0.008	23.607	0.000	-0.034	8.188	0.085
5	0.015	31.617	0.000	0.007	23.632	0.000	0.027	8.587	0.127
6	0.156	58.318	0.000	0.173	40.299	0.000	0.128	17.691	0.007
7	-0.025	59.036	0.000	-0.023	40.592	0.000	-0.028	18.127	0.011
8	0.009	59.133	0.000	0.024	40.908	0.000	-0.014	18.235	0.020
9	-0.020	59.577	0.000	-0.013	41.009	0.000	-0.030	18.725	0.028
10	-0.074	65.616	0.000	-0.095	46.006	0.000	-0.040	19.620	0.033
11	0.026	66.341	0.000	0.023	46.313	0.000	0.025	19.976	0.046
12	0.019	66.744	0.000	0.030	46.829	0.000	-0.005	19.989	0.067

Unit Root Test

Table 3 displays the results of the Augmented Dickey-Fuller test. The return series of all cryptocurrencies reject the null hypothesis at the 1% significance level – the results imply that all cryptocurrencies are highly consistent with stationarity. The ADF test did not detect unit root within any of the return series. Thus, the level of the root seems to be smaller than 1. This indicates that as time elapses, the mean and variance of the observations are becoming more constant. This is inconsistent with the random walk model, as it means that the characteristics of the sequence are predictable. Future returns are thus foreseeable, and investors are capable of beating the market by using information that is implicit in the price statistic.

Table. 3. Results of the Augmented Dickey-Fuller test for all sample periods.

Currency	Full Sample	First Subsample	Second Subsample
BTC	-33.112***^{a}	-23.434***^{a}	-23.313***^{b}
	(0.000)	(0.000)	(0.000)
ETH	-35.006***^{a}	-26.330***^{b}	-22.579***^{b}
	(0.000)	(0.000)	(0.000)
XRP	-17.310***^{b}	-23.584***^{b}	-11.168***^{b}
	(0.000)	(0.000)	(0.000)
LTC	-32.619***^{b}	-18.380***^{b}	-23.130***^{b}
	(0.000)	(0.000)	(0.000)
XMR	-39.108***^{a}	-28.628***^{b}	-25.838***^{b}
	(0.000)	(0.000)	(0.000)

The table reports the Augmented Dickey-Fuller test statistic. The corresponding critical values are -1.61 (10%), -1.95 (5%), and -2.58 (1%) for models without intercept and trend. The critical values for models that include an intercept are -2.57 (10%), -2.86 (5%), and -3.43 (1%). The MacKinnon p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively. Superscripts *a* and *b* refer to models with an intercept (but without trend), and models without intercept and trend, respectively.

Table 4 includes the results of the Phillips-Perron test. According to these results, none of the cryptocurrencies follow the random walk hypothesis, as the null of unit root is highly rejected in each of the samples. The market's characteristics are therefore predictable, and investors are capable of systematically earning excess return based on the price statistic.

Table. 4. Results of the Phillips-Perron test for all sample periods.

Currency	Full Sample	First Subsample	Second Subsample
BTC	-33.116***^{a}	-23.468***^{a}	-23.325***^{b}
	(0.000)	(0.000)	(0.000)
ETH	-34.950***^{a}	-26.330***^{b}	-22.848***^{b}
	(0.000)	(0.000)	(0.000)
XRP	-38.533***^{b}	-39.088***^{b}	-25.302***^{b}
	(0.000)	(0.000)	(0.000)
LTC	-32.734***^{b}	-23.016***^{b}	-23.193***^{b}
	(0.000)	(0.000)	(0.000)
XMR	-38.707***^{a}	-28.856***^{b}	-25.766***^{b}
	(0.000)	(0.000)	(0.000)

The table reports the Phillips-Perron test statistic. The corresponding critical values are -1.62 (10%), -1.94 (5%), and -2.57 (1%) for models without intercept and trend. The critical values for models that include an intercept are -2.57 (10%), -2.87 (5%), and -3.44 (1%). The MacKinnon p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively. Superscripts *a* and *b* refer to models with an intercept (but without trend), and models without intercept and trend, respectively.

Runs Test

Table 5 shows the output of the Wald-Wolfowitz runs test. Overall, the total amount of runs for each of the currencies lie reasonably close to one another. A possible explanation for this may be that Bitcoin more or less is the leading factor in the entire cryptocurrency market. Throughout the market's history, other currencies seem to follow the same course as Bitcoin. The test results for Bitcoin and Ethereum are similar in the sense that the test cannot reject the null hypothesis in any of the sample periods. The returns on Bitcoin and Ethereum seem to support the random walk theory. Based on these results, investors should not earn any excess return by employing technical analysis. A somewhat remarkable observation, however, but when one looks closely it is clear that the p-values for Ethereum are larger than the p-values for Bitcoin. This may seem odd at first, as Bitcoin is often (perhaps erroneously) viewed as the most efficient cryptocurrency out of all. The observation that the certainty of not rejecting the null for Ethereum overthrows the level of certainty for Bitcoin might be somewhat counterintuitive. Ripple and Litecoin both reject the null hypothesis in the full sample period, as well as in the first subsample period. Nevertheless, the null hypothesis cannot be rejected in the second subsample. Both currencies seem to transition into becoming random over time. Nonetheless, this effect seems negligible when taken as a whole, as the full sample still undoubtedly rejects the null hypothesis. As in the case of Ripple and Litecoin, Monero also rejects the null hypothesis in the full sample period. However, unlike these two cryptocurrencies, Monero seems to be doing worse in terms of randomness. In the first subsample the test cannot reject its null hypothesis, whereas in the second subsample the opposite is the case. Based on the runs test, no general statement can be made on the degree of randomness on the cryptocurrency market. Some currencies are in line with the random walk theory, whereas others are not.

Table. 5. Results of the Wald-Wolfowitz runs test for all sample periods.

Currency	Full Sample	First Subsample	Second Subsample
BTC	561	292	269
	(0.449)	(0.115)	(0.612)
ETH	545	274	274
	(0.866)	(0.760)	(0.900)
XRP	597***	318***	283
	(0.000)	(0.000)	(0.218)
LTC	576**	313***	285
	(0.020)	(0.000)	(0.322)
XMR	586***	284	300**
	(0.009)	(0.157)	(0.029)

The table reports the results of the Wald-Wolfowitz runs test for randomness. The table presents the amount of total runs within the entire time series sequence. The p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively.

Brock-Dechert-Scheinkman Test

The BDS test is conducted at multiple levels of distance ε, ranging from 0.50σ to 1.50σ with increments of 0.25σ. The amount of embedding dimension m is set at 4 (2-5). For every cryptocurrency, the BDS test is rejected at each and every level. The test results are highly significant. This means that the data generating process seems extremely inconsistent with the random walk theory. The return series are not drawn from an i.i.d. process. There is reason to believe that there exists some type of nonlinear dependence within each of the return series[18]. The results share similarities within all distances. According to these results, investors would be making profits when trading on the price statistic.

[18] Additional diagnostic tests can clarify the type of nonlinearity that exists within the series. However, in this research it is irrelevant to execute additional tests on the type of nonlinearity. The purpose of this analysis is to find out whether or not dependence exists, not to find out exactly what form of dependence this constitutes.

Table 6.a. Brock-Dechert-Scheinkman test results (Bitcoin).

	ε				
m	0.50	0.75	1.00	1.25	1.50
Full Sample					
2	9.382***	8.898***	8.078***	7.563***	6.858***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	13.323***	12.278***	11.100***	10.123***	8.803***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	18.126***	15.456***	13.204***	11.726***	10.118***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	25.555***	19.616***	15.520***	13.227***	11.281***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
First Subsample					
2	4.952***	5.140***	5.491***	5.341***	5.354***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	5.765***	5.578***	5.904***	5.803***	5.798***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	6.866***	6.405***	6.565***	6.266***	5.913***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	8.488***	7.358***	7.299***	6.717***	6.062***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Second Subsample					
2	2.176**	2.450**	2.473**	2.387**	2.454**
	(0.030)	(0.014)	(0.013)	(0.017)	(0.014)
3	3.552***	4.025***	3.959***	3.569***	3.454***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
4	4.702***	5.286***	5.098***	4.566***	4.494***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	5.319***	6.247***	5.988***	5.400***	5.224***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

The BDS test statistic is presented in the table above (under an approximate normal distribution). The p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively.

Table 6.b. Brock-Dechert-Scheinkman test results (Ethereum).

	ε				
m	0.5	0.75	1.00	1.25	1.50
Full Sample					
2	9.660*** (0.000)	10.156*** (0.000)	10.539*** (0.000)	10.286*** (0.000)	10.081*** (0.000)
3	11.645*** (0.000)	12.106*** (0.000)	12.375*** (0.000)	12.033*** (0.000)	11.930*** (0.000)
4	13.470*** (0.000)	13.508*** (0.000)	13.551*** (0.000)	13.096*** (0.000)	12.892*** (0.000)
5	15.685*** (0.000)	15.207*** (0.000)	14.757*** (0.000)	14.019*** (0.000)	13.573*** (0.000)
First Subsample					
2	9.241*** (0.000)	9.877*** (0.000)	10.031*** (0.000)	9.328*** (0.000)	8.778*** (0.000)
3	11.325*** (0.000)	11.607*** (0.000)	11.435*** (0.000)	10.557*** (0.000)	9.954*** (0.000)
4	14.292*** (0.000)	13.426*** (0.000)	12.737*** (0.000)	11.561*** (0.000)	10.743*** (0.000)
5	18.247*** (0.000)	15.614*** (0.000)	14.094*** (0.000)	12.479*** (0.000)	11.434*** (0.000)
Second Subsample					
2	4.636*** (0.000)	4.736*** (0.000)	4.765*** (0.000)	4.708*** (0.000)	4.646*** (0.000)
3	5.586*** (0.000)	5.665*** (0.000)	5.911*** (0.000)	5.968*** (0.000)	6.146*** (0.001)
4	5.429*** (0.000)	5.743*** (0.000)	6.265*** (0.000)	6.324*** (0.000)	6.691*** (0.000)
5	5.033*** (0.000)	5.903*** (0.000)	6.581*** (0.000)	6.706*** (0.000)	7.034*** (0.000)

The BDS test statistic is presented in the table above (under an approximate normal distribution). The p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively.

Table 6.c. Brock-Dechert-Scheinkman test results (Ripple).

	ε				
m	0.5	0.75	1.00	1.25	1.50
Full Sample					
2	13.581***	13.779***	14.086***	14.215***	13.924***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	16.591***	15.770***	15.260***	14.756***	14.107***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	19.707***	17.285***	15.931***	14.846***	13.812***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	23.668***	19.175***	16.752***	15.093***	13.736***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
First Subsample					
2	10.931***	11.390***	11.821***	12.262***	11.980***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	14.078***	13.616***	13.387***	13.188***	12.369***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	16.833***	15.275***	14.248***	13.474***	12.216***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	19.969***	17.172***	15.413***	14.110***	12.431***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Second Subsample					
2	7.670***	7.953***	8.577***	9.085***	9.032***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	8.427***	8.588***	9.039***	9.447***	9.356***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	9,044***	9.014***	9.285***	9.475***	9.173***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	9.828***	9.306***	9.407***	9.439***	9.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

The BDS test statistic is presented in the table above (under an approximate normal distribution). The p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively.

Table 6.d. Brock-Dechert-Scheinkman test results (Litecoin).

	ε				
m	0.5	0.75	1.00	1.25	1.50
Full Sample					
2	14.201***	12.428***	10.593***	8.935**	7.589***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	19.793***	16.094***	13.614***	11.344***	9.423***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	27.115***	19.826***	15.781***	12.971***	10.682***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	37.262***	24.031***	18.033***	14.545***	11.812***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
First Subsample					
2	5.234***	5.692***	5.873***	6.051***	6.171***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	7.286***	7.315***	7.230***	7.209***	7.290***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	9.186***	8.552***	8.160***	7.941***	7.634***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	10.823***	9.678***	8.964***	8.453***	7.840***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Second Subsample					
2	4.983***	4.718***	4.361***	4.102***	4.151***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	6.487***	5.809***	5.173***	4.814***	4.831***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	7.736***	6.711***	5.717***	5.243***	5.120***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	8.776***	7.530***	6.218***	5.740***	5.440***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

The BDS test statistic is presented in the table above (under an approximate normal distribution). The p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively.

Table 6.e. Brock-Dechert-Scheinkman test results (Monero).

	ε				
m	0.5	0.75	1.00	1.25	1.50
Full Sample					
2	7.947***	7.833***	8.104***	8.667***	8.760***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	11.320***	10.643***	10.579***	10.841***	10.702***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	13.109***	11.897***	11.430***	11.478***	11.066***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	15.017***	13.101***	12.117***	11.809***	11.146***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
First Subsample					
2	7.895***	7.758***	8.195***	8.316***	7.543***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
3	10.228***	10.041***	10.080***	9.775***	8.772***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	10.935***	10.538***	10.307***	9.771***	8.609***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	12.188***	11.310***	10.637***	9.752***	8.395***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Second Subsample					
2	3.413***	3.115***	2.805***	3.038***	3.399***
	(0.001)	(0.002)	(0.000)	(0.002)	(0.001)
3	6.239***	5.237***	4.500***	4.490***	4.859***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
4	8.429***	6.707***	5.635***	5.359***	5.671***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
5	9.971***	7.635***	6.389***	5.731***	5.998***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

The BDS test statistic is presented in the table above (under an approximate normal distribution). The p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively.

Autoregressive Model of Order 5

Table 7.a reports the AR(5) results for the Bitcoin market. Individually speaking, nearly all of the lag-coefficients are insignificant. The lagged returns do not seem to have any impact on the current Bitcoin return. However, the p-value as described within the parentheses are based on the t-statistic, in which the hypotheses of H_0: $\beta_2 = 0$ to $\beta_6 = 0$ are tested separately. In order to make sense of whether the coefficients are jointly significantly different from zero, we need to conduct an F-test. The F test's null hypothesis cannot be rejected in any of the samples. This confirms the statement that the lagged Bitcoin returns have no impact at all on Bitcoin's current return. Past returns cannot be used for prediction and the Bitcoin market seems to be consistent with the random walk hypothesis. Testing up to five lags coincides with testing the joint significance of more than five lags, as the results suggested the same.

Table 7.a. AR(5) regression for Bitcoin.

Bitcoin	Full Sample	First Subsample	Second Subsample
Constant	0.302**	0.302**	0.302
	(0.016)	(0.019)	(0.158)
$R_{BTC,t-1}$	-0.000	-0.008	0.003
	(0.988)	(0.852)	(0.935)
$R_{BTC,t-2}$	0.012	-0.083*	0.050
	(0.695)	(0.055)	(0.249)
$R_{BTC,t-3}$	0.005	-0.026	0.013
	(0.882)	(0.546)	(0.758)
$R_{BTC,t-4}$	-0.035	0.080*	-0.080*
	(0.248)	(0.063)	(0.064)
$R_{BTC,t-5}$	0.044	0.016	0.052
	(0.148)	(0.706)	(0.230)
F-stat.	0.724	1.701	1.243
	(0.605)	(0.133)	(0.288)

Regression results of the AR(5) model. The coefficients are presented along the individual p-value (within parentheses). *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively. The F-statistic for the lags' joint significance is presented underneath (H_0: $\beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = 0$).

The test results for Ethereum (Table 7.b), however, do reject the null hypothesis in the full sample period, as well as in the first subsample period. The one-day lag and the three-day lag are weakly (*) to strongly (***) significant. In the joint hypothesis test, the null hypothesis is rejected. At least one of the independent variables seems to have an effect on the current return (as is already evident from looking at the lags individually). Ethereum seems to become random over time, as the null hypothesis cannot be rejected in the second subsample. Overall, it is still possible to employ profitable trading strategies when looking at the past return sequence.

Table 7.b. AR(5) regression for Ethereum.

Ethereum	Full Sample	First Subsample	Second Subsample
Constant	0.462**	0.484	0.402
	(0.036)	(0.143)	(0.167)
$R_{ETH,t-1}$	-0.055*	-0.120***	0.025
	(0.070)	(0.006)	(0.569)
$R_{ETH,t-2}$	0.016	0.009	0.021
	(0.607)	(0.833)	(0.620)
$R_{ETH,t-3}$	0.091***	0.119***	0.044
	(0.003)	(0.006)	(0.308)
$R_{ETH,t-4}$	0.008	-0.002	0.018
	(0.790)	(0.962)	(0.681)
$R_{ETH,t-5}$	-0.004	-0.065	0.073*
	(0.890)	(0.132)	(0.091)
F-stat.	2.487**	3.609***	1.012
	(0.030)	(0.003)	(0.410)

Regression results of the AR(5) model. The coefficients are presented along the individual p-value (within parentheses). *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively. The F-statistic for the lags' joint significance is presented underneath (H_0: $\beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = 0$).

In the case of Ripple, the F-statistic rejects the null hypothesis in each of the sample periods. The p-value shows that the results are highly significant, i.e., at least one of the lags has a significant impact on the current return. Table 7.c shows that the 1-day, 2-day, and 3-day lags are the main causes for rejection of the F-test. The return on Ripple is inconsistent with the random walk model.

Table 7.c. AR(5) regression for Ripple.

Ripple	Full Sample	First Subsample	Second Subsample
Constant	0.330	-0.114	0.610
	(0.215)	(0.675)	(0.175)
$R_{XRP,t-1}$	-0.161***	-0.478***	-0.068
	(0.000)	(0.000)	(0.113)
$R_{XRP,t-2}$	0.082***	-0.223***	0.150***
	(0.008)	(0.000)	(0.001)
$R_{XRP,t-3}$	0.093***	-0.066	0.097**
	(0.002)	(0.169)	(0.026)
$R_{XRP,t-4}$	0.008	0.015	-0.045
	(0.785)	(0.744)	(0.296)
$R_{XRP,t-5}$	0.036	0.018	0.041
	(0.234)	(0.671)	(0.335)
F-stat.	9.612***	24.950***	4.442***
	(0.000)	(0.000)	(0.001)

Regression results of the AR(5) model. The coefficients are presented along the individual p-value (within parentheses). *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively. The F-statistic for the lags' joint significance is presented underneath (H_0: $\beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = 0$).

The results for Litecoin (Table 7.d) show that Litecoin transitioned from a non-random return sequence into a random return sequence. In general, Litecoin seems to support the random walk model, as the null hypothesis cannot be rejected in the full sample. The first subsample shows that the property of non-randomness is due to lags 2 and 3. In the full sample and the second subsample, none of the lags are significant.

Table 7.d. AR(5) regression for Litecoin.

Litecoin	Full Sample	First Subsample	Second Subsample
Constant	0.259	0.052	0.462
	(0.146)	(0.691)	(0.167)
$R_{LTC,t-1}$	0.009	0.014	0.005
	(0.760)	(0.742)	(0.910)
$R_{LTC,t-2}$	-0.014	-0.109**	0.000
	(0.646)	(0.012)	(0.995)
$R_{LTC,t-3}$	0.230	-0.105**	0.042
	(0.448)	(0.015)	(0.325)
$R_{LTC,t-4}$	0.033	0.026	0.032
	(0.272)	(0.539)	(0.461)
$R_{LTC,t-5}$	0.010	-0.007	0.007
	(0.737)	(0.862)	(0.876)
F-stat.	0.447	2.758**	0.315
	(0.816)	(0.018)	(0.904)

Regression results of the AR(5) model. The coefficients are presented along the individual p-value (within parentheses). *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively. The F-statistic for the lags' joint significance is presented underneath (H_0: $\beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = 0$).

Last but not least, Monero displays similar results to that of Ripple. The test rejects the null hypothesis in the full sample period and the first subsample period. Over the years, Monero changes from being non-random to becoming somewhat random. The current return has a significant and direct link to its 1-day lagged value. Trading on historical information is a profitable strategy in the case of Monero.

Table 7.e. AR(5) regression for Monero.

Monero	Full Sample	First Subsample	Second Subsample
Constant	0.581**	0.750*	0.393
	(0.021)	(0.055)	(0.221)
$R_{XMR,t-1}$	-0.150***	-0.185***	-0.102**
	(0.000)	(0.000)	(0.018)
$R_{XMR,t-2}$	-0.016	-0.042	0.011
	(0.596)	(0.337)	(0.796)
$R_{XMR,t-3}$	0.007	-0.027	0.048
	(0.824)	(0.536)	(0.266)
$R_{XMR,t-4}$	-0.013	-0.014	-0.022
	(0.676)	(0.734)	(0.614)
$R_{XMR,t-5}$	0.010	0.002	0.019
	(0.738)	(0.965)	(0.656)
F-stat.	4.980***	3.711***	1.571
	(0.000)	(0.003)	(0.166)

Regression results of the AR(5) model. The coefficients are presented along the individual p-value (within parentheses). *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively. The F-statistic for the lags' joint significance is presented underneath (H_0: $\beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = 0$).

The verdict on weak form market efficiency

Table 8 provides an overview of the test results with respect to the first null hypothesis. Bitcoin passes three out of six tests, which indicates that the results are somewhat ambiguous. Even though the Bitcoin market seems to be somewhat in line with random walk, it cannot be purely classified as an efficient market. None of the other cryptocurrencies manages to pass even more than 1 test. With a few minor exceptions in some of the sample periods, it seems rather evident that the test results for Ethereum, Ripple, Litecoin, and Monero are not in accordance with the random walk theory. Based on these results it seems that the cryptocurrency market is fairly predictable and inefficient. The verdict on Bitcoin is not entirely conclusive, as it does not strictly pass of fail the majority of the tests.

Table 8. Overview of diagnostic tests of randomness on the cryptocurrency market.

Currency	LB	ADF	PP	Runs	BDS	AR(5)
BTC	✓	×	×	✓	×	✓
ETH	×	×	×	✓	×	×
XRP	×	×	×	×	×	×
LTC	×	×	×	×	×	✓
XMR	×	×	×	×	×	×

This table presents an overviews of the statistical tests on the return series. The cryptocurrencies considered are Bitcoin (BTC), Ethereum (ETH), Ripple (XRP), Litecoin (LTC), and Monero (XMR). The table shows the results of the following tests: the Ljung-Box (LB) test; the Augmented Dickey-Fuller (ADF) test; the Phillips-Perron (PP) test; the Wald-Wolfowitz runs test; the Brock-Dechert-Scheinkman (BDS) test; and an autoregressive model of order five (AR(5)). The check mark (✓) indicates that the results are consistent with the random walk model. The cross (×) indicates the results reject the random walk model.

6.3. Volume-Price Relationship

The first step of the Granger-causality test is to make sure the variables are stationary. In case of non-stationarity, the series first has to be differenced up to the point it becomes stationary, i.e., eliminating the unit root. The table in Appendix C shows that all volume series are stationary. Next, the research runs the actual Granger-causality test with the null of no Granger-causality. Table 8 reports the test results[19]. Ethereum, Ripple, Litecoin, and Monero cannot reject the null hypothesis in any of the samples. The joint effect of the lagged values of the return series and the trading volume series are insignificantly different from zero. Trading volume data does not Granger-cause returns. The results for Bitcoin imply the same, except for the first subsample where the null is rejected. Bitcoin volume data is useful for the prediction of returns only in the first subsample period. The overall combined effect shows that investors generally are incapable of using trading volume to generate excess return. Technical analysis based on trading volume data does not lead to extra profits on top of the market return.

[19] Including less or more than five lags leads to the same results, i.e., rejecting the null hypothesis for all cryptocurrencies in all samples (or in the case of Bitcoin: not rejecting the null in the first subsample).

Table 9. Granger-causality test.

Currency	Full Sample	First Subsample	Second Subsample
BTC	1.617	2.574**	0.361
	(0.153)	(0.026)	(0.875)
ETH	1.740	1.653	1.187
	(0.123)	(0.144)	(0.314)
XRP	0.607	0.581	1.047
	(0.694)	(0.715)	(0.389)
LTC	0.278	0.893	0.160
	(0.925)	(0.486)	(0.977)
XMR	0.996	0.814	0.528
	(0.419)	(0.540)	(0.755)

Granger-causality test results. The F-statistic is presented in the table and the p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively.

7. Conclusion

This paper sets out to examine whether the cryptocurrency market resides in a state of efficiency, as proposed according to the Efficient Market Hypothesis. With use of daily price data and daily trading volume data, weak form efficiency is tested for the following cryptocurrencies: Bitcoin, Ethereum, Ripple, Litecoin, and Monero. The analysis separates two hypotheses.

In the first hypothesis, a set of diagnostic tests are executed in order to see whether historical cryptocurrency values can help predict the future. As it turns out, Ethereum, Ripple, Litecoin, and Monero defy nearly all randomness tests. These cryptocurrencies are not random, thus indicating that they are inefficient. The results for Bitcoin are quite ambiguous. Bitcoin passes the Ljung-Box test for serial correlation, the Wald-Wolfowitz runs test, and the autoregressive process of order 5. On the other hand, Bitcoin fails to support the Augmented Dickey-Fuller test and Phillips-Perron test for unit root, and the Brock-Dechert-Scheinkman test. In general, it seems that technical analysis is a rather profitable move when applied on the cryptocurrency market. In the second hypothesis, trading volume data plays a part in hypothesis testing. The Granger-causality test examines whether trading volume is useful for the prediction of cryptocurrency returns. The test results cannot reject the null hypothesis. Trading volume data does not predict cryptocurrency returns. The final statement concludes that the cryptocurrency market is rather inefficient, and that the Efficient Market Hypothesis is rejected. Historical prices can be used to generate excess profits. However, the use of trading volume data does not systematically lead to any excess return.

As mentioned before, the cryptocurrency market is still very much in its infancy. This may have been a drawback in this research. The timespan in this research only consists of three years of data in the full sample. The reason for choosing this sample period is simply due to the inability to acquire any sensible aggregated data from the period prior to the start of the sample period. The results may turn out to be rather different in a few years' time, when there is more data available. Further research on this topic may shed light on the characteristics of the cryptocurrency market. A suggestion for further research is the inclusion of a larger time period as well as possibly including other (well-established) cryptocurrencies. The debate on efficiency may also lead to more research on the possible existence of anomalies within the cryptocurrency market.

References

Alexander, S. (1961). Price movements in speculative markets: trends or random walks. *Industrial Management Review, 2*(2), 7-26.

Alexander, S. (1964). Price movements in speculative markets: trends or random walks, number 2. *Industrial Management Review, 5*(2), 25-46.

Antoniou, A., Ergul, N., Holmes, P., & Priestley, R. (1997). Technical analysis, trading volume and market efficiency: evidence from an emerging market. *Applied Financial Economics, 7*(4), 361-365.

Antonopoulos, A. (2015). Mastering Bitcoin: unlocking digital cryptocurrencies. In A. Antanopoulos, *Mastering Bitcoin: unlocking digital cryptocurrencies* (p. 185). Sebastopol, California, The United States of America: O'Reilly Media Inc.

Böhme, R., Christin, N., Edelman, B., & Moore, T. (2015). Bitcoin: economics, technology, and governance. *Journal of Economic Perspectives, 29*(2), 213-238.

Balcilar, M., Bouri, E., Gupta, R., & Roubaud, D. (2017). Can volume predict Bitcoin returns and volatility? A quantiles-based approach. *Economic Modelling, 64*, 74-81.

Bariviera, A. (2017). The inefficiency of Bitcoin revisited: a dynamic approach. *Economics Letters, 161*, 1-4.

Bartos, J. (2015). Does Bitcoin follow the hypothesis of efficient market? *International Journal of Economic Sciences, IV*(2), 10-23.

Blume, L., Easley, D., & O'Hara, M. (1994). Market statistics and technical analysis: the role of volume. *The Journal of Finance, 49*(1), 153-181.

Box, G., & Ljung, G. (1978). On a measure of lack of fit in time series models. *Biometrika, 65*(2), 297-303.

Brauneis, A., & Mestel, R. (2018). Price discovery of cryptocurrencies: Bitcoin and beyond. *Economics Letters, 165*, 58-61.

Brock, A., Dechert, D., LeBaron, B., & Scheinkman, J. (1996). A test for independence based on the correlation dimension. *Econometric Reviews, 15*(3), 197-235.

Brock, W., Hsieh, D., & LeBaron, B. (1991). Nonlinear dynamics, chaos, and instability: statistical theory and economic evidence. In W. Brock, D. Hsieh, & B. LeBaron, *Nonlinear dynamics, chaos, and instability: statistical theory and economic evidence* (pp. 41-54). Cambridge, Massachusetts, United States: The MIT Press.

Caporale, G., & Plastun, A. (2017). *The day of the week effect in the cryptocurrency market.* Working Paper, Brunel University London, Department of Economics and Finance, London.

Caporale, G., Gil-Alana, L., & Plastun, A. (2018). Persistence in the cryptocurrency market. *Research in International Business and Finance, 46*, 141-148.

Chu, J., & Nadarajah, S. (2017). On the inefficiency of Bitcoin. *Economics Letters, 150*, 6-9.

Dickey, D., & Said, S. (1984). Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika, 71*(3), 599-607.

Fama, E. (1970). Efficient capital markets: a review of theory and empirical work. *The Journal of Finance, 25*(2), 383-417.

Fama, E., & Blume, M. (1966). Filter rules and stock-market trading. *The Journal of Business, 39*(1), 226-241.

Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica, 37*(3), 424-438.

Jiang, Y., Nie, H., & Ruan, W. (2018). Time-varying long-term memory in Bitcoin market. *Finance Research Letters, 25*, 280-284.

Karpoff, J. (1987). The relation between price changes and trading volume: a survey. *Journal of Financial and Quantitative Analysis, 22*(1), 109-126.

Kurihara, Y., & Fukushima, A. (2017). The market efficiency of Bitcoin: a weekly anomaly perspective. *Journal of Applied Finance & Banking, 7*(3), 57-64.

MacKinnon, J. (1996). Numerical distribution functions for unit root and cointegration tests. *Journal of Applied Econometrics, 11*, 601-618.

Malkiel, B. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives, 17*(1), 59-82.

Nakamoto, S. (2008). *Bitcoin: a peer-to-peer electronic cash system.* Retrieved from www.bitcoin.org: https://bitcoin.org/bitcoin.pdf

Perron, P., & Phillips, P. (1988). Testing for a unit root in time series regression. *Biometrika, 75*(2), 335-346.

Urquhart, A. (2016). The inefficiency of Bitcoin. *Economics Letters, 148*, 80-82.

Wald, A., & Wolfowitz, J. (1940). On a test whether two samples are from the same population. *The Annals of Mathematical Statistics, 11*(2), 147-162.

Appendix A.

Unit Root and Non-Stationarity

Consider, for simplicity, the random walk equation as described in Section 3.1.3. (Equation (7), which does not include the intercept or trend term):

$$y_t = y_{t-1} + \varepsilon_t$$

As the random walk states that the best forecast is simply the preceding value, it can be safely said that the observation at time *t* is nothing more than the previous observation plus an error term. This implies that for every increase in time *t* the forecast of y_t remains the same (y_t equals *y*'s original value, y_0), with an increase in the error term (as it is the sum of all preceding error terms). The random walk model only applies to situations in which the process contains a unit root. This becomes evident when rewriting the equation:

$$y_t = \varphi y_{t-1} + \varepsilon_t$$

For *y* at time *t-1*, the equation becomes:

$$y_{t-1} = \varphi y_{t-2} + \varepsilon_{t-1}$$

and so on, where φ denotes the level of the root, here assumed to be either 1 (H_0: φ =1, unit root) or smaller than 1 (H_a: $\varphi < 1$). Continuing the process further down the line (substituting y_{t-1} into y_t) leads to the following:

$$y_t = \varphi(\varphi y_{t-2} + \varepsilon_{t-1}) + \varepsilon_t$$

$$y_t = \varphi^2 y_{t-2} + \varphi\, \varepsilon_{t-1} + \varepsilon_t$$

When φ reaches unity, y_t equals y_{t-1}, or equivalently y_{t-2}, plus the sum of the error term at time *t-1* and *t*. Equivalently, for *T* amount of previous time periods:

$$y_t = \varphi^T y_0 + \varphi\, \varepsilon_{t-1} + \varphi^2 \varepsilon_{t-2} + \varphi^3 \varepsilon_{t-3} + \ldots + \varphi^T \varepsilon_0 + \varepsilon_t$$

Therefore, in case of unit root, the level of y_t is actually determined by the sum of the error terms (every future *y* increases due to ε). Moreover, as *y* increases for every increment of *t*, the variance of *y* increases as well. An ever-increasing variance can be problematic under the circumstances of hypothesis testing, as this results in heteroskedasticity, or volatility clustering, which can lead to unreliable results. In this case, the next variance ultimately will be underestimated. In all other cases (with φ smaller than 1), the variance of y_t falls to zero. The faster the variance converges to zero, the better predictable the properties of the series are (stationarity). Under normal circumstances, models require series to be stationary. However, when searching for random walk, non-stationarity is required, as it implies that the mean and variance cannot be predicted. Statistically, the aforementioned equation is not used in the Dickey-Fuller test. Alternatively, the regression uses the first-differenced equation.

$$y_t - y_{t-1} = (\varphi-1)\, y_{t-1} + \varepsilon_t$$

$$\Delta y_t = \delta y_{t-1} + \varepsilon_t$$

where δ stands for $(\varphi - 1)$. The Dickey-Fuller test is a one-sided test with H_0: $\delta = 0$, unit root; H_a: $\delta < 0$. However, the Dickey-Fuller test only examines the case of stationarity versus non-stationarity. The exceptional case of $\delta = -1$ (or $\varphi = 0$) is not explicitly mentioned. For $\delta = -1$ there is a white noise process in which:

$$\Delta y_t = -y_{t-1} + \varepsilon_t$$

$$y_t = \varepsilon_t$$

On another note, to account for possible serial correlation within the series, this research uses the Augmented Dickey-Fuller test. The Augmented Dickey-Fuller test equation uses lags as regressors. With the inclusion of an additional lag, the number of observations decreases. Therefore, one should be careful when deciding the amount of lags to include. In this research the number of lags is chosen under the Schwarz Information Criterion (SIC).

Appendix B.

Cryptocurrency LN Return Plots

- *Full Sample (09/01/2015 – 08/31/2018)*

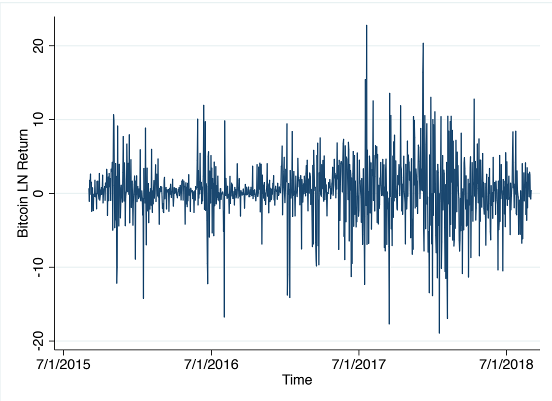

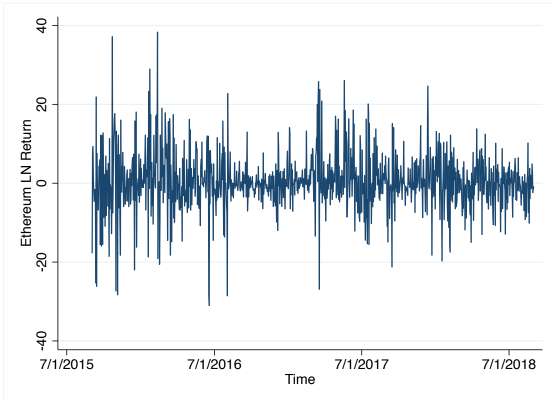

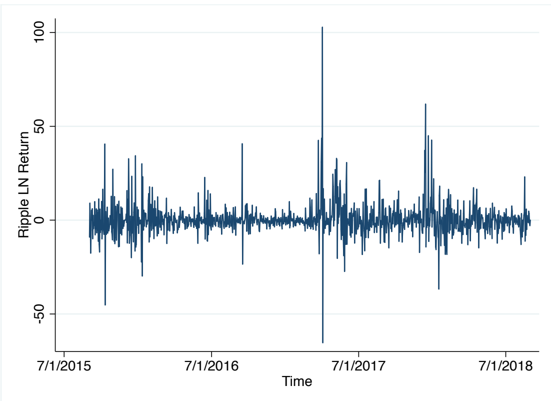

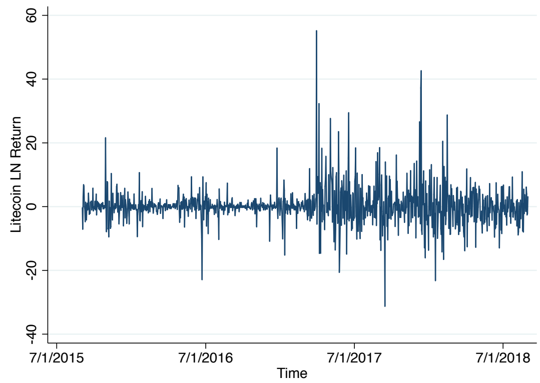

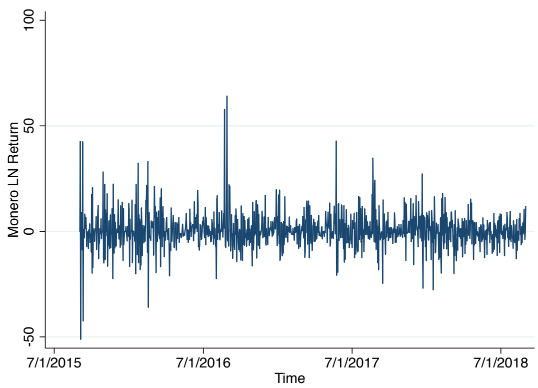

- *First Subsample (09/01/2015 – 02/28/2017)*

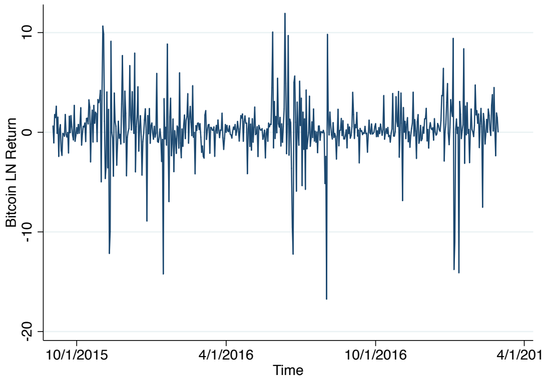

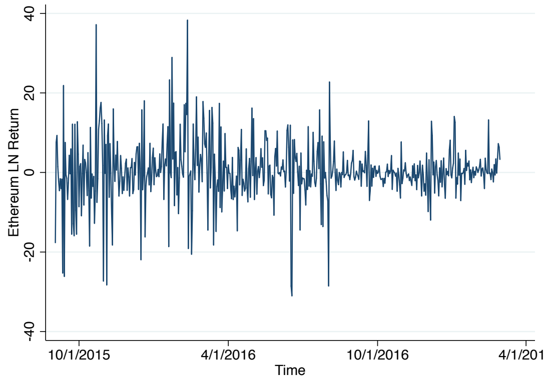

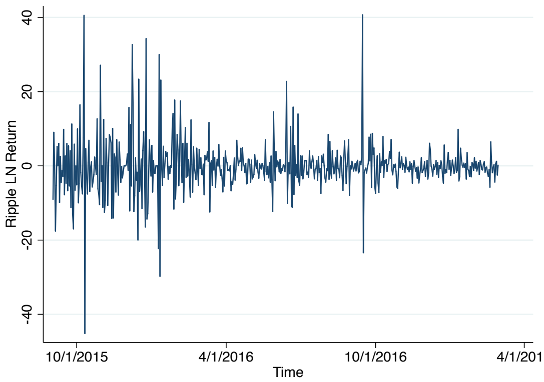

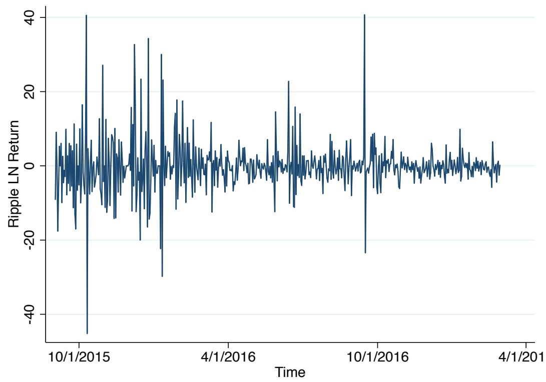

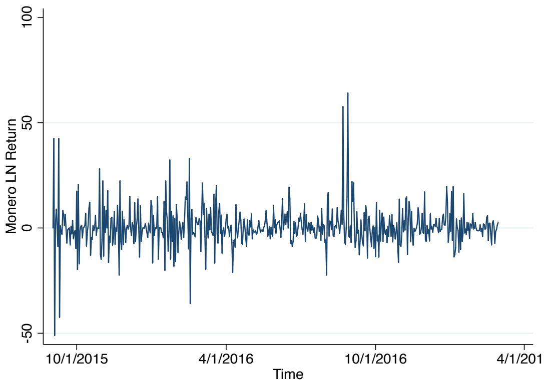

- *Second Subsample (03/01/2017 – 08/31/2018)*

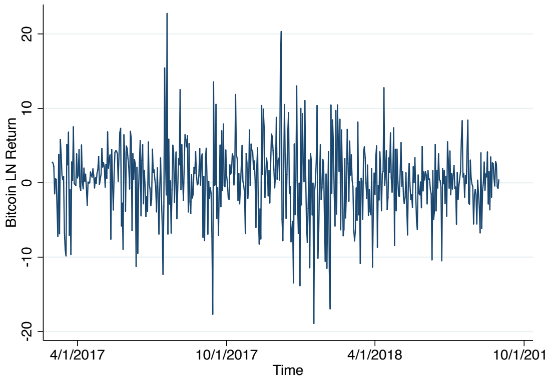

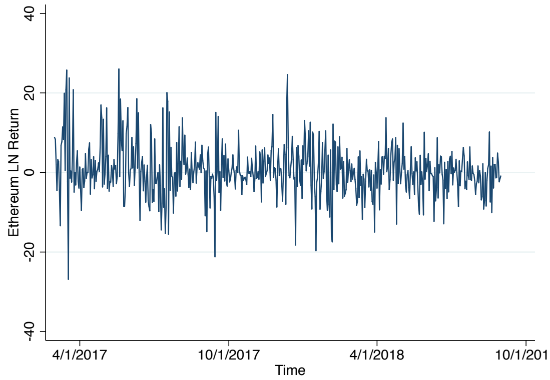

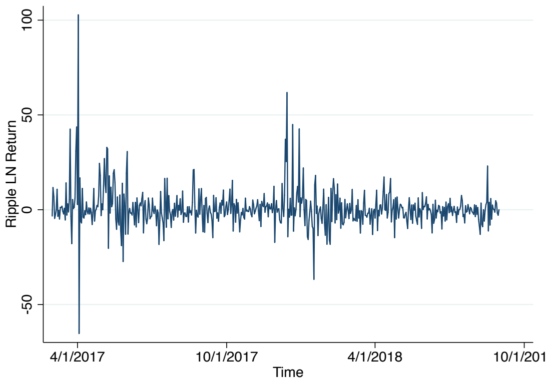

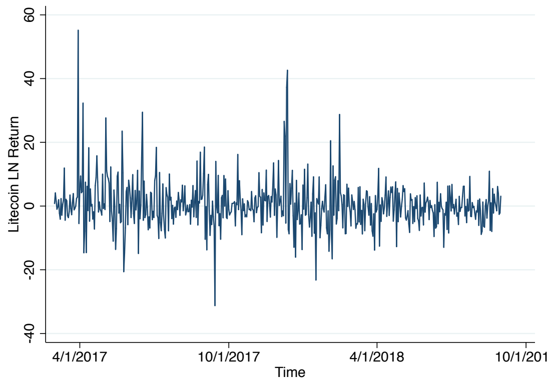

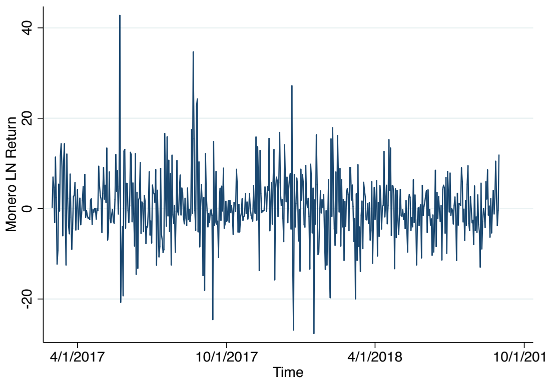

Appendix C.

Augmented Dickey Fuller Test Results of the Trading Volume Time Series

Table. Ap.C. Results of the Augmented Dickey-Fuller test for all sample periods on trading volume data.

Currency	Full Sample	First Subsample	Second Subsample
BTC	-22.593***	-16.125***	-17.194***
	(0.000)	(0.000)	(0.000)
ETH	-20.581***	-14.700***	-15.827***
	(0.000)	(0.000)	(0.000)
XRP	-21.370***	-14.429***	-16.252***
	(0.000)	(0.000)	(0.000)
LTC	-22.247***	-16.217***	-14.873***
	(0.000)	(0.000)	(0.000)
XMR	-19.189***	-13.683***	-19.400***
	(0.000)	(0.000)	(0.000)

The table reports the Augmented Dickey-Fuller test statistic, and the MacKinnon p-value is reported within parentheses. *, **, and *** mark the significance levels at 10%, 5%, and 1%, respectively.