Technische Universiteit Delft
Faculteit Elektrotechniek, Wiskunde en Informatica
Delft Institute of Applied Mathematics

Tail Risk in Cryptocurrencies

Verslag ten behoeve van het
Delft Institute of Applied Mathematics
als onderdeel ter verkrijging

van de graad van

BACHELOR OF SCIENCE
in
TECHNISCHE WISKUNDE

door

Linda Michelle Leeuwestein

Delft, Nederland
Juli 2018

BSc verslag TECHNISCHE WISKUNDE

“Tail Risk in Cryptocurrencies”

Linda Michelle Leeuwestein

Technische Universiteit Delft

Begeleider

Dr. P. Cirillo

Overige commissieleden

Dr. W.M. Ruszel | Dr. J.G. Spandaw

July, 2018 | Delft

Abstract

In this research, the returns of four cryptocurrencies (Bitcoin, Litecoin, Ripple and Ethereum) were analyzed in order to answer the following research question: "How do the returns of Bitcoin and other altcoins behave over time, and what can we say about extreme values for losses and profits?" With respect to volatility, cryptocurrencies can still be considered extremely volatile. For Bitcoin, the least volatile of the four, we found an annual volatility of approximately 70% based on daily exchange rates. For Ethereum, the most volatile of all four, this percentage was closer to 130%. Also, several distributions were fitted on the returns. It is shown that the Generalized Hyperbolic Distribution is the best fit for all four cryptocurrencies, apart from the tails in some cases.

The tails were investigated seperately by using Extreme Value Analysis and by looking into both empirical and theoretical risk quantities (the Value at Risk and Expected Shortfall). Bitcoin appears to be the least risky of all four cryptocurrencies, but also the least profitable, whereas Ripple appears to be the most risky and also the most profitable.

Compared to previous research, Bitcoin has also become less risky, showing a less fat tail for the losses than before. For Litecoin and Ripple, the reverse is true, as they appear to have become riskier. For Ethereum, no comparisons could be made, as this is a relatively new cryptocurrency that has not been investigated much yet. When tested for Paretianity, the left tails of Litecoin and Ripple appear to Pareto distributed: the losses seem to exhibit heavy tail behavior. For the profits, the tails turned out to be even heavier and can therefore also be considered Paretian. These results were confirmed by Maximum to Sum ratio plots, indicating infinite third and fourth moments for the losses and profits of Litecoin and Ripple, but not for Bitcoin and Ethereum. The results have implications for investment and risk management purposes.

Contents

1 Introduction **5**

2 Methodology **9**
2.1 Descriptive Statistics . . . 9
2.2 Volatility . . . 11
2.3 Value at Risk and Expected Shortfall . . . 11
2.4 Fitting of distributions . . . 12
2.5 Testing for Paretianity . . . 17
2.6 Maximum to Sum . . . 18

3 Results **20**
3.1 Descriptive Statistics . . . 20
3.1.1 Exchange Rates . . . 20
3.1.2 Log returns . . . 24
3.1.3 Summary Statistics . . . 27
3.2 Volatility . . . 29
3.3 Value at Risk and Expected Shortfall . . . 31
3.4 Fitting of distributions . . . 37
3.4.1 Visual Explanations . . . 37
3.4.2 Kolmogorov-Smirnov and Chi-Square statistics . . . 42
3.4.3 Theoretical Value at Risk and Expected Shortfall . . . 45
3.5 Testing for Paretianity . . . 47
3.5.1 Losses . . . 47
3.5.2 Profits . . . 52
3.6 Maximum to Sum . . . 54

4 Conclusion **62**

5 Discussion **65**

6 Appendix **69**
6.1 Datasets . . . 69
6.2 Plots . . . 69
6.3 Source code . . . 83

1 Introduction

Some say that cryptocurrencies such as Bitcoin are the future. Others believe that Bitcoin is the biggest scam in history.
What's certain is that Bitcoin has gained a lot of attention from the media, academics and the finance industry and is therefore a hot topic of investigation.

Most people today have heard of Bitcoin, a digital currency and electronic peer-to-peer payment network (meaning there is no third party such as a bank involved), with its price determined purely by market demand and supply [27]. In 2009, Bitcoin was developed by (an) anonymous hacker(s) (with pseudonym Satoshi Nakamoto) [19] and became the first decentralized cryptocurrency. Within four years, the value of Bitcoin already reached levels beyond $ 1,000,- by the end of 2013.
At present, five years later, Bitcoin has experienced a high of around $20,000,- but has, in the past few months, dropped to values around $6,500,- again.

Since Bitcoin was created in 2009, many other cryptocurrencies have been created.
Coinmarketcap [3] shows, as of 18 June 2018, over 1600 of cryptocurrencies, ordering them in terms of their so called 'Market Cap'.
Market Cap or Market Capitalization is a way to rank the relative size of a cryptocurrency. It's calculated by multiplying the spot price of the cryptocurrency by the so called 'Circulating Supply', which is the best approximation of the number of coins that are circulating in the market and in the general public's hands (as opposed to 'Total Supply', indicating the total number of coins in existence right now, and the 'Max Supply' which is the best approximation of the maximum number of coins that will ever exist in the lifetime of the cryptocurrency).
Even though over 1600 cryptocurrencies exist today, the largest part of these currencies have an extremely low value and will probably become inactive shortly after they are launched [10]. In fact, 90% of the market belongs to the top 30 [3], with number 30 (Bytom) representing only 0.2% of the market, giving an indication of the speed with which the market share (calculated by dividing the currency's market cap by the total market cap) drops when we look further down the list.

In this research, Bitcoin, Ethereum, Ripple and Litecoin will be investigated.
A short description will be given for each of the currencies:

Bitcoin (or BTC) was, as mentioned before, founded in 2009, built upon blockchain technology, and is the first and most popular decentralized currency of all, taking the first position in the ranking based on market capitalization [3]. It was designed with the idea of using cryptography to control the creation and transfer of money, instead of relying on central authorities [2],[19].
What makes Bitcoin fundamentally different from fiat currencies, lies in the following characteristics: first, the technology behind Bitcoin is based on cryptographic proof rather than traditional trust [19]. This provides many possible advantages over traditional payment methods, such as lower transaction costs and pseudo-anonymity) [2]. Also, for investors, Bitcoin can be used as a store of value, which some consider better than gold, given that the number of Bitcoins that can ever be mined is finite (equal to 21 million [2]), whereas for gold this is more complicated to determine.
Furthermore, Bitcoin is regarded as being beyond the government regulation's reach, meaning the government cannot manipulate or interfere with the value of Bitcoin. The downside, however, is that there is no central authority to back the value of Bitcoin, and no one to ensure that the money supply is controlled.

Litecoin (or LTC) was based on the Bitcoin protocol, and developed in 2011 by Charles Lee with some technical differences from that of its ancestor (for technical details, [17] can be consulted). It is considered one of Bitcoin's leading rivals (even though Charlie Lee never wanted to compete with Bitcoin: in his vision Litecoin is a supplement), as it features some implementation improvements compared to Bitcoin. It has, for example, the capacity to process transactions much faster than Bitcoin (i.e. a much shorter confirmation time). However, it's also less secure because the algorithm is solved with less power. Litecoin is currently the fifth largest cryptocurrency in terms of market capitalization [3].

Ripple (with currency component XRP) was founded in 2012 and also has a mathematical foundation like Bitcoin [25]. It is considered a 'high cap' currency, currently at position 3 in the ranking based on market capitalization [3], and a serious contestant for the title of Bitcoin successor. However, it runs on a completely different philosophy. It can be centrally controlled and transactions can be reversed by the Ripple foundation which means it's a centrally controlled currency.

Ethereum (with currency component ETH) was only founded in 2015, and is therefore the 'youngest' cryptocurrency of the four that are considered in this research. Ethereum is a protocol that facilitates the deployment of smart-contracts and distributed applications [8]. The Ethereum network runs on Ether which is the native currency of the network. Ether can be used for rewarding miners for performing transactions on the network but it may also be used as a cryptocurrency like Bitcoin.

Currency crises, stock market crashes and other extreme events that might lead to large losses from an investment's point of view, are not uncommon in traditional financial markets.
Given that cryptocurrencies are characterized by even larger volatility swings[21], these extreme events are expected to be even more extreme for the cryptocurrency market.
Whereas much research has focused on the risks of Bitcoin with respect to security, fraud and criminal activities, only recently the focus has shifted to investing in Bitcoin and the risks from an investor's perspective. For this purpose, it is important to fully understand the underlying risks of cryptocurrencies.

For example, in 2016, Osterrieder, Lorenz and Strika [21] have investigated Bitcoin, Litecoin, Ripple among some other cryptocurrencies (but not Ethereum), with respect to their relative price increases and decreases, also known as returns. They have provided a statistical analysis and an extreme value analysis of these returns.
They based their results on data from June 2014 to September 2016 and found that the returns were extremely volatile and much riskier than traditional fiat currencies, exhibiting heavier tail behavior as well. They also discovered that the volatility itself showed large swings and instabillity over time, resulting in extremely high volatility regularly.
Bitcoin, however, did come out as the least volatile out of all the cryptocurrencies investigated, with an annual volatility of 'only' 62%.
Similar results with respect to volatility were obtained by Chan et.al [4] who found that the range of volatility appeared to be smallest for Bitcoin compared to the other cryptocurrencies).
It will be interesting to see whether these same results will be obtained in the current research, based on data from January 2015 to April 2018 (i.e. with an overlap of 21 months present in the data compared to [21], but 19 new, more recent months).
Also, the authors from [21] did not include Ethereum in their research, as it was a relatively new cryptocurrency at the time, and therefore less interesting and reliable to investigate in terms of volatility and risk. Given that Ethereum is currently a bit older, and also the number 2 based on Market Capitalization, it is included in the current research, with data used from January 2016 to April 2018, and it is therefore interesting to see whether its behavior is similar to that of the other cryptocurrencies.

Osterrieder et. al. also studied the Value at Risk and Expected Shortfall, providing important information with respect to risk analysis. They have provided a detailed overview of the Historical VaR and ES for both the left as well as the right tails of the distributions of log returns. They found that Bitcoin came out with the lowest VaR and ES of all cryptocurrencies, after which came Litecoin and then Ripple (considering only those currencies that are also included in the current research). The values for the losses turned out to be higher for the losses than for the profits, indicating that the left tails were heavier than the right ones. In the current research, these risk measures are also calculated, and so it will be interesting to compare them to these figures. It will also be interesting to see where Ethereum fits within this ranking. Finally, the authors of [21] have fitted extreme value distributions and concluded that Bitcoin, Litecoin and Ripple show very risky behavior.

In 2017, five months after the research described before ([21]) was published, Osterrieder and Lorenz [22] have again investigated the risk of an investment in Bitcoin from a statistical point of view. This time, they used data from September 2013 to September 2016, and only investigated Bitcoin, leaving out the altcoins (which is the general term for all cryptocurrencies other than Bitcoin).
The conclusions with respect to volatility, Value at Risk and Expected Shortfall are very similar to those in [21]: Bitcoin is extremely volatile compared to traditional G10 currencies, with a loss of more than 10% to be expected within 20 days.

While Bitcoin is still by far the largest cryptocurrency, ever since its creation nine years ago, last year, the market has made significant room for other cryptocurrencies, taking a large part of the market share away from Bitcoin.
Only one year ago, 81% of the total market belonged to Bitcoin, meaning its share has halved since then, given that as of 18 June, 2018, Bitcoin represents only 40% of the total market.
Furthermore, 81% of the total market now belongs to the top 12 of all cryptocurrencies [3], giving a clear indication of the relevance to investigate and place more focus on other, 'high cap', cryptocurrencies besides Bitcoin as well.
As of 18 June 2018, Bitcoin, the largest cryptocurrency, represents, with a Market Cap of over 110 billion USD, about 40 % of the total market of cryptocurrencies [3], followed by Ethereum (18%) and Ripple (7%). The fourth to tenth place are taken by Bitcoin Cash, Eos, Litecoin, Stellar, Cardano, IOTA and Tron (note that this order is by no means fixed and in fact, changes almost daily, even for the high cap currencies).
Given that the altcoins have taken up quite a large market share from Bitcoin in only one year time, it will be interesting to see how this has affected their behavior as well as Bitcoin's in terms of the distributions of returns, volatility, Value at Risk and Expected Shortfall.

What is also an interesting research topic with respect to cryptocurrencies, is the fitting of distributions on the returns. Most people would probably agree that financial returns in general, are uncertain or risky. Information with respect to the distribution of the returns, is an attempt to chart this uncertainty. When the underlying distribution is known, probability calculations and predictions for future behavior can be made more easily.
Chan et.al. [4] investigated cryptocurrencies with respect to their return distributions, and discovered that the Generalized Hyperbolic Distribution gave the best fit for Bitcoin and Litecoin, whereas the Normal Inverse Gaussian Distribution and the Laplace Distribution appeared to fit the other, smaller currencies well. The authors used data from June 2014 to February 2017.
It will be interesting to see whether these results can be replicated for more recent data used in the current research, using data from January 2015 to April 2018. Also, Chan et. al. [4] did not include Ethereum, as it only started trading in 2015 and they therefore did not have enough data available. Ethereum was included in the current research and it is therefore interesting to find out what distribution will fit this cryptocurrency best.

When fitting theoretical distributions to data, it often occurs that the distribution fits the data well in mid regions (high density regions), but poorly in the tails, as there are fewer data points in the end regions, and therefore a model will most likely be selected based on the higher density regions[6].
Also, several empirical studies have established that the distributions of speculative asset returns tend to have heavier tails than the Normal Distribution tails [18], [23]. For cryptocurrencies, extremely heavy tails have been shown to be even more common, given their extreme volatility [4], [21],[22].
It will therefore be interesting to look at the behavior of the extreme returns of Bitcoin, Litecoin, Ripple and Ethereum, by examining the tails seperately, using Extreme Value Theory, or, more specifically, by fitting Generalized Pareto Distributions on the tails of the cryptocurrencies.

To summarize, in the current research, first and foremost, a risk analysis will be made for Bitcoin, Litecoin, Ripple and Ethereum, and will, (when possible), be compared to the analyses done in previous research described before [4], [21], [22]. It will be interesting to see how the shift in market share (as the altcoins have taken over a large part of the market share from Bitcoin since the results in the research papers just mentioned), has influenced the distributions of the returns as well as the Value at Risk and

the Expected Shortfall.
Also, more focus will be placed on profits compared to previous research, instead of solely focussing on the risks (losses).
Previous research was also done before the current highs of Bitcoin and the altcoins, which happened in the last few months of 2017. Briefly put, a lot has happened in the cryptocurrency world since this research took place, and it will therefore be interesting to see how distributions and risk measures have changed since then.
Finally, Ethereum was not included in most of the research described before, and given that it is the second largest cryptocurrency today, it will be interesting to get some results for Ethereum as well.

The aim of this research is to answer the following research question:

How do the returns of Bitcoin and other altcoins behave over time, and what can we say about extreme values for losses and profits?

This research paper is organised as following: all methods will be explained thoroughly in the Methodology (section 2), after which the results for each will be covered in the Results (section 3).
Here descriptive statistics are given in section 3.1, giving a clear overview of historical data retrieved from CoinMarketCap [3], for all four cryptocurrencies, together with summary statistics. Then, volatility is covered in section 3.2, showing and comparing volatility between the four cryptocurrencies.
Then, the Value at Risk and Expected Shortfall are introduced as quantities for risk analyses in section 3.3. These will first be calculated using historical data from the same dataset, and then, after distributions are fitted in section 3.4, be compared to the Value at Risk and Expected Shortfall based on these distributions.
This will not only be done for the losses, but also for the profits, after which a comparison will be made between them. Finally, the tails of the distributions will be covered seperately, and be tested for Paretianity in section 3.5. Also, a Maximum to Sum analysis is made and in order to confirm (or refute) the Paretianity results, in section 3.6.

Each subsection ends with a short conclusion, where the most important results are summarized and interpreted.

The research paper is concluded with a Conclusion (section 4), as well as a Discussion (section 5), reflecting on the research and providing ideas for future research.

2 Methodology

In this section, the methods that were used in order to obtain results and answer the research question

"How do the returns of Bitcoin and other altcoins behave over time, and what can we say about extreme values for losses and profits?"

are thoroughly explained, and consist of the following:

- First, descriptive statistics are given to provide an overview of the behavior of the exchange rates of the four cryptocurrencies of interest. Based on these exchange rates, information with respect to the returns can be obtained and analysed. Along with time series plots and histograms, summary statistics will be given and analysed for all four cryptocurrencies in order to form well founded hypotheses with respect to the behavior of the returns.
- Then, volatility distributions will be created in order to provide some insight with respect to how volatile each of the currencies are. This provides useful information with respect to the risk that is taken by investing in cryptocurrencies.
- Other very important terms in risk analysis are the Value at Risk and the Expected Shortfall that will be calculated and compared to previous research.
- Next, some common distributions will be fitted on the log returns, and compared to the empirical distributions of the cryptocurrencies.
- Also, the tails of the distributions will be tested for Paretianity, to discover whether the tails are truly as heavy as expected.
- Finally, the Paretianity results will be checked by creating Maximum to Sum ratio plots.

Each of these methods contributes in its own way to a better understanding of the behavior of the returns of Bitcoin, Litecoin, Ripple and Ethereum, and their tail behavior in specific.

2.1 Descriptive Statistics

Exchange Rates and Log Returns

In this section, descriptive statistics for Bitcoin, Litecoin, Ripple and Ethereum are presented, based on a time span of three years and four months for Bitcoin, Litecoin and Ripple (starting from January 1, 2015 up to April 24, 2018) , and based on a time span of two years and four months for Ethereum (starting from January 1, 2016 up to April 24, 2018), as Ethereum only started trading in Q3 of 2015, and the first few months of trading have not been taken into account (as in general, the first months after a new financial product is released, the exchange rate remains close to zero, hence there is little volatility to speak of, which could lead to a distorted image of pattern of the timeseries overall).

- First, the daily exchange rates (closing prices) are considered. These rates indicate the last known price of each day, according to CoinMarketCap [3]. Time series will be presented in the form of a plot of the exchange rates against time, with t_0 representing January 1, 2015 for Bitcoin, Litecoin and Ripple, and January 1, 2016 for Ethereum.
  This has resulted in $n = 1210$ values of exchange rates for Bitcoin, Litecoin and Ripple, and $n = 845$ values for Ethereum.

- From the daily exchange rates, the log returns can be extracted.
  The reason we will transform the returns into log returns, is because we want to make comparisons between the different cryptocurrencies, and we therefore want to know the relative

returns. (If, for example, the value of Bitcoin were to increase by $10 dollars, going from $7000,- to $7010,-, this would have a very different impact than an increase by $10,- for Ripple, going from $0.40 to $10.40).

The log returns are calculated as follows:

Let S_i be the exchange rate on day i, $i = 1, ...n$, then the log return R_i on day i is found by

$$R_i = \log\left(\frac{S_{i+1}}{S_i}\right) = \log(S_{i+1}) - \log(S_i) \approx \frac{S_{i+1} - S_i}{S_i}, \quad i = 1, ..., n-1.$$

They have also been calculated by considering time spans of weeks and months, instead of just days.
The weekly log returns have been calculated by taking averages over non-overlapping seven consecutive daily log returns. So the first value of the weekly log returns is the average over the first seven daily log returns, the second value of the weekly log returns is the average over values 8 up until 14 of the daily log returns, etc.
The monthly log returns have been calculated in a similar way.

Moving Averages

Also, several moving averages have been computed based on the daily exchange rates, to present a visualization with less 'noise', showing the overall trend more clearly. The $x-$MA denotes the Moving Average where the average has been taken over the past x days.
For Bitcoin, Litecoin, Ripple and Ethereum, moving averages have been computed for $x = 10, x = 30, x = 100$ and $x = 300$. Here, the starting point t_0 is once again January 1, 2015. for Bitcoin, Litecoin and Ripple, and January 1, 2016 for Ethereum.
For the first couple of days, the moving average based on the daily exchange rates has simply been calculated by averaging over the days up until the current day. As soon as more than x days have passed, the average is taken over only the last x days, thereby ignoring data from before x days ago.

Let S_i be the exchange rate on day i and $x-\text{MA}_i$ be the $x-$day Moving Average on day $i, i = 1, ...n$. Then the $x-\text{MA}_i$ is calculated as follows:

$$x - MA_i = \begin{cases} \frac{1}{i}\sum_{k=1}^{i} S_k & \text{if } i < x \\ \frac{1}{x}\sum_{k=i-x}^{i} S_k & \text{if } i \geq x \end{cases} \tag{1}$$

Summary Statistics

Summary statistics have been computed of both daily exchange rates as well as daily log returns, of Bitcoin, Litecoin, Ripple and Ethereum versus the US Dollar.
These statistics are based on the same time spans as before, i.e. 1210 days for Bitcoin, Litecoin and Ripple, and 845 days for Ethereum.

The statistics that have been computed, are the minimum, Q1, mean, Q3, maximum, skewness, kurtosis, standard deviation (sd), variance and Interquartile Range (IQR).
Whereas the meaning of the minimum, median, mean, maximum, standard deviation and variance are rather straight forward, the other statistics may require some more explanation: the first quantile value of the exchange rates (or log returns), or Q1, is the value such that 25% of observations are below this value, and 75% greater than or equal to this value. Likewise, Q3 represents the third quantile value over all days considered. This is the value such that 75% of observations are below this value, and 25% greater than or equal to this value. Based on these two values is the Interquartile Range (IQR), a measure of dispersion, indicating the 'middle 50%' of the data. It is calculated by simply taking the difference between Q3 and Q1.

Then, the skewness is a measure of the asymmetry of the distribution of the exchange rate (or log return) values about their means.
If the data is distributed in a way that there is only one mode (or peak), a negative skew indicates that the tail on the left side of the probability density function is longer or fatter than the right side. Conversely, a positive skew indicates that the tail on the right side is longer or fatter than the left side. On the other hand, the kurtosis is a measure of the tailedness of the distribution. It describes the shape of the distribution, similar to the skewness, except that from the kurtosis value, the presence of possible asymmetry cannot be derived. The kurtosis value is often compared to the value 3, belonging to the kurtosis value for a Normal Distribution. It is common to compare the kurtosis of a distribution to this value. A value below 3 means the distribution produces fewer and less extreme outliers than does the Normal Distribution, whereas a higher value implies more (extreme) outliers.
Again, when a high value for kurtosis is obtained, it is unclear from this value alone whether these extreme values occur mostly on the left, on the right, or on both sides of the distribution.

2.2 Volatility

A cryptocurrency's volatility is the variation in its value over a period of time. For example, two different cryptocurrencies might end up at the same value after a while, but perhaps the first one had a tendency to swing wildly to higher and lower values, whereas the second may have moved in a much steadier, less turbulent way. Both cryptocurrencies may end up at the same exchange rate at the end of the day, but their path to that point could have been very different.

A variable's volatility, σ, is defined as the standard deviation of the log returns of a variable [11].
This standard deviation was also included in the summary statistics described in the section before, but what if the volatility varies throughout time and is not represented well by a fixed parameter? In fact, previous research by [21], also described in the Introduction, suggests the volatility itself is very unstable. It therefore makes sense to treat volatility as a time dependent variable, and examine the distribution of the volatilities.
In order to do this, a window of a certain width must be chosen over which the volatilities can be calculated. Chan et. al.[4] used a window of width 20 days, and we will do the same.
This means that, for each day, the volatility is simply the standard deviation over the past 20 days.
Based on the 20 days volatility, a histogram will be given for the distribution of the volatilities, along with the common measures of central tendency: the mean, median and mode, for the annualized volatility (which is obtained by multiplying the 20 days rolling volatility by $\sqrt{365}$ (here, the mode will be obtained by first rounding the log returns to three decimals). The authors in [22] found an annualized volatility of 77% for Bitcoin. It will be interesting to see whether this has changed much, and how the other three cryptocurrencies compare to this.

2.3 Value at Risk and Expected Shortfall

In risk analysis, two very important quantities are Value at Risk (VaR) and Expected Shortfall (ES) [12],[13]. The former is a measure to make a statement with respect to the confidence level(α) that we will not lose more than V dollars in time T. It is the loss level during a time period of length T that we are $X\%$ certain will not be exceeded [11].
The latter is a measure to make a statement about the expected loss when you know a loss will be incurred. Whereas VaR asks the question:"How bad can things get?" ES asks: "If things do get bad, what is the expected loss?"
Considering the expectation that the log returns of cryptocurrencies show extreme behavior in the tails, VaR alone does not provide sufficient information, and it is especially important to also consider the ES. Both VaR and ES are calculated based on the daily, weekly and monthly log returns.
Note however, that this will not result in a monetary value for the VaR or ES, but a percentage for the VaR and ES instead, making it easier to make comparisons between the cryptocurrencies of interest.

Other than losses, it can also be of interest to consider potential profits. Statements can be made with respect to profits based on the VaR and ES as well, this time considering the 'other tail' of the distribution.

When the theoretical distribution of the log returns is known (for example we know data is normally distributed), The 5% VaR with respect to losses ($\alpha = 0.05$) is simply the value corresponding to the 5th percentile of the distribution of the log returns (assuming negative values denote losses and positive values denote profits: otherwise the 95th percentile needs to be calculated instead).
For the profits, this is then of course the other way around (in general, one should also keep in mind that, whereas the left side of the distribution may sometimes be losses, it doesn't have to be, as some instruments are so profitable that the left 5% is still in the green and therefore the left side refers to profits too).

Mathematically speaking, assuming we are interested in losses, this means that, for F a fitting cumulative distribution function, and α the percentile of interest ($0 < \alpha < 1$), that

$$\text{VaR}(\alpha) = F^{-1}(\alpha) \tag{2}$$

Without a theoretical distribution, it is also possible to calculate VaR, simply by sorting the data, and taking the 5th percentile afterwards.

For the ES the same principle applies, except this time an average will be taken over all values between the 0th and the 5th percentile, and so, for general α,

$$\text{ES}(\alpha) = \frac{1}{\alpha} \int_0^\alpha \text{VaR}(x)\ dx \tag{3}$$

Analyses will be done for all cryptocurrencies of interest with respect to their VaR and ES based on the daily, weekly and monthly log returns. These will then be compared to the theoretical VaR and ES of the best fitting theoretical distribution (more information about fitting of distributions can be found in the next section).
VaR and ES will be calculated both for losses as well as profits. For losses, these values will give a measure of risk, but for profits, they are of course indicative of potential gains.
In both cases, however, the entire distributions are used for computing the VaR and the ES (not the semi-distributions for the losses or profits seperately).

It will be interesting to compare the results to previous research done by Osterrieder et.al [21], [22], who found, for the losses, values around 10% for the Value at Risk, and around 20% for the Expected Shortfall.

2.4 Fitting of distributions

Based on histograms of the log returns, as well as the summary statistics provided, it is possible to make hypotheses about underlying distributions.
As mentioned in the introduction, Chan et. al. [4] have also tried fitting distributions to several cryptocurrencies, and found that some well fitting distributions were the Generalized Hyperbolic Distribution, the Normal Inverse Gaussian Distribution and the Laplace Distribution.
These distributions will also be fitted to the cryptocurrencies of interest in this research, together with the Normal Distribution and the Student's t Distribution.

In order to fit the distributions found below, the method of Maximum Likelihood is used to obtain maximum likelihood estimates (MLE) for the parameters, based on the data at hand.
Let $\underline{\theta}$ be a vector of parameters for a certain distribution.
Then the optimal value for $\underline{\theta}$ is found by maximizing the likelihood function $\mathcal{L}(x|\underline{\theta})$, therefore solving

$$\frac{\partial \mathcal{L}(x|\underline{\theta})}{\partial \underline{\theta}} = 0$$

i.e. solving a system of equations.

While the maximum likelihood estimates are easily obtained analytically for some distributions, for others this is not so straightforward, as it involves solving complex systems of equations.
Therefore, several `R` packages are used in order to obtain the maximum likelihood estimates.
Below, an overview of the fitted distributions can be found, together with information as to how the maximum likelihood estimates for the parameters are obtained.

1. Normal Distribution
   The density for the Normal Distribution, for $X \sim \mathcal{N}(\mu, \sigma)$, is given by
   $$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \tag{4}$$
   Parameters μ and σ are the mean (and also the median and the mode), and the standard deviation respectively. Therefore, the Normal Distribution is defined only by a location and a scale parameter.
   The Normal Distribution is one of the most important distributions in statistics and will therefore be fitted to the data, even though the expectation is that this distribution will not be a good fit at all, as explained before.
   In order to find the maximum likelihood estimates, $\frac{\partial \mathcal{L}(x|\mu,\sigma)}{\partial \mu} = 0$ and $\frac{\partial \mathcal{L}(x|\mu,\sigma)}{\partial \sigma} = 0$ are solved to obtain the following estimates:
   $$\begin{aligned} \widehat{\mu} &= \frac{1}{n}\sum_{i=1}^{n} x_i \\ \widehat{\sigma}^2 &= \frac{1}{n}\sum_{i=1}^{n}(x_i - \widehat{\mu})^2 \end{aligned}$$

2. Student's t Distribution
   The density for the Student's t Distribution [9], that is $X \sim t(\mu, \sigma, \nu)$ is given by
   $$f(x) = \frac{\sqrt{\nu}B(\frac{1}{2}, \frac{\nu}{2})}{\sigma}\left(1 + \frac{(x-\mu)^2}{\sigma^2\nu}\right)^{-\frac{\nu+1}{2}} \tag{5}$$
   for $-\infty < x < \infty, -\infty < \mu < \infty, \sigma > 0$ and $\nu > 0$, where $B(\cdot,\cdot)$ denotes the beta function defined by
   $$B(a,b) = \int_0^1 t^{a-1}(1-t)^{b-1}\ dt.$$
   Here, the parameters μ and σ are location and scale parameters respectively (mean and standard deviation), and ν is the number of degrees of freedom.
   The shape of the probability density function resembles that of the Normal Distribution, except that it is a bit lower and wider (heavier tails), depending on the parameter ν..
   The expectancy is therefore that the Student's t Distribution will fit the data better than the Normal Distribution.
   Maximum likelihood estimates for the Student's t Distribution are obtained using the `R` function `MASS::fitdistr` from the `MASS` package.

   The Student's t Distribution with ν degrees of freedom, is a special case of the Generalized Hyperbolic Distribution (discussed later), with $\lambda = -\frac{\nu}{2}, \alpha = 0, \beta = 0, \delta = \sqrt{\nu}$.

3. Laplace Distribution
   The density for the Laplace Distribution [16], that is for $X \sim \text{Laplace}(\mu, \sigma)$ is given by
   $$f(x) = \frac{1}{2\sigma}\exp\left(-\frac{|x-\mu|}{\sigma}\right) \tag{6}$$

and can also be written as

$$f(x) = \begin{cases} \frac{1}{2\sigma} \exp\left(\frac{\mu - x}{\sigma}\right) & \text{if } x < \mu \\ \frac{1}{2\sigma} \exp\left(\frac{x-\mu}{\sigma}\right) & \text{if } x \geq \mu \end{cases} \tag{7}$$

for $-\infty < x < \infty, -\infty < \mu < \infty, \sigma > 0$.

The Laplace Distribution is basically a combination of two Exponential Distributions, spliced together at the location $x = \mu$, the median of the distribution. σ is again a scale parameter, the absolute standard deviation from the median [20]. Maximum likelihood estimates for the Laplace Distribution are therefore

$$\begin{aligned} \widehat{\mu} &= \text{med}(x_1, ...x_n) \\ \widehat{\sigma} &= \frac{1}{n}\sum_{i=1}^{n} |x_i - \widehat{\mu}| \end{aligned}$$

The Laplace Distribution is the only one out the list of distributions here, that is not heavy tailed.

4. Normal Inverse Gaussian Distribution
   The density for the Normal Inverse Gaussian (NIG) Distribution [1], that is for $X \sim \text{NIG}(\alpha, \beta, \delta, \mu)$, is given by

$$f(x) = \frac{(\gamma/\delta)^{-1/2}\alpha}{\sqrt{2\pi}K_{-1/2}(\delta\gamma)} \exp[\beta(x-\mu)] \left[\delta^2 + (x-\mu)^2\right]^{-1} K_{-1}\left(\alpha\sqrt{\delta^2 + (x-\mu)^2}\right) \tag{8}$$

   for $-\infty < x < \infty, -\infty < \mu < \infty, \delta > 0, \alpha > 0$ and $\beta > 0$, where $\gamma = \sqrt{\alpha^2 - \beta^2}$.

   Here, $K_\nu(\cdot)$ denotes the modified Bessel function of the second kind of order ν defined by

$$K_\nu(x) = \begin{cases} \frac{\pi xsc(\pi\nu)}{2}[1_{-\nu}(x) - I_\nu(x)], & \text{if } \nu \notin \mathbb{Z}, \\ \lim\limits_{\mu \to \nu} K_\mu(x), & \text{if } \nu \in \mathbb{Z} \end{cases}$$

   where $I_\nu(\cdot)$ denotes the modified Bessel function of the first kind of order ν defined by

$$I_\nu(x) = \sum_{k=0}^{\infty} \frac{1}{\Gamma(k+\nu+1)k!}\left(\frac{x}{2}\right)^{2k+\nu},$$

   where $\Gamma(\cdot)$ denotes the gamma function defined by

$$\Gamma(a) = \int_0^\infty t^{a-1}\exp(-t)\ dt$$

   The parameters for the Normal Inverse Gaussian are a location parameter μ, a parameter α for the tail heaviness, an asymmetry parameter β, and a scale parameter δ.
   The expectation is that the Normal Inverse Gaussian will fit the data well, because it has a tail heaviness and an asymmetry component, which we expect to be present in our data as well.
   Also, Chan et. al. [4] found that this distribution was the best fit for Ripple.
   Maximum likelihood estimates for the Normal Inverse Gaussian Distribution are obtained using the `R` function `GeneralizedHyperbolic::nigFit` from the `GeneralizedHyperbolic` package.

   The Normal Inverse Gaussian Distribution is a special case of the Generalized Hyperbolic Distribution (discussed next), with $\lambda = -\frac{1}{2}$.

5. Generalized Hyperbolic Distribution
   The density for the Generalized Hyperbolic Distribution [1], that is for $X \sim \mathrm{GH}(\lambda, \alpha, \beta, \delta, \mu)$, is given by

$$f(x) = \frac{(\gamma/\delta)^{\lambda} \alpha^{1/2-\lambda}}{\sqrt{2\pi} K_{\lambda}(\delta\gamma)} \exp[\beta(x-\mu)] \left[\delta^2 + (x-\mu)^2\right]^{\lambda-1/2} K_{\lambda-1/2}\left(\alpha\sqrt{\delta^2+(x-\mu)^2}\right) \qquad (9)$$

   for $-\infty < x < \infty, -\infty < \mu < \infty, -\infty < \lambda < \infty, \delta > 0, \alpha > 0$ and $\beta > 0$, where $\gamma = \sqrt{\alpha^2 - \beta^2}$. The parameters for the Generalized Hyperbolic Distribution are a location parameter μ, a scale parameter δ, two shape parameters α and λ, and a skewness parameter β.

   Maximum likelihood estimates for the Generalized Hyperbolic Distribution are obtained using the `R` function `ghyp::fit.ghypuv` from the `ghyp` package, providing parameter estimates for $\mu, \delta, \alpha, \beta$ and λ .
   The Generalized Hyperbolic Distribution is often used in economics, with particular application in the fields of modelling financial markets and risk management, due to its semi-heavy tails.
   We expect the Generalized Hyperbolic Distribution to fit our data well. Also, Chan et. al. [4] found that the Generalized Hyperbolic Distribution was the best fit for Bitoin and Litecoin.

In order to find out whether the distribution of interest is a good fit for the dataset, both visualizations (plots of the histograms of log returns together with the corresponding densities, and QQ-plots will be given), as well as the Kolmogorov-Smirnov [15] statistic as a goodness-of-fit test, is used.

Visual explanation

First, visualizations will be used to see whether the distributions to be tested may be good fits. In order to obtain useful visualizations, histograms of the log returns will be plotted (like the ones in the Descriptive Statistics section).
This time, however, using the maximum likelihood parameters (as explained before), densities will be created for each of the distributions, and will be plotted on top of the histograms, thereby providing intuition whether the distribution of interest fits the data well or not.

Also, QQ-plots will be created. A QQ-plot is a Quantile-Comparison plot, which is a very useful tool for investigating whether a sample is from a theoretical distribution [28].
It compares the quantiles of the empirical (cumulative) distribution function

$$\widehat{P} = \frac{\text{Number of data points } \leq x}{n}$$

to the quantiles of one of the theoretical (also cumulative) distribution functions:

$$P(x) = \mathbb{P}(X \leq x)$$

If the sample is generated from the distribution of interest, the plot will be approximately linear with intercept 0 and slope 1.

Histograms and QQ-plots will therefore be used to provide some insight into what distributions might be good fits, but visualizations alone are not enough evidence to base any conclusions on. Therefore, the Kolmogorov-Smirnov statistic will be used to confirm or refute the expectations based on the visualizations.

Kolmogorov-Smirnov Goodness of Fit statistic

Where the visualizations give an intuitive idea as to whether a distribution fits the data well, the KS statistic can be used to see whether the 'match' is significant or not.
Let $\widehat{F}$ be the maximum likelihood estimate of the distribution of interest $F(x)$, and $\mathbb{1}$ the indicator

function, then the Kolmogorov-Smirnov statistic is given by

$$\mathrm{KS} = \sup_x \left| \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{\{x_i \leq x\}} - \widehat{F}(x) \right| \tag{10}$$

What this basically comes down to, is that the empirical and theoretical distributions are compared, and the length of the maximal distance between the two distributions is used as a statistic that will then be compared to a table of thresholds for statistical significance.

In this research, the Kolmogorov-Smirnov statistic was calculated by comparing the empirical distribution of the log returns to one of the aforementioned theoretical distributions, that was created by using the maximum likelihood parameters for the distribution at hand and generating random samples.
For the Normal, Laplace, NIG and Generalized Hyperbolic Distributions, 1000 random samples were drawn, resulting in 1000 $p-$values indicating whether the null hypothesis is to be rejected or not. For the student's t distribution, only 100 random samples were drawn for computational intensity reasons.

The null-hypothesis, in this case, states that the empirical and randomly drawn theoretical values come from the same distribution, and so a $p-$value smaller than 0.05 indicates that there is sufficient evidence to reject the null hypothesis on a critical value corresponding to the 95th percentile of the corresponding distribution, therefore indicating that the distributions are not the same, implying that the data does not likely come from the distribution of investigation.

For each of the distributions, the mean and standard deviation of the $p-$values for the KS statistic will be stated, together with a percentage of $p-$values that imply the values come from the same distribution.
Based on this information, a statement with respect to the best fitting distributions will be made for each of the currencies.

Chi-Square(χ^2) Goodness of Fit statistic

Another goodness of fit test that will be used, is the Chi-Square goodness of fit test [26].
This goodness of fit test is used to find out how the observed value is significantly different from the expected value, by making use of so called 'bins.' A theoretical distribution is compared to the empirical distribution, by dividing the sample data into intervals (where the intervals are not of equal length, but each interval consists of an equal number of values). The number of values that fall into the interval or bin when considering the theoretical distribution of interest, is then compared to the number of values that fall into the same interval in the sample data, for each bin, therefore comparing two vectors of the length of the number of bins.
The null-hypothesis is that the distributions are the same.
When using a cut-off value of $\alpha = 0.05$, and the frequencies in the two vectors are not similar enough so that the Chi-Square test will return a $p-$value smaller than α, the distributions are not likely to be the same.

In case of Bitcoin, where the log returns consist of 1209 values, the boundaries will be determined based on the number of bins that will be used (we will use 5 and 10 bins respectively, and compare the results).
When using 10 bins, for example, the boundaries for the first bin will be determined based on values of the first approximately 120 log returns (when using the sorted log returns, of course).
For the remaining 9 bins, the boundaries will be determined in a similar way. Afterwards, these boundaries will be used as quantiles for the theoretical distributions that are tested, after which the Chi-Square statistic will be calculated and reported.

Empirical vs Theoretical Value at Risk and Expected Shortfall

Finally, the best fitting distributions will be used to calculate the Value at Risk and Expected Short-

fall (therefore calculating a theoretical VaR and ES), and be compared to the empirical (daily) VaR and ES, to get an idea with respect to how well the best fitting distributions capture the extreme values of the log returns as well.
For the comparisons, the entire empirical distributions are used and both the left side of the distributions (the losses) as well as the right side (the profits) will be compared to the theoretical distributions.

2.5 Testing for Paretianity

Given that empirical studies have established that the distribution of financial returns tend to have heavier tails than for example the Gaussian distribution tails [23], it is interesting to examine these tails separately.
Using extreme value theory, we want to give a statistical characterization of the tail properties of the log returns for Bitcoin, Litecoin, Ripple and Ethereum. In this section, the tails of the distributions of the log returns of the daily exchange rates will be tested for Paretianity, that is, a Pareto Distribution, first introduced by Vilfredo Pareto [24], will be fitted on both of the tails of the distributions, one at a time. Fitting Pareto Distributions on the data will provide information with respect to, for example, the heaviness of the tail.
A random variable X is said to follow a Pareto Distribution if its density function $f(x)$ is such that

$$f(x) = \frac{\rho x_0^\rho}{x^{\rho+1}}, \qquad 0 < x_0 \leq x, \tag{11}$$

where ρ is the shape parameter, which measures the heaviness of the tail, and x_0 is a scale parameter [5].

The scale parameter x_0 can be interpreted as some sort of cut-off value (or a minimum) after which the data is assumed to be Pareto distributed (so for $x > x_0$, a Pareto Distribution might hold), but more important is the shape parameter, since the behavior of the distribution is determined by it and interpretations with respect to tail characteristics such as the heaviness of the tail, are based on it.
When fitting a Pareto Distribution on the tail, the following holds: the smaller the shape parameter ρ , the fatter the tail.
Equation 11 comes with a corresponding cumulative distribution function, which is given by

$$F(x) = \mathbb{P}(X \leq x) = 1 - \left(\frac{x}{x_0}\right)^{-\rho}, \qquad 0 < x_0 \leq x \tag{12}$$

The distribution just described is known as the pareto I, but many generalizations have been proposed, among which is also the Generalized Pareto Distribution (GPD).
For the Generalized Pareto Distribution, we have

$$F(x) = \begin{cases} 1 - \left(1 + \frac{\xi(x-\nu)}{\beta}\right)^{-\frac{1}{\xi}} & \xi \neq 0 \\ 1 - \exp\left(-\frac{x-\nu}{\beta}\right) & \xi = 0 \end{cases} \tag{13}$$

where $x \geq \nu$ for $\xi \geq 0, \nu \leq x \leq \nu - \frac{\beta}{\xi}$ for $\xi < 0, \xi \in \mathbb{R}$ and $\sigma > 0$.

Note that the Pareto I Distribution is in fact the Generalized Pareto Distribution for $\frac{1}{\xi} = \rho$ (when $\xi > 0), \nu = \beta\rho$ and $\beta = \frac{x_0}{\rho}$.

In this research, the Generalized Pareto Distribution found in equation 13 will be fitted to the tails (both to the losses and the profits) of the returns for Bitcoin, Litecoin, Ripple and Ethereum, and we will therefore base our conclusions with respect to the heaviness of the tails, on the shape parameters ξ belonging to this distribution.

It is important that the threshold β is chosen high enough for the exceeding values to be well approximated by the GPD, but not so high to substantially increase the variance of the estimator due to

reduction in the sample size (the number of exceedances).
Values for β and ξ are obtained using the `evir::gpd` function from the `evir` package in `R`.
Here, we are interested to find the smallest value for $k \in \mathbb{N}$ such that $\xi \geq \frac{1}{k}$, as this implies an infinite $k-$th moment (and then, automatically infinite l-th moments for all values $l > k$) and will therefore provide information with respect to the tail's heaviness. To summarize, the following holds: the smaller the shape parameter ξ is, the heavier the tail.

In order to visualize the data, the plots that will be given will be based on the survival function (in general, a survival function is a function that gives the probability that a patient, device, or other object of interest will survive beyond any given specified time [14]), and therefore the complement of the cumulative distribution function will be taken instead:

$$\overline{F}(x) = 1 - F(x) = \mathbb{P}(X > x) \tag{14}$$

This survival function is used in the so called Zipf plots introduced in [5], that will be given for the log returns of the cryptocurrencies, which plots the logarithm of this function against the logarithm of the log returns.

Note that, in order to obtain the logarithm of the log returns, when we are interested in potential losses, we need to multiply the log returns by -1 in order to get positive values denote the losses, and examine the tail.

First, Zipf plots will be made for all the losses.
Based on these plots, the values for β can be guessed and compared to the ones found by `R`, as this will be the value where the graph starts to resemble the one of a negative, linear, straight line. Then, new Zipf plots will be created based on the scale parameters obtained, to again provide a visualization of whether the tails are really Pareto distributed.
All methods described above will be performed for the losses and the profits seperately.

2.6 Maximum to Sum

Another very simple tool for detecting heavy tails of a distribution and obtaining a rough estimate of the order of the finite moments of a distribution, is by looking at a Maximum to Sum ratio [7]. In this section, Maximum to Sum plots for the Bitcoin, Litecoin, Ripple and Ethereum log returns will be given.
Let X_i denote the log returns of one of the cryptocurrencies on day i.
Then it is interesting to examine the ratio of the partial (cumulative) maximum value to the partial (cumulative) sum of X:

$$R_n(p) = \frac{M_n(p)}{S_n(p)} \tag{15}$$

where $M_n(p)$ denotes the maximum absolute value, to the power p, up until day n:

$$M_n(p) = \max\{|X_i|^p : 1 \leq i \leq n\} \tag{16}$$

and

$$S_n(p) = \sum_{i=1}^{n} |X_i|^p \tag{17}$$

and see whether the Maximum to Sum ratio converges to 0.
The reason why we are interested in the convergence, has to do with the following equivalence:

$$\frac{M_n(p)}{S_n(p)} \overset{\text{a.s.}}{\to} 0 \Leftrightarrow \mathbb{E}\left(|X|^p\right) < \infty \tag{18}$$

and therefore, the convergence provides us with information with respect to the $p-$th moment of the distribution (R_n provides us with information with respect to the mean of the log returns, $R_n(2)$ with respect to the variance, $R_n(3)$ with respect to the skewness and $R_n(4)$ with respect to the kurtosis).

Equivalence 18 will be used to make statements with respect to the order of the moments based on four plots that will be given for $R_n(p), p \in \{1, 2, 3, 4\}$ against n.
When we observe significant deviations from 0 for large n, this will be indicative of $\mathbb{E}|X|^p$ being infinite. Equations 16 to 18 use the absolute values of the log returns, therefore not distinguishing between losses and profits, but when we look only at the losses or the profits, similar plots can of course be obtained, providing us information with respect to the tails for the losses and profits seperately. If the log returns plots are inconclusive or if there are other reasons to believe it would be interesting to examine the losses and profits seperately, similar plots will be provided for the losses as well as for the profits. In this case, to keep the total number of days constant, when examining the losses, all profits will be set to 0, and when examing the profits, all losses will be set to 0 instead.

3 Results

3.1 Descriptive Statistics

3.1.1 Exchange Rates

A plot of the time series of the Bitcoin exchange rates starting from January 2015 up until the end of April 2018 can be found in figure 1.

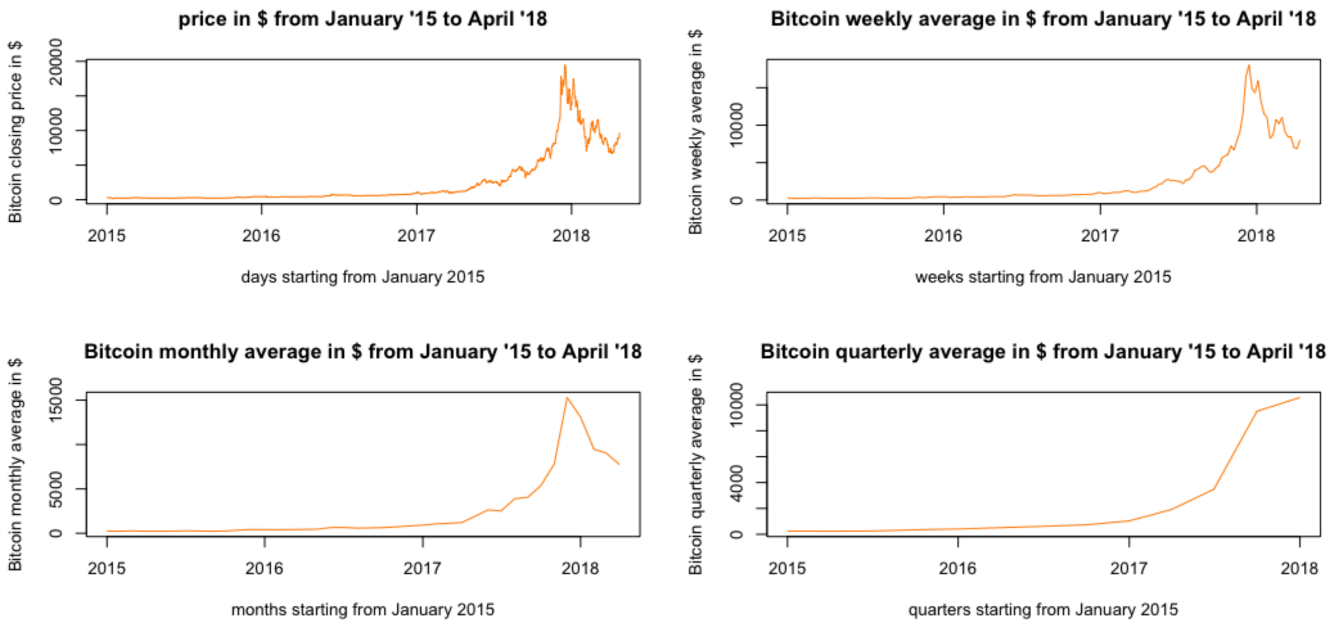

Figure 1: Time series daily Bitcoin Exchange Rates

From figure 1 it is apparent that there is an overall upward trend for the exchange rates. This trend appears to be exponentially increasing until the market experienced what can be denoted as a 'market crash' in January 2018. This is when the value dropped dramatically after a high increase happened during the months November and December 2017.
However, overall we can still speak of an increasing trend, which becomes especially clear by looking at the lower right plot showing the quarterly average exchange rates.
Finally, several moving averages have been computed, more specifically, a 10, 30, 100 and 300 day moving average respectively (from now on referred to as the $x-$MA, with $x = 10, x = 30, x = 100, x = 300$.)

Figure 2 shows a plot of the four moving averages. From these plots, we can see a similar pattern as the time series in figure 1: there is an upward trend apparent, where the sudden drop in value (in January 2018) immediately influences the 10 and 30 day moving averages, whereas for the 100 day moving average, we notice a delay in the drop (as can be expected, of course). The 300 day moving average only appears to be affected by the drop slightly, in that the increase appears to be dampened slightly.

For Litecoin, similar plots are made as was done for Bitcoin.
A plot of the time series of the Litecoin exchange rates, together with Moving Averages starting from January 2015 up until the end of April 2018 can be found in figure 3.

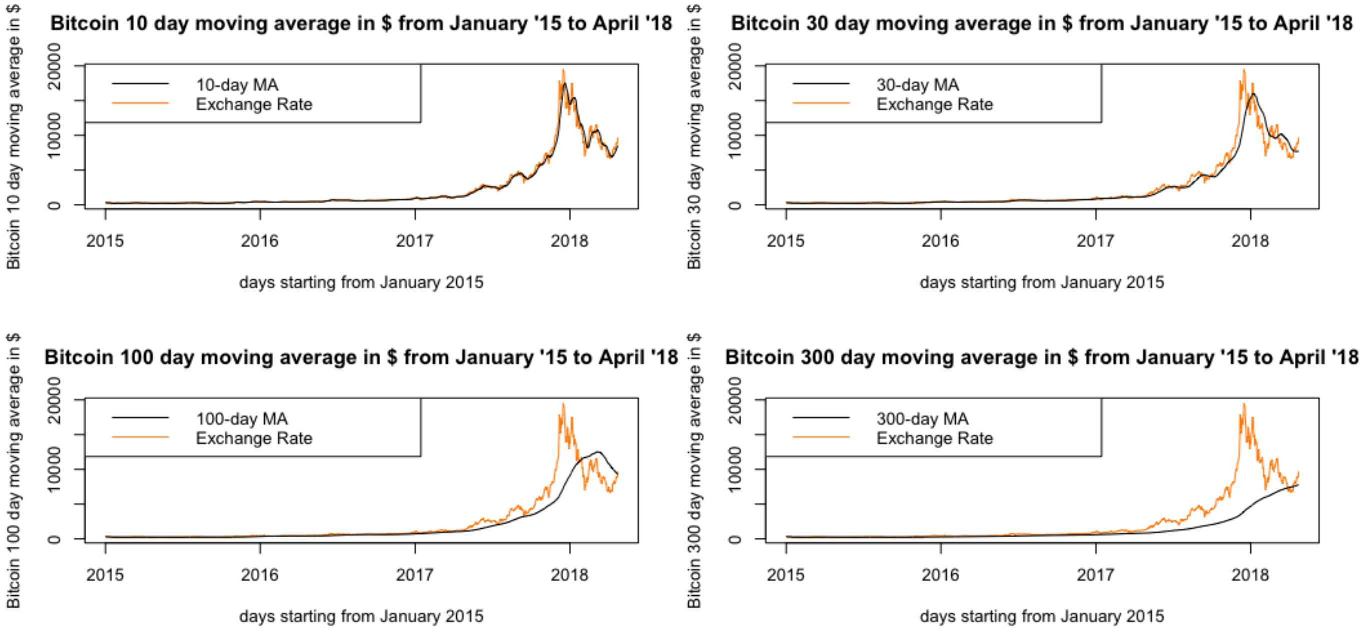

Figure 2: Bitcoin Moving Averages of Daily Exchange Rates

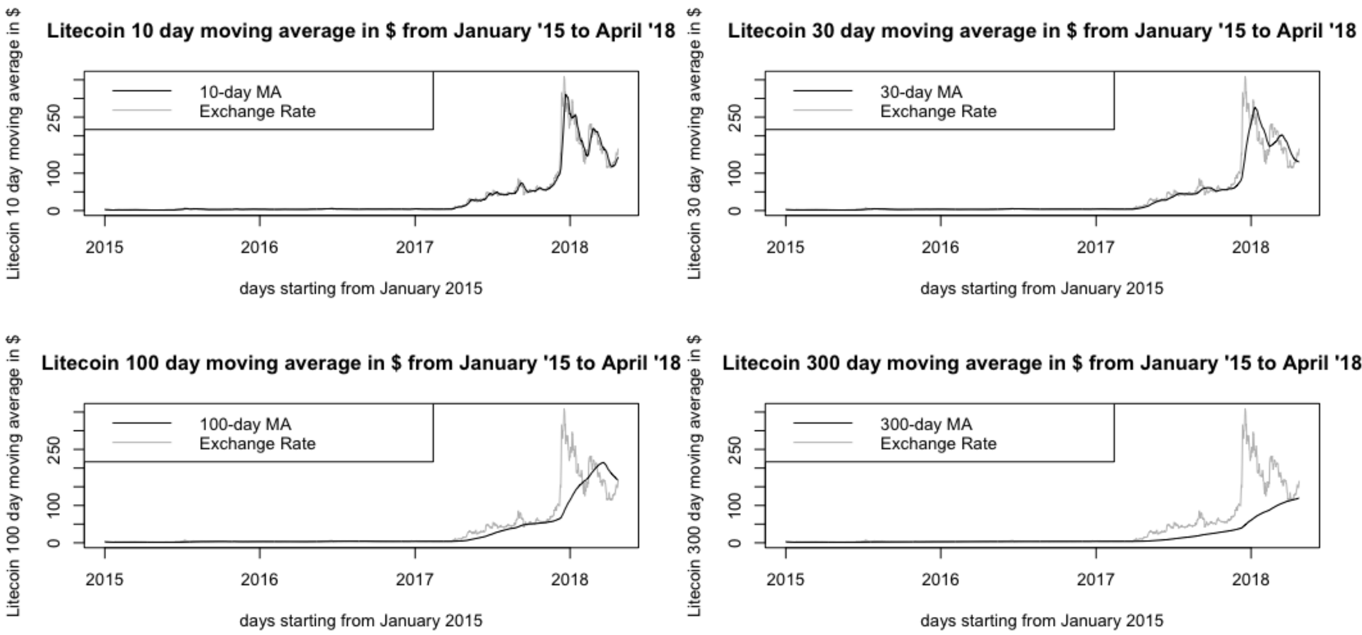

Figure 3: Time series and Moving Average daily Litecoin Exchange Rates

The plots in figure 3 show a similar trend as was observed for Bitcoin, except that the sudden increase around November and December 2017 appears to be even steeper.
Considering the fact that Litecoin is often traded in Bitcoin as opposed to US Dollar directly, signifying that an increase (or decrease) in the value for Bitcoin will automatically translate to an increase (or decrease) in Litecoin as well, and the fact that for cryptocurrencies, overall market increases and decreases are not uncommon, a similar to Bitcoin overall pattern of the exchange rate was to be expected.

Also, descriptive statistics for Ripple are presented, similar to the ones for Bitcoin and Litecoin.
A plot of the time series of the Ripple exchange rates together with Moving Averages, starting from January 2015 up until the end of April 2018 can be found in figure 4.

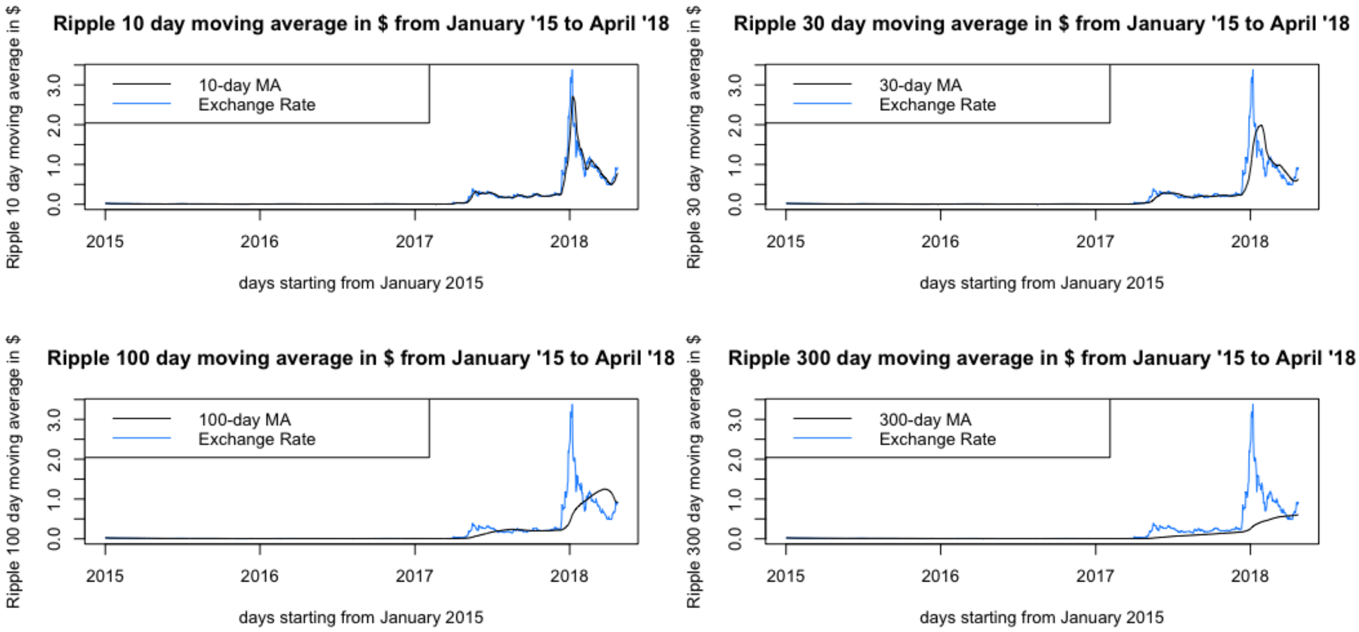

Figure 4: Time series and Moving Average daily Ripple Exchange Rates

The time series presented in figure 4 show similar patterns as the ones for Bitcoin and especially Litecoin. The same reasoning with respect to the Bitcoin trading pairs, as for Litecoin applies here.
As for Litecoin, here we also notice a very steep increase around November and December 2017.

Finally, descriptive statistics for Ethereum are presented. Since Ethereum started trading much later than Bitcoin, Litecoin, and Ripple, data has been collected for a shorter timespan, two years and four months to be exact, starting from January 2016, up until April 2018).

A plot of the time series of the Ethereum exchange rates together with Moving Averages, starting from January 2016 up until the end of April 2018, can be found in figure 5.

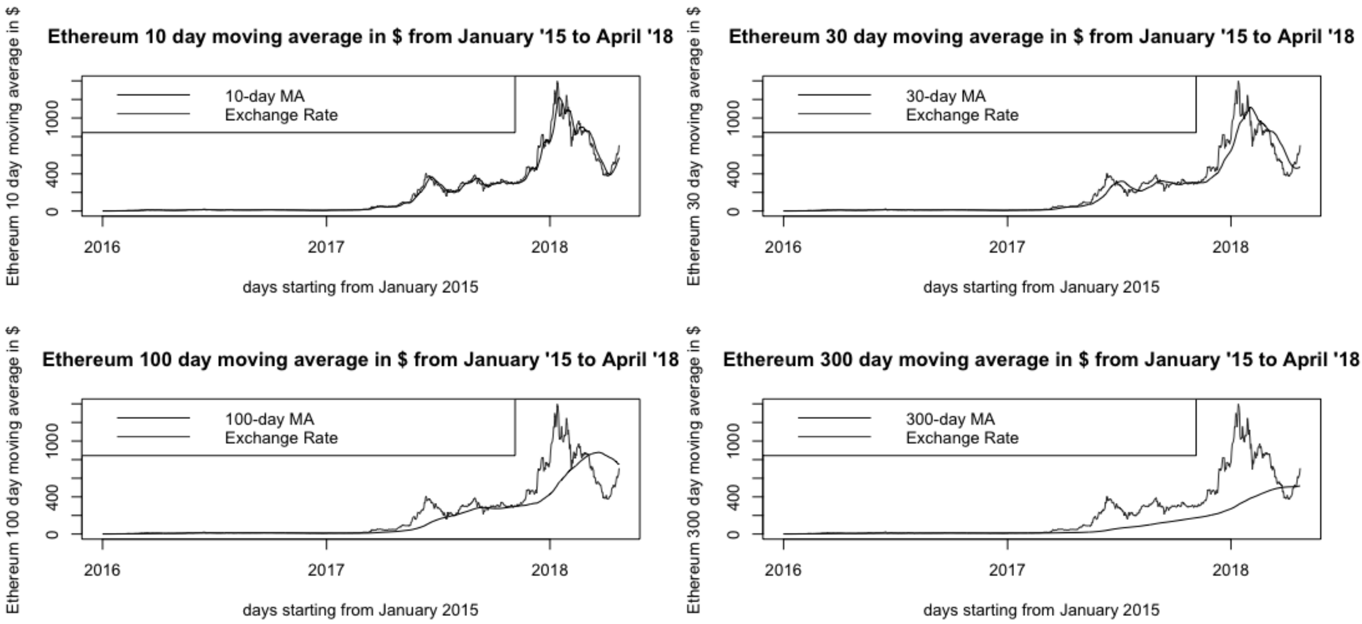

Figure 5: Time series and Moving Average daily Ethereum Exchange Rates

The time series plots for Ethereum look different from the ones for Bitcoin, Litecoin and Ripple. While the overall trend is still increasing, with a sudden exponential increase around the same time as for the other currencies, the increase overall appears to be much steadier compared to Litecoin and Ripple, and even somewhat steadier than Bitcoin.
The comparisons need to made with care, seeing as exchange rates have been taken into account starting in January 2016, as opposed to January 2015 as was the case for the other three currencies.

Short Conclusion

From the time series and Moving Averages plots, it appears that all four cryptocurrencies show an overall increasing trend: they appear to be moving in the same direction.
More detailed analyses with respect to this subject, is beyond the scope of this research, however. In future research, correlations between the cryptocurrencies could be examined, and Time Series Analysis could be done more explicitly, as this would have consequences in terms of risk diversification.

3.1.2 Log returns

A histogram of the daily Bitcoin log returns starting from January 2015 up until the end of April 2018 can be found in figure 6. Here, positive returns denote profits, whereas negative returns denote losses.

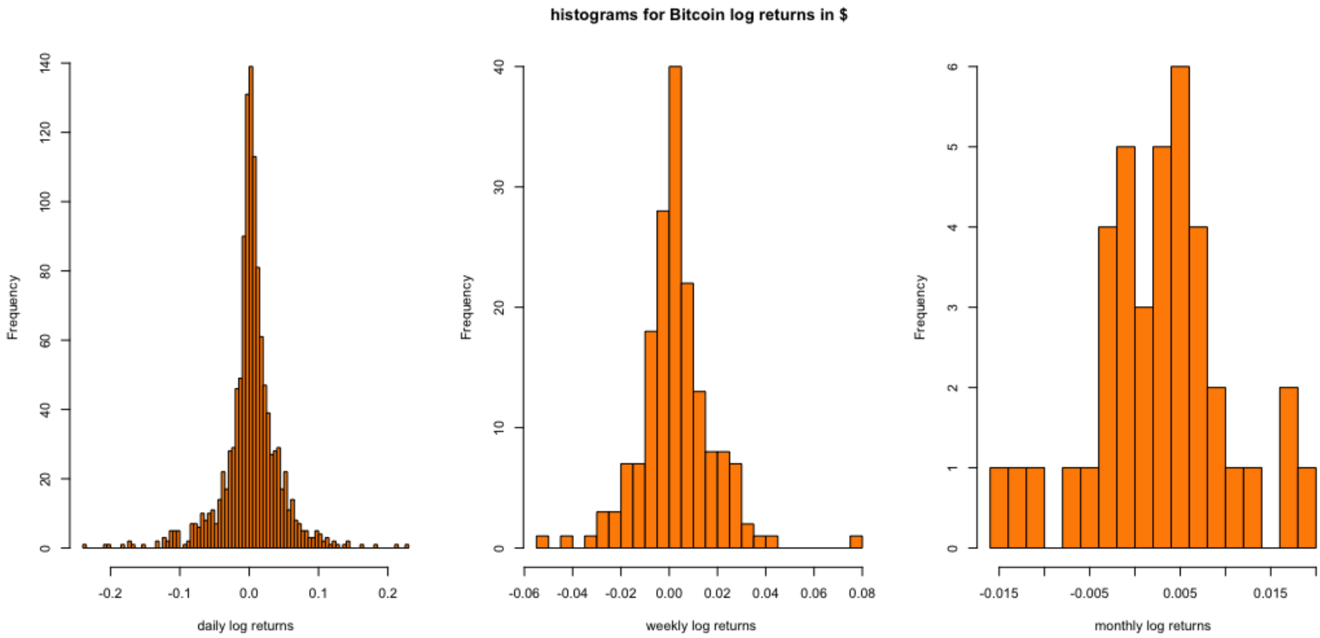

Figure 6: Bitcoin Log Returns

From the histograms in figure 6 the distribution appears to be somewhat symmetrical around 0: profits and losses appear to be evenly distributed. From these histograms there does not appear to be either a clear negative or positive skew present (summary statistics will be given later, from which more reliable conclusions can be drawn).
What also catches the eye, is that the monthly log returns show more than one peak (a so called bimodal distribution, where two local maxima occur).

Now, for Litecoin, a histogram of the daily Litecoin log returns starting from January 2015 up until the end of April 2018 can be found in figure 7. Here, positive returns denote profits, whereas negative returns denote losses.

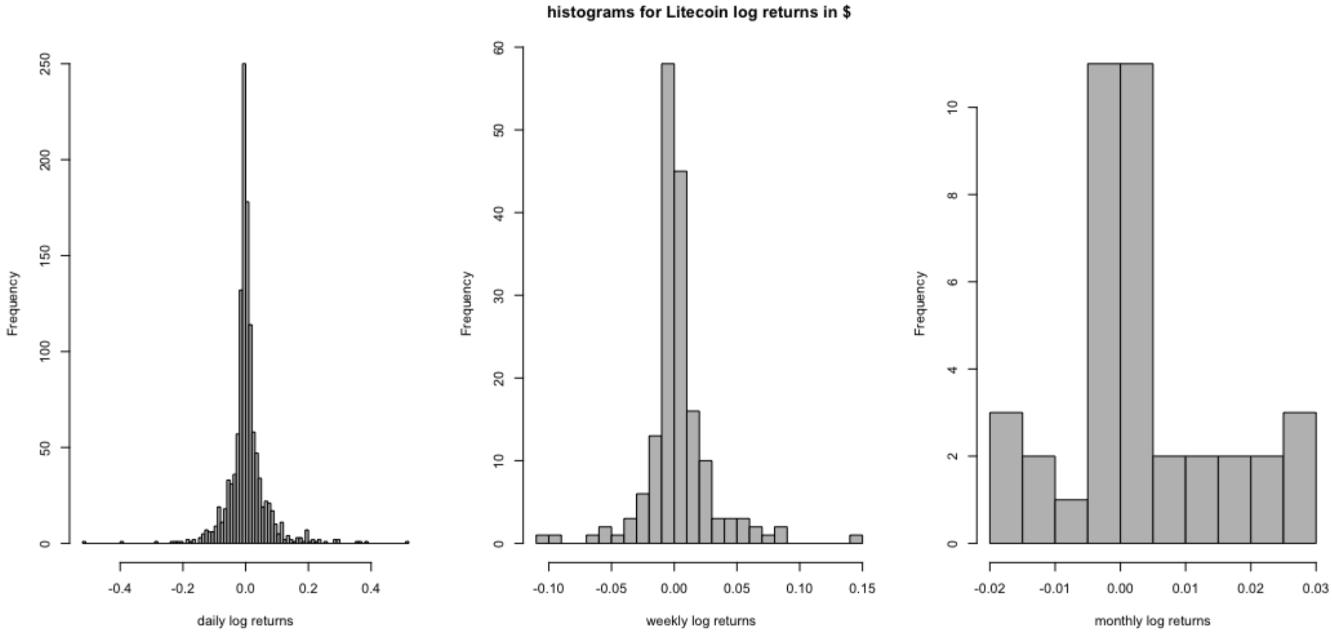

Figure 7: Litecoin Log Returns

From these plots, there appears to be a slight positive skew (the right tail seems to be fatter than the left one).
Other than that, an interesting observation can be made by considering the monthly plot, taking into account the monthly log returns. Here it can be observed that both the left tail and the right one are rather heavy: extreme log returns are not uncomon.

Then, for Ripple, a histogram of the daily Ripple log returns starting from January 2015 up until the end of April 2018 can be found in figure 8. Here, once again, positive returns denote profits, whereas negative returns denote losses.

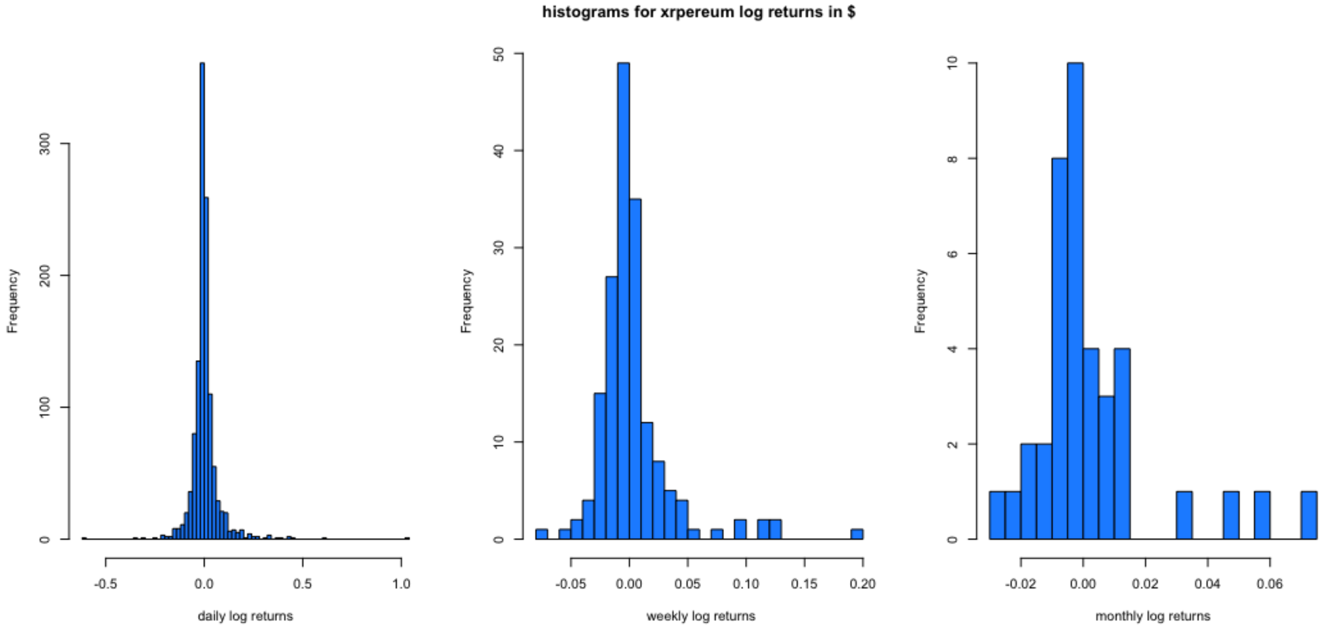

Figure 8: Ripple Log Returns

What is immediately apparent from these plots, is the fat right tail, that is especially clear in the histograms for the weekly and monthly log returns.
The extreme positive log returns undoubtedly originate from the last two months of 2017 as well.

Finally, for Ethereum, a histogram of the daily Ethereum log returns starting from January 2016 up until the end of April 2018 can be found in figure 9. Here, once again, positive returns denote profits, whereas negative returns denote losses.

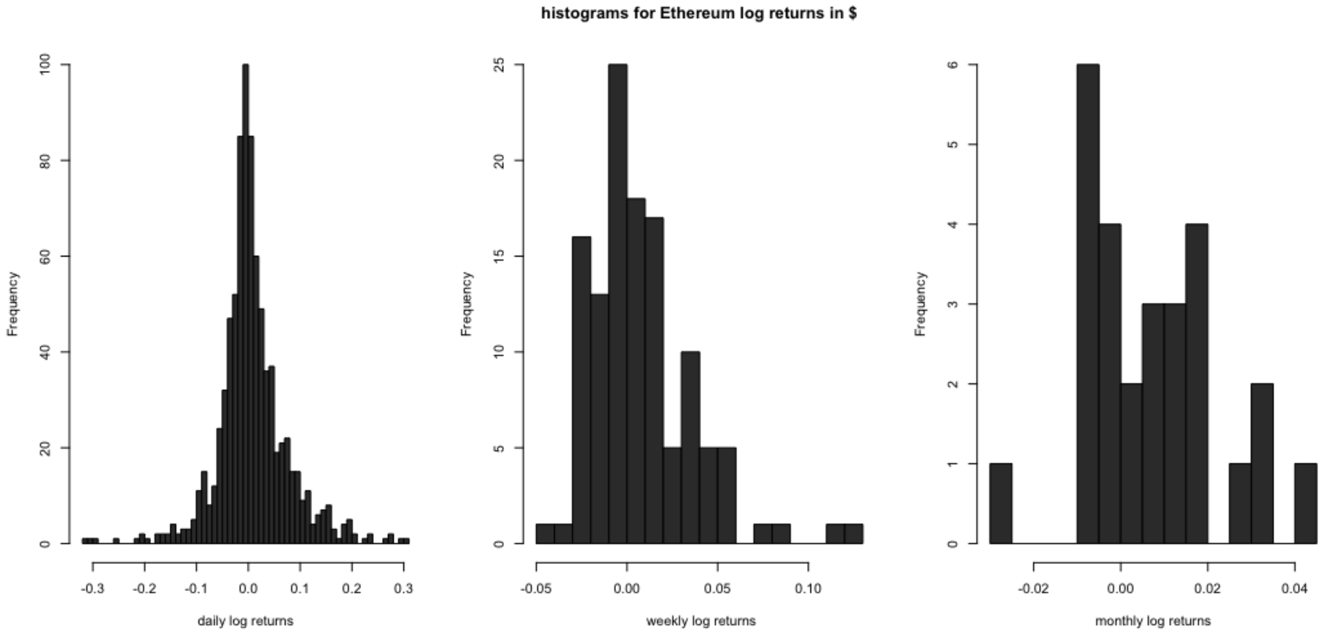

Figure 9: Ethereum Log Returns

What immediately jumps out, is that there appears to be a positive skew. This is especially apparent for the weekly log returns, visible in the upper right plot of figure 9. Overall, the monthly log returns show us that the log returns are spread much more evenly than the other currencies. It is, however, important to take into consideration that a shorter time span has been taken into account for Ethereum, as opposed to the other currencies where one extra year of data has been considered.

Short conclusion:

From the histograms of the log returns, what jumps out most of all is that there appears to be a positive skew for all four cryptocurrencies but Bitcoin: the right tail appears to be fatter than the left one. This is especially obvious for Ripple.
This leads us to believe that out of all four cryptocurrencies, the returns for Ripple produce most extreme values, especially on the right side of the distribution (i.e. the profits), and so fat tails are expected for Ripple and possibly also for the other three cryptocurrencies.

3.1.3 Summary Statistics

An overview of the daily exchange rates of Bitcoin, Litecoin, Ripple and Ethereum versus the US Dollar is given in table 1. Please note that these values are not to be compared mindlessly without taking into consideration the supply of the currency.

Statistics	Bitcoin	Litecoin	Ripple	Ethereum
Minimum	178.1	1.16	0.00409	0.9371
Q1	314.4	3.23	0.0065	0.1083
Median	631.8	3.9	0.0082	0.1735
Mean	2402	33.79	0.1686	207.1
Q3	2546	41.11	0.1895	307.9
Maximum	19500	358.3	3.38	1396
Skewness	2.194	2.504	4.0195	1.6802
Kurtosis	4.31	5.88	20.298	2.2144
SD	3754.76	63.44	0.394	294.7
Variance	14098211	4025.16	0.155	86848
IQR	2231.6	37.9	0.183	297.1

Table 1: Summary statistics of daily exchange rates of Bitcoin, Litecoin, Ripple and Ethereum versus the US Dollar

Table 1 give a simple reflection of the 'worth' of the cryptocurrencies in US Dollar. It is clear that the value of Ripple is lowest, whereas that of Bitcoin, the first and most popular cryptocurrency, has the highest value (here, it must be taken into consideration that the supply of Ripple is also much higher, and a lower value is therefore to be expected). This is true for the minimum, first quartile, median, mean, third quartile and maximum. The exchange rates of all four cryptocurrencies are positively skewed, which was also clear from the upward trend we saw in the time series plots earlier. Ripple shows the highest skewness, which was to be expected based on the extreme increase in the final quarter of 2017.
When it comes to kurtosis, the highest values can be observed for Ripple as well. As explained in the methodology section, statements are generally made with respect to a comparison to the Normal Distribution. We see that Bitcoin, Litecoin and Ripple show larger kurtosis than the Normal Distribution, whereas Ethereum shows little kurtosis compared to the Normal Distribution.
Finally, the exchange rates of Ripple show a smaller variance than the other coins, which can obviously for a large part be explained by the lower exchange rates, meaning comparisons between the currencies do not make much sense here. More useful information about the variance can be extracted from the log returns of the daily exchange rates, from which the same summary statistics can be extracted.

Summary statistics for the log returns of all four currencies are presented in table 2.

Statistics	Bitcoin	Litecoin	Ripple	Ethereum
Minimum	-0.238	-0.514	-0.6163	-0.3155
Q1	-0.010	-0.015	-0.020	-0.024
Median	0.003	0	-0.004	0.001
Mean	0.003	0.003	0.003	0.008
Q3	0.019	0.017	0.017	0.036
Maximum	0.225	0.510	1.027	0.303
Skewness	-0.417	0.710	3.132	0.268
Kurtosis	5.618	13.964	41.637	3.534
SD	0.041	0.063	0.075	0.069
Variance	0.002	0.004	0.006	0.005
IQR	0.029	0.032	0.037	0.060

Table 2: Summary statistics of log returns of daily exchange rates of Bitcoin, Litecoin, Ripple and Ethereum versus the US Dollar

From table 2 it can be seen that the results are a bit different compared to the daily exchange rates. Whereas Ripple still has the lowest minimum value, it also shows the largest maximum value now.
For Bitcoin, the mean and the median are equal to each other and are slightly positive. For Litecoin, the median is (almost) equal to 0, whereas the mean is also slightly positive. Ripple is the only currency with a negative median. Ethereum is the clear winner when it comes to the mean log returns, showing the highest value of all four.

Only the log returns of Bitcoin are negatively skewed, whereas all others are positively skewed with Ripple the most significant. This is consistent with our expectations based on the visualizations from the previous section: the histograms of the log returns.
A different way to measure the skewness can be done by comparing the quartiles by making use of the following formula:

$$\text{Skewness Measure} = \frac{\text{Upper hinge - Median}}{\text{Median- Lower hinge}} \tag{19}$$

Resulting in skewness levels 1.26 for Bitcoin, 1.14 for Litecoin, 1.22 for Ripple and 1.4 for Ethereum.
Note that the Lower hinge is similar to Q1, the upper hinge is similar to Q3. Surprisingly, this measure for skewness results in a different order for the four currencies, seeing as Ethereum now has a greater value for skewness than Ripple.
This can be explained by the fact that, this different skewness measure does not take into account the extreme values (it only compares the range between the median and Q3 to to the range between the median and Q1).
For this different skewness measure, it is intuitively clear that the closer the value is to 1, the more symmetric the distribution is. A value greater than 1 indicates that there is more mass in the range from the median to Q3 than there is in the range from the median to Q1.
This is indeed the case for all four currencies.
Also, this different skewness measure also shows that, when extreme values are left out, Bitcoin also shows a positive skew (as a value greater than 1 was obtained).
When it comes to the peakedness, from the value of the kurtosis it is clear that Ripple, once again, shows the highest peakedness. All four coins, however, show a greater peakedness compared to the Normal Distribution (remember that a value greater than 3 means that the distribution produces more extreme outliers than the Normal Distribution would).
This leads to believe that, especially for Ripple, there is relatively more weight in the tails as opposed to the center of the distribution, compared to the other currencies. This is in accordance with Ripple having fatter tails than the other cryptocurrencies. More about this will be covered in later sections (Value at Risk and Expected Shortfall, and Testing for Paretianity).

For the log returns, it makes more sense than for the exchange rates, to look at the values for the variance and make comparisons between the currencies.
It is clear that the log returns of Bitcoin have the lowest variance (and therefore the lowest standard deviations as well). The highest variance can be observed for Ripple.

The summary statistics can also be visualized using Boxplots. Boxplots, where a distinction has been made between daily, weekly and monthly losses, can be found in the appendix.

Short Conclusion:

From the summary statistics, we were particularly interested in the measures for skewness and kurtosis of the log returns, as they provide us with important information with respect to the distribution of the log returns and the tails of the distributions.
The summary statistics have shown that the distributions for Litecoin, Ripple and Ethereum all appear to be positively skewed, whereas for Bitcoin this only appears to be the case when extreme values are not taken into account (only taking into consideration the values between the upper hinge and the lower hinge).
The values for Kurtosis have shown us that all four cryptocurrencies come from distributions that produce relatively many extreme outliers (at least compared to the Normal Distribution). This is especially true for Ripple.

3.2 Volatility

Volatility, or simply the standard deviation of the log returns, might not be fixed, but may vary over time, as explained in the Methodology section.
Therefore, volatilities could be calculated over moving intervals of 20 days and could then be presented using histograms, to get a clear picture of the several values for volatility that were obtained.

Below, histograms for these 20-day-window volatilities are shown for Bitcoin, Litecoin, Ripple and Ethereum, in figure 10:

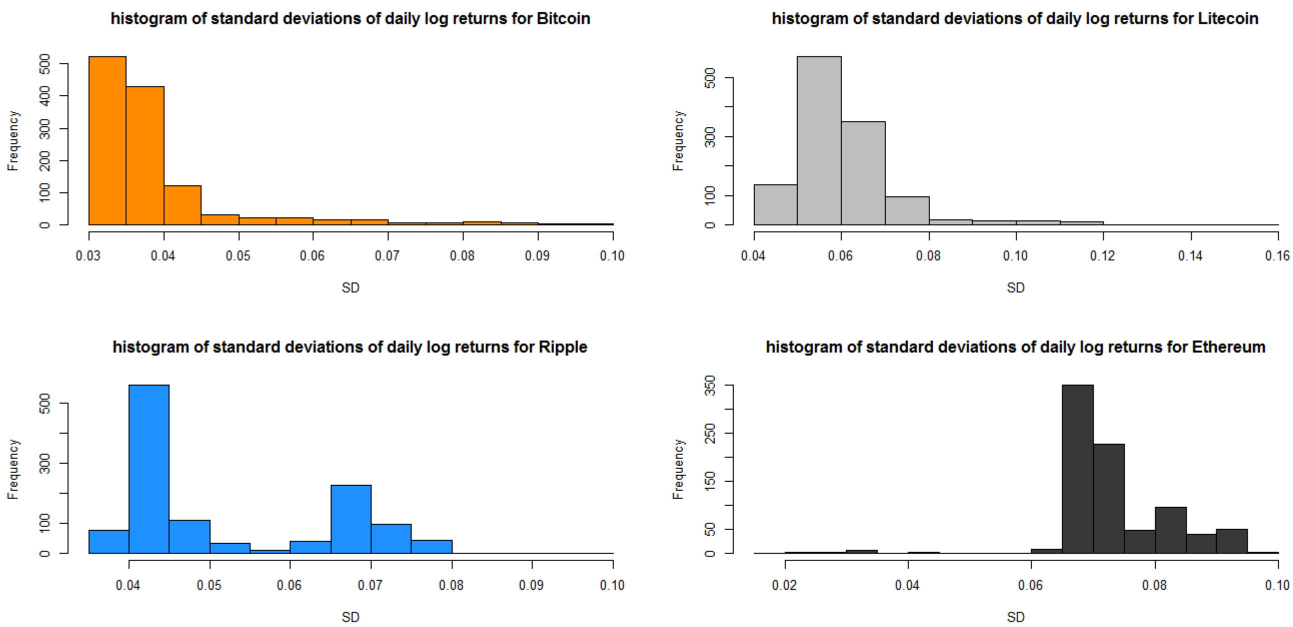

Figure 10: Histograms of standard deviations of daily log returns of the exchange rates over windows of width 20 days

What stands out is that the frequencies of higher volatility levels are lower than the ones for lower volatility levels. For Ripple, the distribution of volatilities appears to be bimodal. Bimodality might indicate the presence of two 'groups', or in this case, two time spans that behave very differently. Perhaps the extreme gains originating from the end of 2017 are responsible for deviating behavior when it comes to the volatility, resulting in their own distribution of volatilities, so to speak. Even though the other currencies also experienced extreme gains during this period, for Ripple this turned out to be even more extreme.

Furthermore, Litecoin appears to have the highest range of volatility, whereas this appears to be smallest for Ripple.
However, given the relatively low frequencies for some of the higher standard deviations, if these were to be excluded from the analysis, Ripple would actually show the highest range, whereas Ethereum would have the lowest range.

What also strikes as odd, is the fact that the distribution for the volatilities of Ethereum looks very different from the other distributions. Higher volatility rates (greater than 0.06) appear to be much more frequent for Ethereum than the lower volatility rates, a phenomenon that is not present in the other currencies.

When we convert the obtained volatility rates into annualized volatilities, we obtain the following measures of central tendency in table 3:

	Bitcoin	Litecoin	Ripple	Ethereum
Mean	0.736	1.159	0.999	1.392
Median	0.688	1.118	0.854	1.342
Mode	0.688	1.198	0.797	1.320

Table 3: Central tendency measures for the annualized volatility based on a 20-day rolling volatility

When we compare the measures for Bitcoin to the annual volatility that was obtained in [22] (an annual volatility of 77%), our results show that the annualized volatility has decreased slightly in the past 19 months (remember the authors of [22] used data up until September 2016, whereas our data extends to April 2018).

What immediately jumps out from table 3 is that the other cryptocurrencies show much higher volatility rates, which is especially true for Ethereum with an annualized volatility of over 130%. This is most likely due to its shorter existence, given that Ethereum was founded only in September 2015.

What also jumps out is that, while still more volatile than Bitcoin, Ripple appears to be less volatile than Litecoin and Ethereum. It is also noticeable that the mean, median and mode differ most for Ripple. This once again confirms the conjecture that Ripple has the most extreme outliers (also with respect to the volatility), as extremely volatile periods are likely to have influenced the mean significantly. On the long run, however, Ripple appears to be less volatile than Litecoin and Ethereum.

Short Conclusion:

From the volatility analyses, it can be concluded that the annual volatility for Bitcoin has dropped slightly compared to approximately two years ago, and is currently equal to about 70%. The other three cryptocurrencies are more volatile than Bitcoin, with Ethereum the most volatile, as its annual volatility comes down to more than 130%.
It is noteworthy that Ripple appears to be less volatile than Litecoin and Ethereum, even though in table 2 it was clear that Ripple showed the highest volatility (the highest values for the standard deviation and the variance). The reason for this lower volatility must therefore lie within the 20-day moving windows that were used.

3.3 Value at Risk and Expected Shortfall

Here, for the log returns of the daily exchange rates, Values at Risk and Expected Shortfall values have been calculated for 80%, 90%, 95% and 99% respectively.
This has been done both for losses as well as profits, meaning that for the losses we have used $\alpha = 0.01, 0.05, 0.1$ and 0.2, whereas for the profits we have used $\alpha = 0.8, 0.9, 0.95$ and 0.99, because the log returns have been used and so negative returns are losses(the left tail), and positive returns are profits(the right tail).

Bitcoin

First, the value for VaR (in dollars) has been calculated for several values of T and α in table 4. Here, T denotes the timespan between returns that was considered (so for T =1 week, weekly returns are considered, which are calculated by taking an average over the past 7 days of daily returns).

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 1$ day	-0.016	-0.039	-0.066	-0.122
$T = 1$ week	-0.007	-0.015	-0.022	-0.054
$T = 1$ month	-0.004	-0.010	-0.015	-

Table 4: Bitcoin Values at Risk (VaR) for several values of α and T: Losses

These values can be interpreted as follows:
For $\alpha = 0.2, T = 1$ day, the value -0.016 means that with 80% certainty, we will not lose more than 1.6% of our investment in Bitcoin in a period of 1 day.
As the value for α decreases, and so our confidence interval gets larger, the risk will of course increase, therefore leading to potentially greater losses. For example, if we want to make a statement with respect to the percentage we will probably not lose with 99% confidence, the value is equal to 12.2%, a much larger portion of our investment.
We also notice that the numbers are decreasing when we look at larger time intervals, like weeks or months.
This means that, based on the past data for Bitcoin, the Value at Risk actually decreases when a larger timespan is considered.
Note that the 99% confidence level ($\alpha = 0.01$) for time intervals of 1 month has been left out, as not enough months have been taken into account to provide an accurate value here.

Besides the Value at Risk, the Expected Shortfall has also been computed for Bitcoin, and is presented, first for the losses, in table 5 for the same values of α and T:

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 1$ day	-0.051	-0.077	-0.104	-0.167
$T = 1$ week	-0.017	-0.025	-0.033	-0.054
$T = 1$ month	-0.008	-0.013	-0.015	-

Table 5: Bitcoin Expected Shortfall (ES) for several values of α and T: Losses

The values in table 5 can be interpreted as follows:
For $\alpha = 0.2, T = 1$ day, the value -0.051 means that if we lose at least 1.6% of our investment (which is the corresponding Value at Risk), the expected loss is equal to 5.1%.
As the value for α decreases, and so our confidence interval gets larger, the expected losses will of course increase as well, as they are based on the corresponding Values at Risk, which we noticed were also increasing. With the same reasoning, the numbers are decreasing when we look at larger time intervals, like weeks or months.

Taking into account only the daily Value at Risk and Expected Shortfall, and looking only at the 5% quantile for now, we can say that we can expect to lose more than 6.6% of our investment in Bitcoin in one day, about once every 20 days (5% of all days considered). Provided that we find ourselves in one of those 'worst' days, we can expect to lose about 10.4% on that day.

But what about the profits?

For the profits, a similar analysis was made, for which table 6 can be consulted for the Value at Risk, and table 7 for the Expected Shortfall. Please note that the terms 'Value at Risk' and 'Expected Shortfall' are not really appropriate in this context, considering we are looking at profits instead of risks, but using these two quantities makes it nevertheless clear how the results were obtained).

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
$T = 1$ day	0.025	0.045	0.063	0.111
$T = 1$ week	0.012	0.020	0.026	0.034
$T = 1$ month	0.008	0.010	0.016	-

Table 6: Bitcoin Values at Risk (VaR) for several values of α and T: Profits

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
$T = 1$ day	0.054	0.074	0.095	0.145
$T = 1$ week	0.023	0.029	0.035	0.048
$T = 1$ month	0.012	0.015	0.017	-

Table 7: Bitcoin Expected Shortfall (ES) for several values of α and T: Profits

First of all, when we compare the profits to the losses, the Values at Risk appear to be higher for the profits than for the losses when looking at the 20%/10% and corresponding 80% /90% quantiles, but not for the 5%/1% and corresponding 95%/99% quantiles. This is exactly in correspondence with our earlier findings that Bitcoin appeared to positively skewed only when not taking into account the ends of the tails, therefore ignoring the most extreme values (remember we obtained a negative skew in table 2, but the different skewness measure, taking into account only the 'mid region' of the log returns, returned a positive skew).

For the profits, it can be concluded that, when taking into account daily profits and considering the 95% quantile, we can expect to gain more than 6.3% of our Bitcoin investments in one day, about once every 20 days. Provided that we find ourselves in one of those 'best' days, we can expect to gain more than 9.5%.
Note that this implies that Bitcoin is actually a rather risky cryptocurrency, where you can expect to lose more money than you gain during extreme events. However, the mean return could make up for the extremes and end you in the green, so these results should be interpreted with caution.

Litecoin

For Litecoin, the same analyses can be made as for Bitcoin, resulting in the following four tables for the VaR (tables 8, 9) and ES (table 10,11):

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 1$ day	-0.021	-0.052	-0.083	-0.164
$T = 1$ week	-0.009	-0.016	-0.034	-0.102
$T = 1$ month	-0.005	-0.015	-0.019	-

Table 8: Litecoin Values at Risk (VaR) for several values of α and T: Losses

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
$T = 1$ day	0.026	0.061	0.091	0.201
$T = 1$ week	0.012	0.027	0.050	0.078
$T = 1$ month	0.010	0.021	0.027	-

Table 9: Litecoin Values at Risk (VaR) for several values of α and T: Profits

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 1$ day	-0.068	-0.100	-0.136	-0.245
$T = 1$ week	-0.025	-0.040	-0.059	-0.102
$T = 1$ month	-0.012	-0.017	-0.019	-

Table 10: Litecoin Expected Shortfall (ES) for several values of α and T:Losses

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
$T = 1$ day	0.082	0.125	0.173	0.297
$T = 1$ week	0.037	0.054	0.072	0.099
$T = 1$ month	0.021	0.026	0.028	-

Table 11: Litecoin Expected Shortfall (ES) for several values of α and T:Profits

Making similar statements as for Bitcoin, we can conclude from tables 8 to 11 that, when taking into account daily losses and considering the 5% quantile, we can expect to lose more than 8.3% of our Litecoin investments in one day, about once every 20 days. Provided that we find ourselves in one of those 'worst' days, we can expect to lose more than 13.6%.
For the profits, we can expect to gain more than 9.1% once every 20 days, whereas on such a particular day, the expected gain is equal to about 17.3%.
First of all, note that these losses and gains are more extreme than the ones we found for Bitcoin. This is consistent with our earlier findings with respect to the volatility rates, where Litecoin appeared to be much more volatile than Bitcoin.
Also, we note that for Litecoin, there is a clear positive skew (also consistent with earlier findings). For each of the quantiles, the gains are higher than the (absolute) losses. So even though the risk when investing in Litecoin is greater than for Bitcoin, the potential gains are also much greater.

Ripple

For Ripple, the same analyses can be made as for Bitcoin and Litecoin, resulting in the following four tables for the VaR (table 12) and ES (table 14):

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 1$ day	-0.029	-0.054	-0.078	-0.159
$T = 1$ week	-0.016	-0.025	-0.033	-0.071
$T = 1$ month	-0.010	-0.018	-0.030	-

Table 12: Ripple Values at Risk (VaR) for several values of α and T: Losses

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
$T = 1$ day	0.024	0.058	0.102	0.245
$T = 1$ week	0.011	0.026	0.047	0.117
$T = 1$ month	0.009	0.013	0.045	-

Table 13: Ripple Values at Risk (VaR) for several values of α and T: Profits

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 1$ day	-0.070	-0.100	-0.137	-0.251
$T = 1$ week	-0.027	-0.035	-0.044	-0.071
$T = 1$ month	-0.017	-0.023	-0.029	-

Table 14: Ripple Expected Shortfall (ES) for several values of α and T: Losses

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
$T = 1$ day	0.092	0.140	0.210	0.410
$T = 1$ week	0.046	0.070	0.099	0.140
$T = 1$ month	0.029	0.044	0.058	-

Table 15: Ripple Expected Shortfall (ES) for several values of α and T: Profits

Making similar statements as for Bitcoin and Litecoin, we can conclude from tables 12 to 15 that, when taking into account daily losses and considering the 5% quantile, we can expect to lose more than 7.8% of our Ripple investments in one day, about once every 20 days. Provided that we find ourselves in one of those 'worst' days, we can expect to lose more than 13.7%.
For the profits, we can expect to gain more than 10.2% once every 20 days, whereas on such a particular day, the expected gain is equal to about 21%.
First of all, note that these gains are more extreme than the ones we found for Bitcoin and Litecoin, but the losses are not (they are slightly worse than the ones for Bitcoin, but not compared to Litecoin). This is consistent with our earlier findings with respect to the skewness, as the skewness level for Ripple was highest for all four cryptocurrencies(see table 2).
Also, this extreme positive skew can be found when comparing the quantiles. For each of the quantiles, except for the 20% and corresponding 80% ones, the gains are (much) higher than the (absolute) losses. This is especially apparent at the far end of the tails: whereas the 1% quantile shows an expected loss of 25.1%, the expected gains at the 99% quantile are equal to 41%!

Ethereum

Finally, for Ethereum, the same analyses can be made as well, resulting in the following four tables for the VaR (table 16) and ES (table 18):

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 1$ day	-0.034	-0.056	-0.091	-0.196
$T = 1$ week	-0.013	-0.024	-0.027	-0.045
$T = 1$ month	-0.007	-0.009	-0.027	-

Table 16: Ethereum Values at Risk (VaR) for several values of α and T: Losses

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
$T = 1$ day	0.047	0.089	0.134	0.202
$T = 1$ week	0.029	0.040	0.053	0.072
$T = 1$ month	0.016	0.027	0.034	-

Table 17: Ethereum Values at Risk (VaR) for several values of α and T: Profits

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 1$ day	-0.076	-0.109	-0.146	-0.249
$T = 1$ week	-0.024	-0.029	-0.034	-0.045
$T = 1$ month	-0.012	-0.018	-0.027	-

Table 18: Ethereum Expected Shortfall (ES) for several values of α and T: Losses

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
$T = 1$ day	0.106	0.144	0.182	0.252
$T = 1$ week	0.049	0.063	0.0767	0.097
$T = 1$ month	0.027	0.035	0.037	-

Table 19: Ethereum Expected Shortfall (ES) for several values of α and T: Profits

For Ethereum, we can conclude from tables 16 to 19 that, when taking into account daily losses and considering the 5% quantile, we can expect to lose more than 9.1% of our Ethereum investments in one day, about once every 20 days. Provided that we find ourselves in one of those 'worst' days, we can expect to lose more than 14.6%.
For the profits, we can expect to gain more than 13.4% once every 20 days, whereas on such a particular day, the expected gain is equal to about 18.2%.
From these numbers, we can conclude once again that the potential gains seem to outweigh the potential losses. For each of the percentiles, the (absolute) VaR and ES values are higher for the profits than for the losses. This is consistent with our earlier findings with respect to the positive skewness(see table 2).
Furthermore, from these tables it once again becomes clear that when it comes to the losses, Ethereum has the highest VaR and ES compared to the other three cryptocurrencies. For the profits, this difference is less obvious, but, especially compared to Bitcoin and Litecoin, certainly present. This is in accordance with the highest volatility rates we found for Ethereum in the Volatility section.

Overview of all four cryptocurrencies

In order to make a good comparison between the four cryptocurrencies of interest, and also to compare profits and losses in an organized fashion, tables 20 and 21 have been created. They provide a ranking based on both the Value at Risk and the Expected Shortfall, when considering $\alpha = 0.05$ for losses and $\alpha = 0.95$ for profits, therefore considering equally large portions of both tails.
Note that the values in both tables correspond to the ones from tables 4 to 19, each time taking the absolute values of the ones found in the third columns from tables 4 to 19, corresponding to $\alpha = 0.05$ for losses (or 0.95 for profits), converted to a percentage (multiplied by 100) and rounded to integers.

	VaR			ES		
Risk Rank	Daily VaR	Weekly VaR	Monthly VaR	Daily ES	Weekly ES	Monthly ES
1	Ethereum(9)	Litecoin(3)	Ripple(3)	Ethereum(15)	Litecoin(6)	Ripple(3)
2	Litecoin(8)	Ripple(3)	Ethereum(3)	Ripple(14)	Ripple(4)	Ethereum(3)
3	Ripple(8)	Ethereum(3)	Litecoin(2)	Litecoin(14)	Bitcoin(3)	Litecoin(2)
4	Bitcoin(7)	Bitcoin(2)	Bitcoin(2)	Bitcoin(10)	Ethereum(3)	Bitcoin(2)

Table 20: Ordering of cryptocurrencies in Value at Risk and Expected Shortfall at $\alpha = 0.05$, with loss percentages (order from most risky to least risky)

	VaR			ES		
Profit Rank	Daily VaR	Weekly VaR	Monthly VaR	Daily ES	Weekly ES	Monthly ES
1	Ethereum(13)	Ethereum(5)	Ripple(5)	Ripple(21)	Ripple(10)	Ripple(6)
2	Ripple(10)	Litecoin(5)	Ethereum(3)	Ethereum(18)	Ethereum(8)	Ethereum(4)
3	Litecoin(9)	Ripple(5)	Litecoin(3)	Litecoin(17)	Litecoin(7)	Litecoin(3)
4	Bitcoin(6)	Bitcoin(3)	Bitcoin(2)	Bitcoin(10)	Bitcoin(4)	Bitcoin(2)

Table 21: Ordering of cryptocurrencies in Value at Risk and Expected Shortfall at $\alpha = 0.95$ with gain percentages (order from most profitable to least profitable)

A quick glance at tables 20 and 21 is enough to immediately notice that the values for profits are overall higher than the ones for losses (once again confirming the positive skew present in the datasets), especially for Ripple, Ethereum and Litecoin.
Also, what jumps out is that, while on a daily basis the numbers differ greatly between the four cryptocurrencies, the differences on a weekly and monthly basis are much less extreme, especially for the losses.

Apparently, heavy daily fluctuations are compensated rather quickly, making for much less extreme losses or profits on a monthly, or even weekly, basis.

Also, in table 20, we see that for the losses, on a daily basis, Bitcoin shows the lowest VaR (7%) and ES(10%), with only a relatively small difference. The differences are a bit bigger for the other three currencies, Ethereum going from a VaR of 9% to an ES of 15%, and Litecoin and Ripple going from a VaR of 8% to an ES of 14%.
For the profits, found in table 21, these differences are even more extreme. This once again implies that Litecoin, Ripple and Ethereum show heavier tails than Bitcoin, which is in accordance with earlier findings. Especially the values for the profits for Ripple, with a VaR equal to 10% and an ES of 21% (meaning the ES is more than twice as large as the VaR), are indicative of extreme tail behavior for the right tail. This result is once again in accordance with earlier findings, such as the values of Kurtosis we saw in table 2, where Ripple showed a Kurtosis value of 41.6, more than three times as large as the second largest value (Litecoin).

When we compare our findings to those of Osterrieder et. al [21], who found that Bitcoin was least risky, after which came Litecoin, and then Ripple, (both for the VaR and ES, and also both for the losses as well as the profits), this is somewhat consistent with our own findings.

When it comes to the weekly losses, the order of the currencies with respect to VaR and ES are very similar, with the small exception that, looking at the ES, Bitcoin and Ethereum are about equally risky (with a small difference of 0.1 percent, not visible in the table). On a monthly basis, the order of the currencies for the VaR and the ES are exactly the same.

Something else that immediately stands out is that, looking solely at the Expected Shortfall, there is a clear rank for the cryptocurrencies: Ripple has turned out most profitable, after which comes Ethereum, then Litecoin, and finally Bitcoin.
However, when taking into consideration the Expected Shortfall, Ripple has surpassed Ethereum in rank. This can most likely be explained by some extreme values (remember that Ripple has shown some extreme gains in December 2017, with a maximum increase of 102 % in one day [3]. These extreme values are likely to have caused the interaction effect that does not appear to occur for the other currencies.
Ignoring Ripple, though, the order for Bitcoin, Litecoin and Ethereum is the same for the VaR and ES, concluding that, when it comes to these three currencies, Ethereum is the most profitable, after which comes Litecoin, and finally Bitcoin.

Short Conclusion:

When taking into account the 5th and 95th quantiles of the distributions for the log returns, it is clear that Bitcoin is the least risky (but also the least profitable of all four cryptocurrencies), with a daily VaR of about 7% and an ES of about 10%. This is in accordance with findings from previous research by Osterrieder et. al. [21].
On a monthly basis, the values for the VaR and ES are only about 2%, which shows that Bitcoin is not as risky on a monthly basis.
For the other three cryptocurrencies of interest, the order is not so clear. Whereas Ethereum appears to be most volatile on a daily basis, showing the most extreme losses and profits, on a monthly basis Ripple shows the highest losses and profits (which is in accordance with earlier findings with respect to volatility), but the difference with Bitcoin is, especially for the losses, rather small (about 1%). For the profits, the difference between the most profitable (Ripple) and the least profitable (Bitcoin) is a bit larger (about 4%), confirming once again the existence of the fat right tail for Ripple.
Overall, the four cryptocurrencies appear to be much more risky for short term investments, than they are when investing for longer periods of time.
It could therefore be considered a wise decision to hold on to your cryptocurrencies a little while longer when losses have been incurred, as the market tends to correct itself more on the long run and you might end up in the green again.

3.4 Fitting of distributions

In this section, the distributions mentioned in the Methodology section, will be fitted on the log returns of the daily exchange rates of the four cryptocurrencies that are investigated.

3.4.1 Visual Explanations

First we will see how the distributions fit the data for the log returns of Bitcoin. The plots from figures 11 and 12 show the histogram corresponding to the Bitcoin log returns of the daily exchange rates, a total of six times. In each of the subplots, a different distribution has been fitted, using the maximum likelihood parameters for the corresponding distribution.

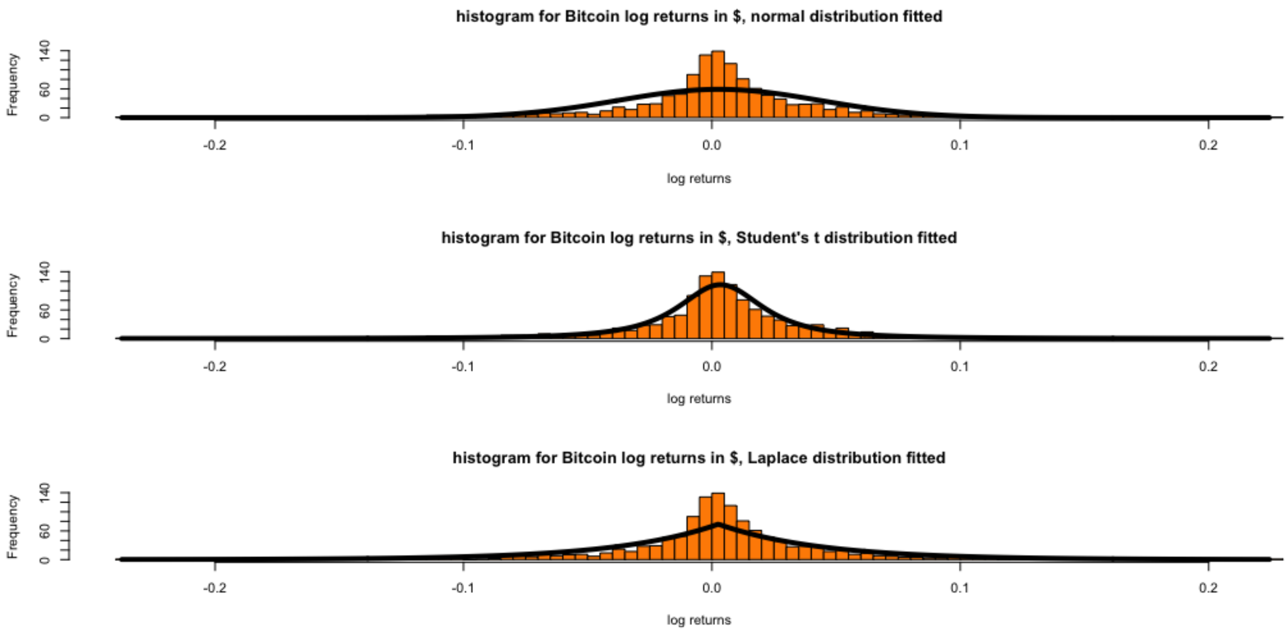

Figure 11: Fitting of the Normal, Student's t and Laplace Distributions on Bitcoin log returns on the daily exchange rates

Looking at figure 11, the first plot shows the Normal Distribution fit. What is immediately apparent, is that this distribution does not capture the data well at all. The peak appears to be much higher than the Normal Distribution would suggest. Also, the tails seem to be heavier.
Clearly, this is no surprise, given the earlier results with respect to the tails, and the summary statistics providing information with respect to the kurtosis compared to the Normal Distribution.
When we take a look at the second plot, the Student's t Distribution, this appears to be a decent fit (remember from the Introduction that Chan et. al. [4] also found this to be a well fitting distribution for some cryptocurrencies). The Laplace Distribution, though perhaps slightly better than the Normal Distribution, does also not seem to capture the data well, which is especially true for the peak.

If we then move on to the next two distribution fits, we see that the first plot in figure 12 (the Normal Inverse Gaussian Distribution) and final plot (the Generalized Hyperbolic Distribution) also fit the data well, as expected based on the findings by [4].

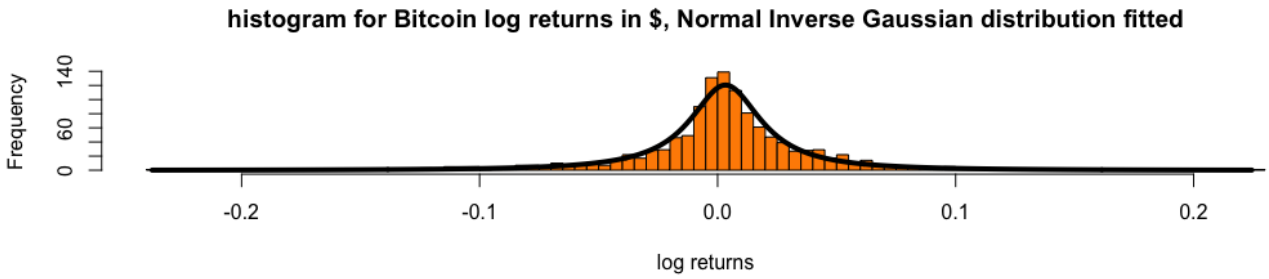

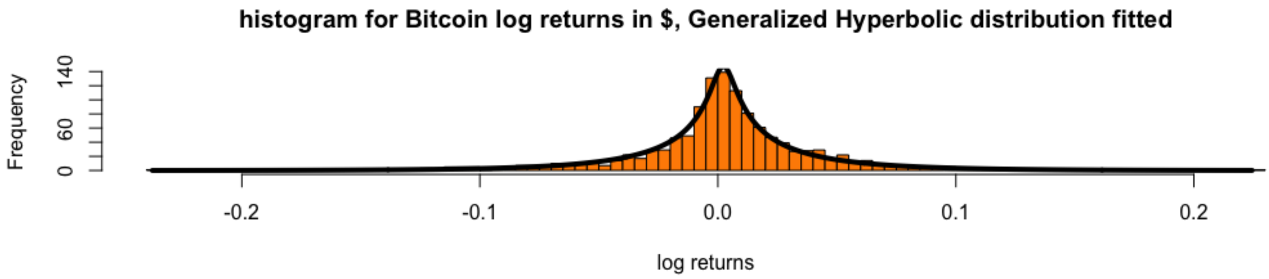

Figure 12: Fitting of the Normal Inverse Gaussian and Generalized Hyperbolic Distributions on Bitcoin log returns on the daily exchange rates

From these plots, it is clear that both appear to be decent fits for the data.
In total, this means that the Student's t, Normal Inverse Gaussian, and Generalized Hyperbolic Distribution seem to fit the data best. Given the fact that these three distributions (all derived from the Generalized Hyperbolic superclass), use more parameters than simply a location and a scale or dispersion parameter, this is no surprise, as these aren't necessarily symmetrical distributions, contrary to the Normal and the Laplace Distribution.
The Student's t, Normal Inverse Gaussian and Generalized Hyperbolic Distributions all come with (a) skewness and/or shape parameter(s), which appear to be very relevant indeed.

Given that the plots in figures 11 and 12 provide some information with respect to how well a distribution fits the data, a fit can be better visualized by making use of Quantile Comparison (QQ) plots.

Figures 13,14 and 15 show the Quantile-Comparison plots for the corresponding distributions. A line (with intercept 0 and slope 1) was plotted as well to give a better visualization of whether the data falls on this straight line and therefore give insight as to whether the sample comes from the theoretical distribution).
From the QQ-plots, it is clear that the data 'in the middle' appears to fall on the straight lines, but the tails do not seem to, even for the distributions that seemed to fit the data well based on figures 11 and 12. This is consistent with our expectation, stated in section 1 where we explained why, when fitting theoretical distributions to data, the fit will mostly be based on the mid region.

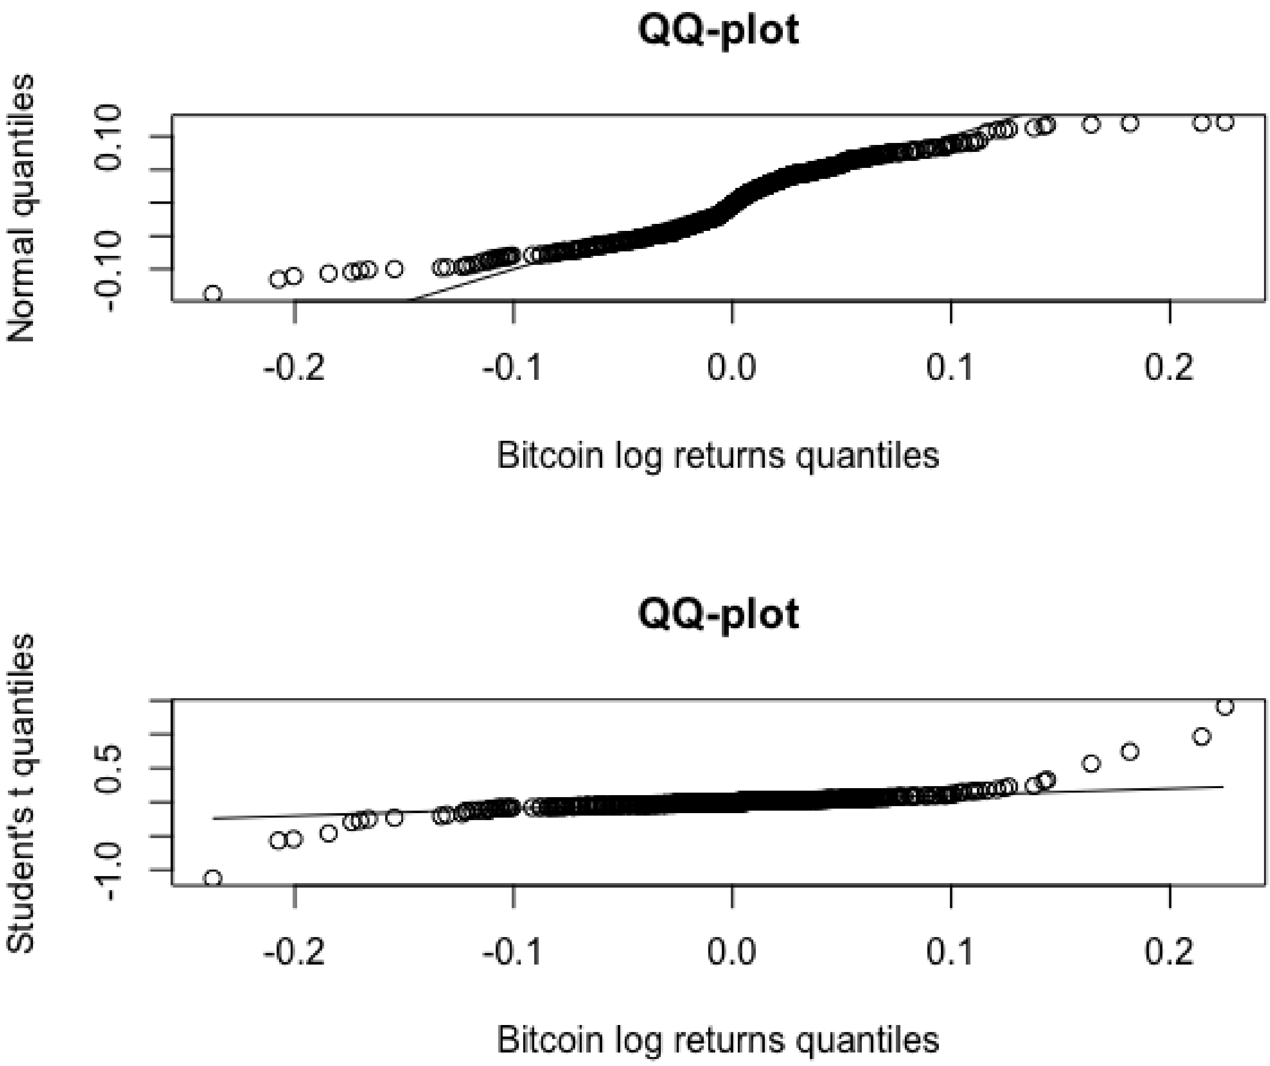

Figure 13: Corresponding QQ-plots for Normal and Student's t fits on Bitcoin log returns

From these figures, it is clear that the points appear to fall on the straight line most of all for the Generalized Hyperbolic and the Normal Inverse Gaussian Distributions.

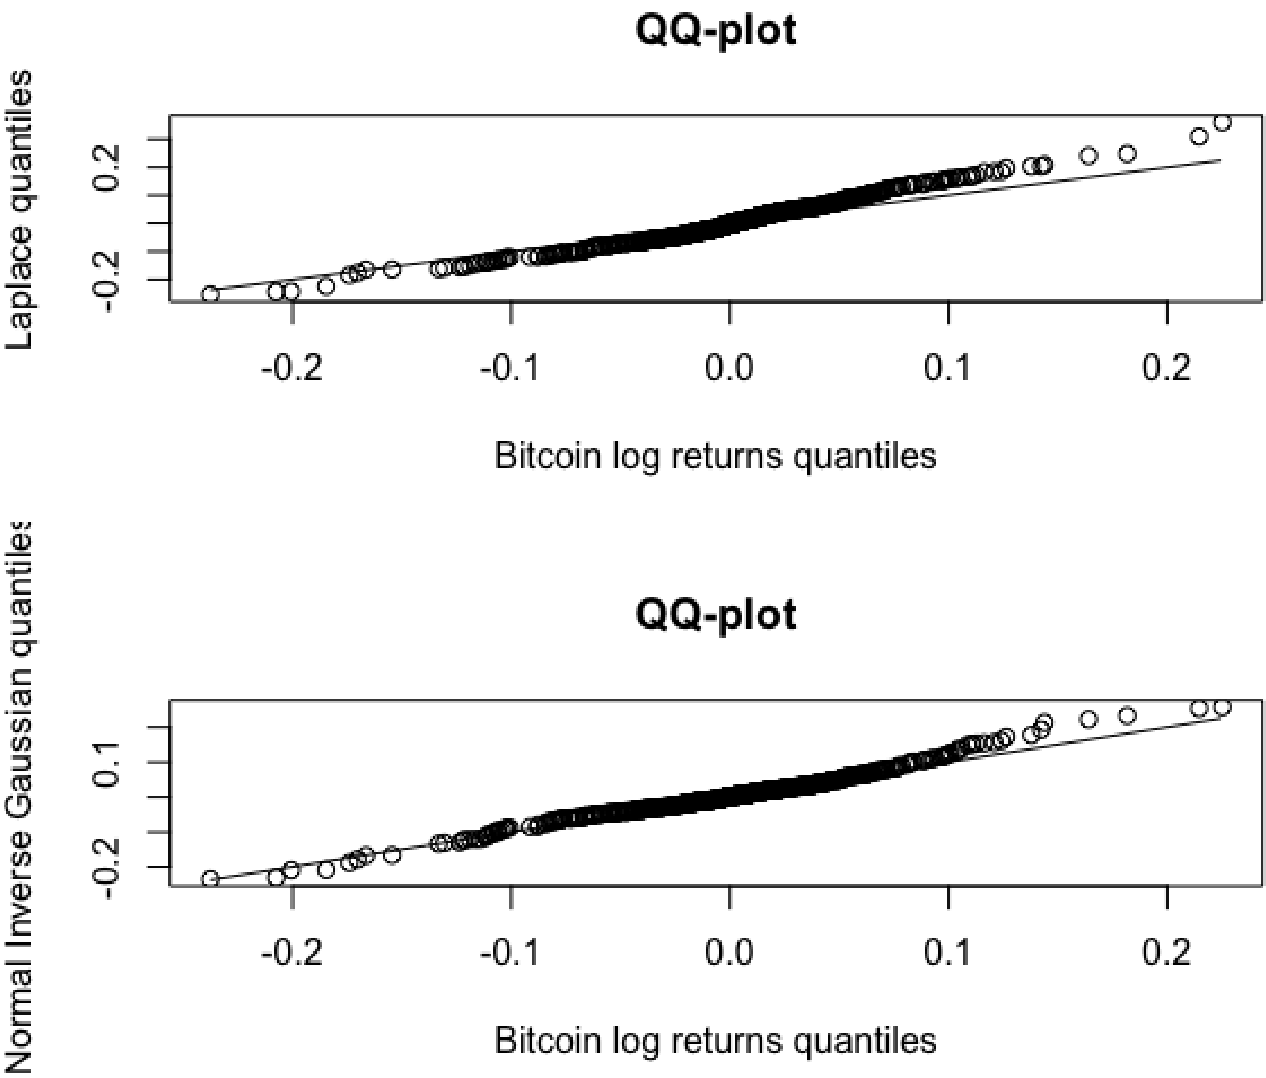

Figure 14: Corresponding QQ-plots for Laplace and Normal Inverse Gaussian fits on Bitcoin log returns

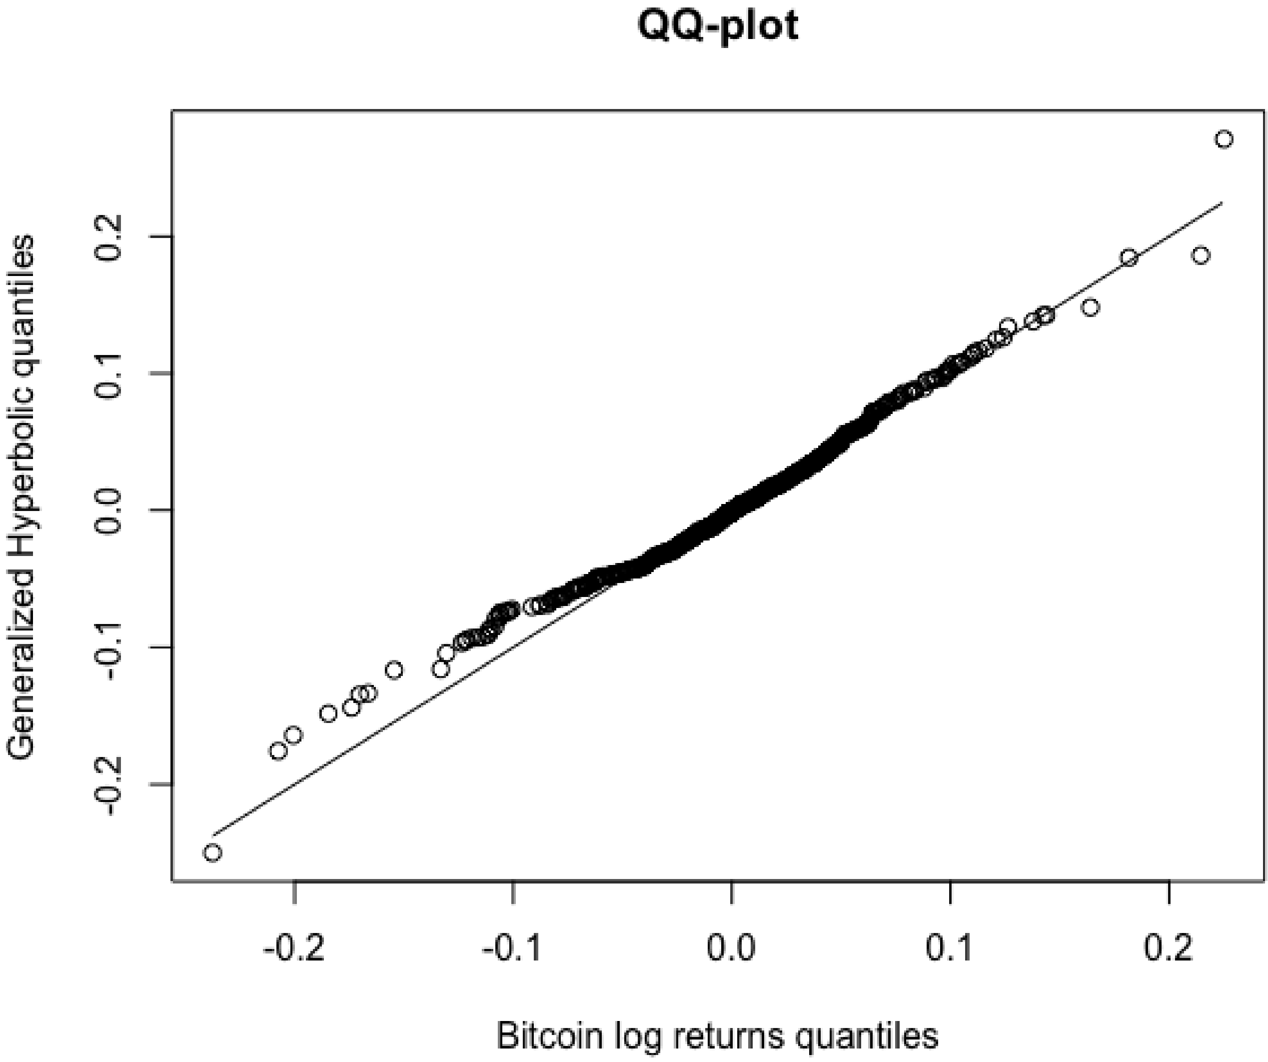

Figure 15: Corresponding QQ-plots for the Generalized Hyperbolic fit on Bitcoin log returns

For Litecoin, Ripple and Ethereum, similar plots are obtained (which can be found in the appendix). The same three distributions (Student's t, Normal Inverse Gaussian and Generalized Hyperbolic) appear to fit the data best, apart from deviations at the left and right tails.

Apart from visualizations, actual statistical tests will provide more information with respect to the significance of the distribution fits.
In the next section, the Kolmogorov-Smirnov and Chi-Square statistics will be computed to see whether the assumptions based on these plots, can be confirmed.

Short Conclusion

From the visualizations we can conclude that the distributions with strong peakedness and asymmetrical considerations work better than the ones that don't have these characteristics. The best fitting distributions appear to be the Generalized Hyperbolic Distribution and the Normal Inverse Gaussian Distribution.
This is visible in both the histogram with density plots, as well as the QQ-plots. We do notice,

especially from the QQ-plots, that the tails are captured by the distributions least of all.

3.4.2 Kolmogorov-Smirnov and Chi-Square statistics

In this section, the assumptions based on the plots from the previous section, will first be tested using the Kolmogorov-Smirnov statistic.
As explained in the Methodology section, 1000 random samples were drawn from the theoretical distributions each, after which they were compared to the data. Averages of the $p-$values, along with standard deviations are reported in table 22. Also, the percentage of $p-$values for which $p > 0.05$, implying the values come from this distribution, is reported here.

	Bitcoin	**Litecoin**	**Ripple**	**Ethereum**
	$\mu \approx 2.9 \cdot 10^{-8}$	$\mu \approx 7.3 \cdot 10^{-14}$	$\mu \approx 1.2 \cdot 10^{-15}$	$\mu \approx 0.3 \cdot 10^{-3}$
Normal Distribution	$\sigma \approx 5.1 \cdot 10^{-7}$	$\sigma \approx 3.8 \cdot 10^{-13}$	$\sigma \approx 1.8 \cdot 10^{-14}$	$\sigma \approx 0.6 \cdot 10^{-3}$
	$\% = 0$	$\% = 0$	$\% = 0$	$\% = 0$
	$\mu \approx 0.21$	$\mu \approx 0.23$	$\mu \approx 0.29$	$\mu \approx 0.19$
Student's t Distribution	$\sigma \approx 0.16$	$\sigma \approx 0.15$	$\sigma \approx 0.21$	$\sigma \approx 0.14$
	$\% \approx 83.8$	$\% \approx 90.1$	$\% \approx 88.1$	$\% \approx 86.7$
	$\mu \approx 6.3 \cdot 10^{-7}$	$\mu \approx 6.0 \cdot 10^{-12}$	$\mu \approx 3.7 \cdot 10^{-14}$	$\mu \approx 0.3 \cdot 10^{-3}$
Laplace Distribution	$\sigma \approx 3.7 \cdot 10^{-6}$	$\sigma \approx 3.2 \cdot 10^{-11}$	$\sigma \approx 3.4 \cdot 10^{-13}$	$\sigma \approx 0.6 \cdot 10^{-3}$
	$\% = 0$	$\% = 0$	$\% = 0$	$\% = 0$
	$\mu \approx 0.30$	$\mu \approx 0.28$	$\mu \approx 0.53$	$\mu \approx 0.47$
NIG Distribution	$\sigma \approx 0.20$	$\sigma \approx 0.16$	$\sigma \approx 0.25$	$\sigma \approx 0.24$
	$\% \approx 91.6$	$\% \approx 94.1$	$\% \approx 98.8$	$\% \approx 97.1$
	$\mu \approx 0.46$	$\mu \approx 0.37$	$\mu \approx 0.64$	$\mu \approx 0.69$
G. Hyperbolic Distribution	$\sigma \approx 0.28$	$\sigma \approx 0.21$	$\sigma \approx 0.25$	$\sigma \approx 0.24$
	$\% \approx 94.6$	$\% \approx 94.3$	$\% \approx 99.4$	$\% = 100$

Table 22: Overview of Kolmogorov-Smirnov p-value statistics for all distributions and all cryptocurrencies considered

From the table, it becomes clear that the best fitting distributions are the Normal Inverse Gaussian Distribution and the Generalized Hyperbolic Distribution. For Bitcoin, approximately 94.6% of all 1000 random drawings from the Generalized Hyperbolic Distribution match the data, resulting in an average $p-$ value of approximately 0.46.
For Litecoin, this percentage is similar: approximately 94.3 % of the random drawings are significant matches, resulting in an average $p-$ value of 0.37.
For Ripple and Ethereum, however, (almost) all $p-$values indicated the data was drawn from the Generalized Hyperbolic Distribution.
Note that for Ethereum, the percentage of 100 was not rounded.
The second best fitting distribution (in terms of the percentages of $p > 0.05$) is the Normal Inverse Gaussian distribution, and the third best fit is the Student's t distribution. This is the case for all four cryptocurrencies.

Note that the three best fits are exactly the three distributions that were assumed to fit the data best based on the plots from figures 11, 12,13 and 14.

From table 22 it is also very clear that the Normal Distribution and the Laplace Distribution are not good fits at all.
Based on the plots in figures 11, 12, 13 and 14 this was to be expected, but in previous research by Chan et. al. [4] the Laplace Distribution did appear to be a good fit for one of the currencies they investigated (MaidSafeCoin, which was not included in this research, but still, it is odd that here, this distribution does not seem to fit at all).

The results just described are partly confirmed by the Chi-Square goodness of fit test.
The Chi-square test was performed, using both 10 bins and 5 bins.
When using 10 bins, for Bitcoin, none of the distributions appeared to be significant fits, according to the Chi-Square test. The same can be said for Litecoin. For Bitcoin and Litecoin, $p-$values below 10^{-5} were obtained for all of the tests. For Ripple and Ethereum, however, the only two significant results that were obtained, were obtained for the Normal Inverse Gaussian Distribution: here, the hypotheses that the observed values were drawn from the NIG Distribution, was not rejected, as $p-$values of 0.08 and 0.11 were obtained for Ripple and Ethereum, respectively.

Interestingly, the Generalized Hyperbolic Distribution resulted in $p-$values below 10^{-9}, and therefore, according to the Chi-Square test, this distribution is not a good fit for any of the currencies, including Ripple and Ethereum.

When making use of only 5 bins, different results were obtained: this time, for the Normal Inverse Gaussian Distribution, $p-$values of 0.07, 0.1, 0.61 and 0.13 were obtained for Bitcoin, Litecoin, Ripple and Ethereum respectively, implying the Normal Inverse Gaussian Distribution is a decent fit for the data.
The Generalized Hyperbolic Distribution, however, resulted in significantly good fits for all four cryptocurrencies, with $p \approx 0.1, p \approx 0.15, p \approx 0.5$ and $p \approx 0.59$ for Bitcoin, Litecoin, Ripple and Ethereum respectively.

This is consistent with the results of the Kolmogorov-Smirnov results, where the Generalized Hyperbolic Distribution fitted the data best for all four cryptocurrencies as well.
For Ripple, the Chi-Squared statistic turned out to be lower for the Normal Inverse Gaussian Distribution, resulting in a higher $p-$value. Therefore, one could argue that perhaps the Normal Inverse Gaussian is a better fit for Ripple after all (remember that we also obtained a good fit based on the Kolmogorov-Smirnov statistic in table 22).

Overall, considering the fact that the Generalized Hyperbolic Distribution has performed rather well in both cases (for the Kolmogorov-Smirnov, this distribution also performed well for all four cryptocurrencies), it can be concluded that the Generalized Hyperbolic Distribution is the best fit for Bitcoin, Litecoin, Ripple and Ethereum (even though the results for Ripple were somewhat ambiguous (here, both the Generalized Hyperbolic and the Normal Inverse Gaussian) could be considered the best fits).

Table 23 gives an overview of the best fitting distributions for the cryptocurrencies of interest, together with parameter estimates that were obtained by `R`.

Cryptocurrency	Best Fitting distribution	Parameter Estimates
Bitcoin	Generalized Hyperbolic	$\widehat{\mu} = 0.002$ $\widehat{\delta} = 0.005$ $\widehat{\alpha} = 21.448$ $\widehat{\beta} = 0.308$ $\widehat{\lambda} = 0.288$
Litecoin	Generalized Hyperbolic	$\widehat{\mu} = 0.000$ $\widehat{\delta} = 0.008$ $\widehat{\alpha} = 9.970$ $\widehat{\beta} = 1.077$ $\widehat{\lambda} = 0.003$
Ripple	Generalized Hyperbolic	$\widehat{\mu} = -0.005$ $\widehat{\delta} = 0.016$ $\widehat{\alpha} = 6.820$ $\widehat{\beta} = 1.805$ $\widehat{\lambda} = -0.285$
Ethereum	Generalized Hyperbolic	$\widehat{\mu} = -0.005$ $\widehat{\delta} = 0.009$ $\widehat{\alpha} = 16.207$ $\widehat{\beta} = 2.797$ $\widehat{\lambda} = 0.515$

Table 23: Overview of best fitting distributions and parameter estimates

When we compare our findings to previous research done by Chan et. al. [4], who also found the Generalized Hyperbolic Distribution to be the best fit for Bitcoin and Litecoin, it is interesting to see that approximately one year later, these distributions still fit the data best.
The parameter estimates are slightly different, however: especially the value for $\widehat{\alpha}$ has decreased. The other four values have increased slightly, both for Bitcoin and Litecoin.

For the distributions, this means that the location has shifted somewhat to the right, and the tails have become a bit fatter. Also, a stronger positive skew is now present, indicating the right tail is fatter than the left one.

For Ripple, they found that the Normal Inverse Gaussian Distribution fitted the data best. Here, this turned out to be a very good fit as well (perhaps even better, depending on whether the Kolmogorov-Smirnov or the Chi-Square test gets most emphasis).

Short Conclusion

From the Kolmogorov-Smirnov statistic, the best fits for the returns of all four cryptocurrencies appear to be the Normal Inverse Gaussian and the Generalized Hyperbolic Distribution, which is consistent with what we saw in the distribution plots and QQ-plots earlier. For all four cryptocurrencies, the Generalized Hyperbolic Distribution appears to fit slightly better than the Normal Inverse Gaussian. Based on the Chi-Square statistic, the Generalized Hyperbolic Distribution turned out to fit the data best as well. We had to limit the number of bins to 5 to get a significant result, though.
It stood out that the Normal Inverse Gaussian Distribution appeared to fit Ripple and Ethereum better when using 10 bins, though.
Compared to previous research, the distributions for Bitcoin and Litecoin have changed slightly:

fatter tails along with a stronger positive skew than before, were observed.

3.4.3 Theoretical Value at Risk and Expected Shortfall

Using the parameters obtained in table 23, the theoretical values for the Value at Risk and the Expected Shortfall can now be obtained and compared to the empirical ones. This has only been done for the daily log returns of the exchange rates (not for the weekly, monthly or quarterly averages, as opposed to before).

Table 24 provides an overview of the empirical VaR losses we found earlier (provided in tables 4, 8, 12 and 16) for the daily log returns, whereas in table 25 the theoretical losses are presented based on the Generalized Hyperbolic Distribution with corresponding parameters for each of the currencies.

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
Bitcoin	-0.016	-0.039	-0.066	-0.122
Litecoin	-0.021	-0.052	-0.083	-0.164
Ripple	-0.029	-0.054	-0.078	-0.159
Ethereum	-0.034	-0.056	-0.091	-0.196

Table 24: Empirical daily Values at Risk (VaR: losses) for several values of α for all four cryptocurrencies

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
Bitcoin	-0.018	-0.038	-0.061	-0.120
Litecoin	-0.022	-0.048	-0.081	-0.177
Ripple	-0.027	-0.049	-0.079	-0.174
Ethereum	-0.033	-0.061	-0.091	-0.165

Table 25: Theoretical daily Values at Risk (VaR: losses) for several values of α for all four cryptocurrencies, based on the best fitting distributions

The same can be done for the Expected Shortfall: Table 26 provides an overview of the empirical VaR losses we found earlier (provided in tables 5, 10, 14 and 18) for the daily log returns, whereas in table 27 the theoretical losses are presented based on the Generalized Hyperbolic Distribution with corresponding parameters for each of the currencies.

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
Bitcoin	-0.051	-0.077	-0.104	-0.167
Litecoin	-0.068	-0.100	-0.136	-0.245
Ripple	-0.070	-0.100	-0.137	-0.251
Ethereum	-0.076	-0.109	-0.146	-0.249

Table 26: Empirical daily Expected Shortfall (ES: losses) for several values of α for all four cryptocurrencies

	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
Bitcoin	-0.052	-0.076	-0.101	-0.159
Litecoin	-0.068	-0.102	-0.143	-0.256
Ripple	-0.068	-0.100	-0.140	-0.254
Ethereum	-0.075	-0.105	-0.136	-0.215

Table 27: Theoretical daily Expected Shortfall (ES: losses) for several values of α for all four cryptocurrencies, based on the best fitting distributions

When we compare the theoretical Values at Risk to the empirical ones (tables 24 and 25), we notice that, for $\alpha = 0.2, \alpha = 0.1$, and $\alpha = 0.05$, the values are quite similar to each other for all four cryptocurrencies. However, as we go further into the tail, greater discrepancies are starting to show for some of the currencies. For Bitcoin, the Generalized Hyperbolic Distribution appears to fit quite well. For Litecoin

and Ripple, all percentiles except for the 0.01 percentile, seem to give rather similar values, whereas the 0.01 percentile appears to slightly overestimate the risk by 1.3% and 1.5% respectively. A similar statement can be made with respect to the Expected Shortfall for Litecoin, but for Ripple, the values for the Expected Shortfall are again very similar, with a maximal difference of 0.3%, as can be seen in tables 26 and 27.
This indicates that at the tip of the tail, the best fitting distribution for Litecoin does not seem to apply as well as for the other parts of the distribution: the tail is slightly less fat than the theoretical distribution implies.
For Ethereum, however, the reverse seems to be true: at the tip of the tail, the Generalized Hyperbolic Distribution appears to underestimate the Value at Risk as well as the Expected Shortfall, by about 3%. To further investigate behavior at the tail, Paretianity tests will be performed in the next section.

But first, for the profits, the empirical Values at Risk will also be compared to the theoretical ones. Table 28 shows the empirical VaR for the profits, whereas table 29 shows the theoretical VaR for the profits, based on the Generalized Hyperbolic Distributions with corresponding parameters for each of the currencies:

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
Bitcoin	0.025	0.045	0.063	0.111
Litecoin	0.026	0.061	0.091	0.201
Ripple	0.024	0.058	0.102	0.245
Ethereum	0.047	0.089	0.134	0.202

Table 28: Empirical daily Values at Risk (VaR: profits) for several values of α for all four cryptocurrencies

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
Bitcoin	0.023	0.044	0.068	0.128
Litecoin	0.026	0.058	0.099	0.218
Ripple	0.025	0.059	0.107	0.265
Ethereum	0.046	0.087	0.130	0.235

Table 29: Theoretical daily Values at Risk (VaR: profits) for several values of α for all four cryptocurrencies, based on the best fitting distributions

Again, the same can be done for the Expected Shortfall:
Table 30 shows the empirical ES for the profits, whereas table 31 shows the theoretical ES for the profits, based on the Generalized Hyperbolic Distributions with corresponding parameters for each of the currencies:

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
Bitcoin	0.054	0.074	0.095	0.145
Litecoin	0.082	0.125	0.173	0.297
Ripple	0.092	0.140	0.210	0.410
Ethereum	0.106	0.144	0.182	0.252

Table 30: Empirical daily Expected Shortfall (ES: profits) for several values of α for all four cryptocurrencies

	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.99$
Bitcoin	0.056	0.080	0.106	0.173
Litecoin	0.080	0.120	0.167	0.297
Ripple	0.089	0.140	0.202	0.384
Ethereum	0.105	0.147	0.190	0.294

Table 31: Theoretical daily Expected Shortfall (ES: profits) for several values of α for all four cryptocurrencies, based on the best fitting distributions

When we again compare the theoretical Values at Risk to the empirical ones (tables 28 and 29), this time for the profits, we notice that, for $\alpha = 0.8, \alpha = 0.9$ and $\alpha = 0.95$, we observe the same phenomenon as we did for the losses: the Generalized Hyperbolic Distributions appear to give rather accurate values for the potential gains at the beginning of the tail. However, at the tip of the tail at $\alpha = 0.99$, the potential gains appear to be overestimated based on the VaR, this time for all four cryptocurrencies.
When we look at the Expected Shortfall in tables 30 and 31, however, while the same conclusion can be made for Bitcoin and Ethereum, the gains appear to be represented rather accurately for Litecoin, however, and for Ripple they even appear to be underestimated by the Generalized Hyperbolic Distribution at the 0.99 percentile.
The gains for Bitcoin and Ethereum appear to be overestimated by 2.8% and 4.2% respectively, whereas the ones for Litecoin are represented accurately, and finally the ones for Ripple appear to be underestimated by 2.6%.

To further investigate behavior at the tail, Paretianity tests will be performed in the next section.

Short Conclusion

Overall, the Generalized Hyperbolic Distributions appear to fit the data quite well when we consider the Value at Risk and the Expected Shortfall and make comparisons between the empirical and the theoretical ones.
Slight discrepancies show, however, at the tips of the tails. For the losses, Bitcoin and Ripple appear to be modeled very well by the Generalized Hyperbolic Distributions fitted, but the risks for Litecoin are slightly overestimated and the risks for Ethereum are grossly underestimated.
When it comes to the profits, however, the gains for Litecoin appear to be captured really well by the Generalized Hyperbolic Distribution fitted, whereas the gains for Bitcoin and Ethereum appear to be grossly overestimated and the gains for Ripple appear to be grossly underestimated.
This is once again indicative of a fat right tail for Ripple.

The results are consistent with the QQ-plots that were obtained in section 3.4 for Bitcoin and the QQ-plots for the other cryptocurrencies that can be found in the appendix.
Overall, these results confirm that the far ends of the tails are not always captured well by the best fitting distributions, and that we need to be careful with our interpretations.
Whereas the Generalized Hyperbolic Distributions fit the data well, the far end of the tails need to be examined seperately in order to make conclusions with respect to the risks and the potential gains that can be made.

A closer look at the tails will be taken in the next section: Testing for Paretianity.

3.5 Testing for Paretianity

In this section, Bitcoin, Litecoin, Ripple and Ethereum will be tested for Paretianity, that is, Pareto Distributions will be fitted on both the tail for losses as well as the tail for profits, for all four cryptocurrencies.

3.5.1 Losses

First, the Bitcoin losses will be analyzed. Based on earlier results, like the excess kurtosis (a value > 3) that was obtained in section 3.1, table 2, indicating relatively many extreme values, as well as the negative skewness measure that was obtained, the expectancy is that fat tails will be observed, especially for the losses. Also, in section 3.3 we noticed that the losses were more extreme than the profits.

First, a Zipf plot is given for Bitcoin, based on the log losses of the daily exchange rates in figure 16:

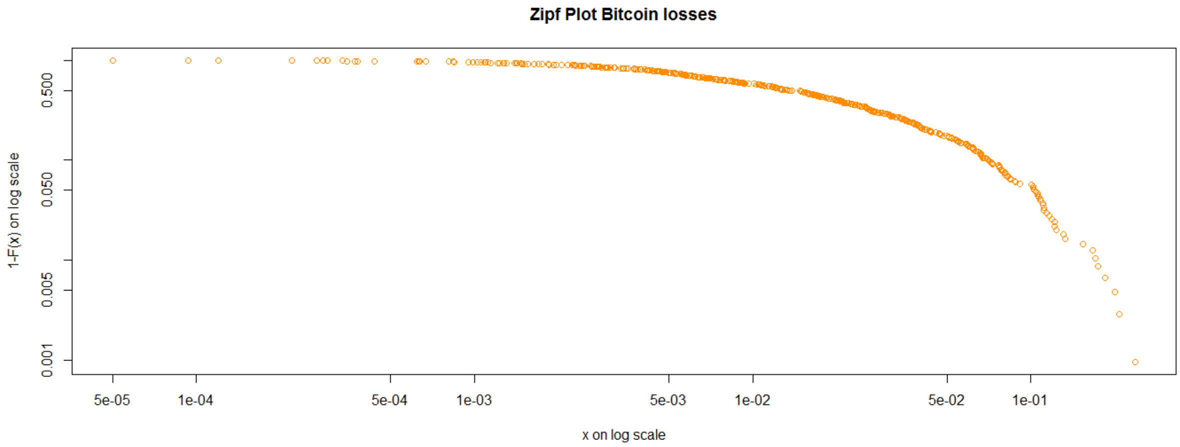

Figure 16: A Zipf plot for Bitcoin losses

Looking at this plot, it is visible from the shape of the graph that (at least a large part of) the data does not appear to be Pareto distributed.
Remember that in order to speak of a fat tail, a linear straight line is expected to show in the Zipf plot.

When we use R to obtain parameters for β and ξ, we obtain, using the maximum likelihood parameters, the values $\beta \approx 0.028$ (with a standard error of 0.003) and $\xi \approx 0.143$ (with a standard error of 0.075). Given these parameter estimates and standard errors that were obtained, the shape parameter ξ does not indicate an extremely fat tail here: at least the first 7 moments are expected to be finite, and therefore there is no reason to believe that the Bitcoin losses are Pareto distributed.

Taking into consideration only those losses that are greater than $\beta \approx 0.03$, we obtain the Zipf plot in figure 17, strengthening the hypothesis that Paretianity does not seem to be present for Bitcoin losses, as there is no negative straight line visible in the plot.

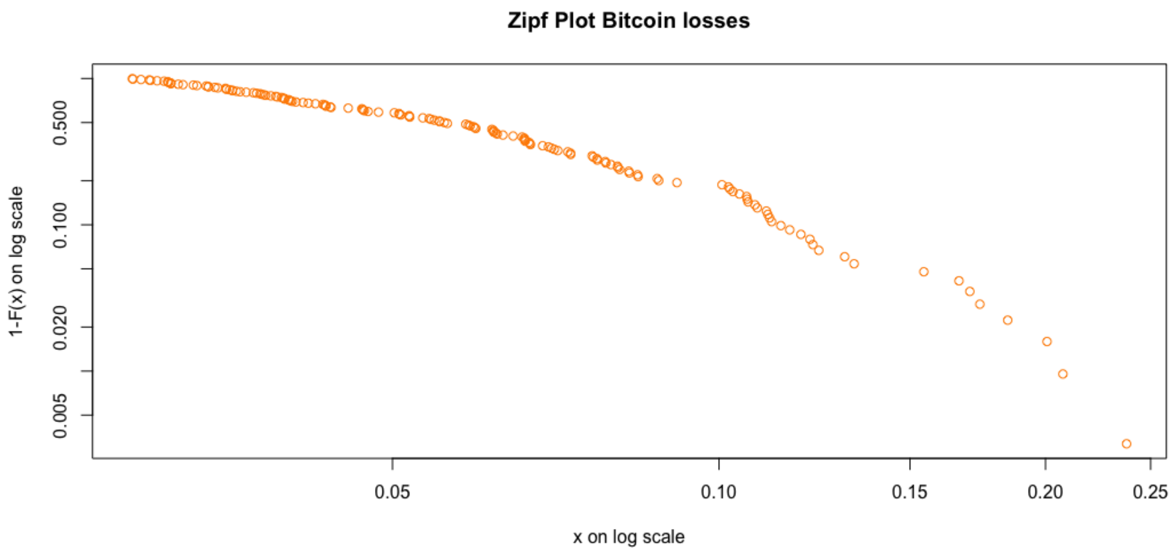

Figure 17: A Zipf plot for Bitcoin losses above threshold $\beta \approx 0.03$

It therefore appears that, even though extreme values were expected for Bitcoin losses, the values are not extreme enough to be modeled using a Pareto Distribution.

Also, when we compare this result to the one found by [21], who found values $\beta \approx 0.02$ and $\xi \approx 0.27$, we can conclude that the tail has become less heavy since 19 months ago, as we obtained a smaller value for the shape parameter.
If we now look at the other cryptocurrencies, we obtain for Litecoin the following Zipf plot in figure 18

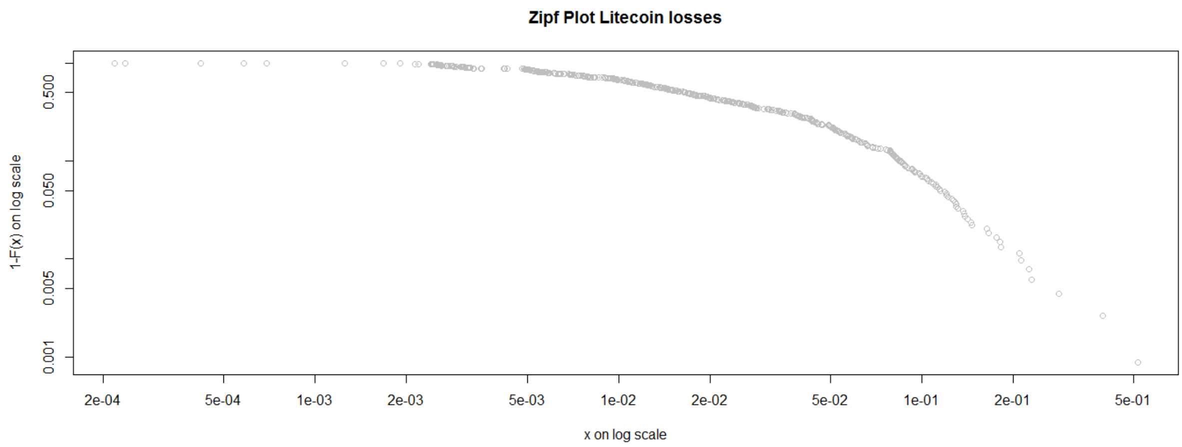

Figure 18: A Zipf plot for Litecoin

Based on earlier results for Litecoin in sections 3.1 and 3.3, we would expect Litecoin to be Pareto distributed. The level of kurtosis, for example, turned out to be relatively high, also compared to Bitcoin. Also, Litecoin came out second in terms of volatility in section 3.2 and the Values at Risk and values for the Expected Shortfall in section 3.3 were indicative of some extreme tail behavior as well.
Looking at this plot, the tail for the losses might seem to be Pareto distributed.

When we use `R` to obtain parameters for β and ξ, we obtain, using the maximum likelihood parameters, the values $\beta \approx 0.027$ (with a standard error of 0.002) and $\xi \approx 0.299$ (with a standard error of 0.075). Given the parameter estimates and standard errors, indicating at least an infinite fourth moment (and perhaps even an infinite third moment), Paretianity appears to be present in the tail for the losses.
It also stands out that the value obtained for the shape parameter is larger than the one for Bitcoin, therefore exhibiting heavier tail behavior than Bitcoin, which is consistent with earlier findings. Taking into consideration only those losses that are greater than $\beta \approx 0.03$, we obtain the Zipf plot in figure 19, strengthening the hypothesis that Paretianity might be present for Litecoin losses.

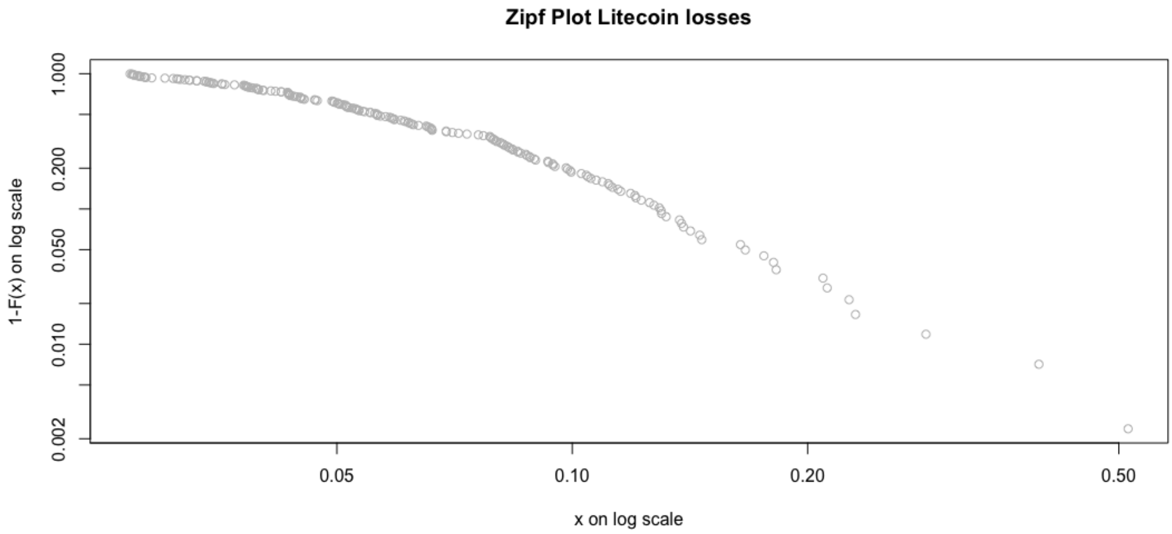

Figure 19: A Zipf plot for Litecoin losses above threshold $\beta \approx 0.03$

When we compare the results to those by [21], who found values $\beta \approx 0.03$ and $\xi \approx 0.23$, we observe that the tail for the Litecoin losses has become somewhat heavier in the last 19 months, as opposed to Bitcoin that became less heavy.

Then, moving on to Ripple, we obtain figure 20

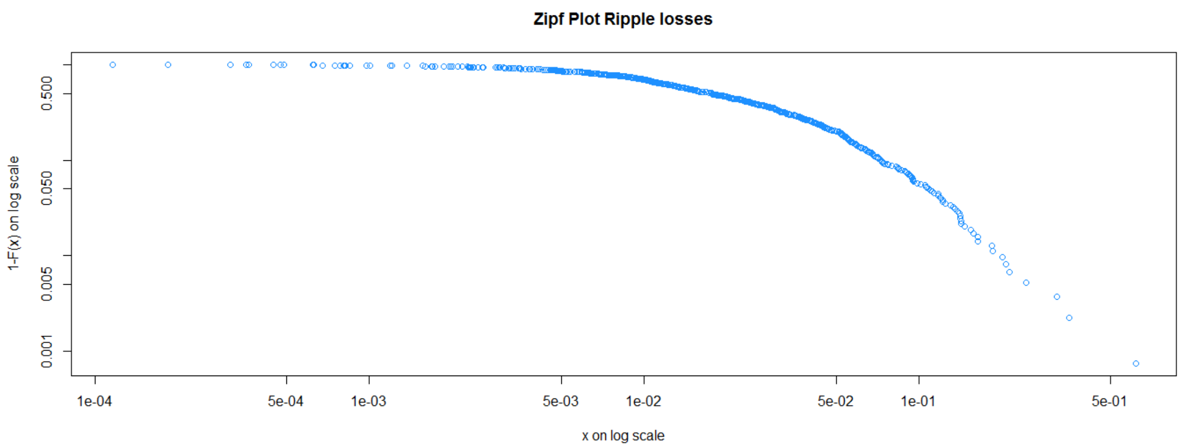

Figure 20: A Zipf plot for Ripple

Out of all four cryptocurrencies, Ripple stood out most with respect to its extreme returns. Remember we obtained for the log returns, a value of kurtosis of approximately 41.6, indicating the presence of extreme values (even compared to Bitcoin and Litecoin). Also, in section 3.3 we obtained that, for the losses, Ripple appeared to produce somewhat more extreme values than Litecoin. Therefore, considering Litecoin has a Paretian tail, the same might be the case for Ripple.
When we take a look at the Zipf plot in figure 20, we observe similar behavior as we did for Litecoin. If we then estimate the parameters once again, we obtain the maximum likelihood estimators $\beta \approx 0.024$ (with standard error of 0.002) and $\xi \approx 0.300$ (with a standard error of 0.065), similar to the parameters we obtained for Litecoin indeed, with infinite fourth and maybe even third moment. If we then zoom in to the part where Paretianity might possible be observed, we obtain figure 21,

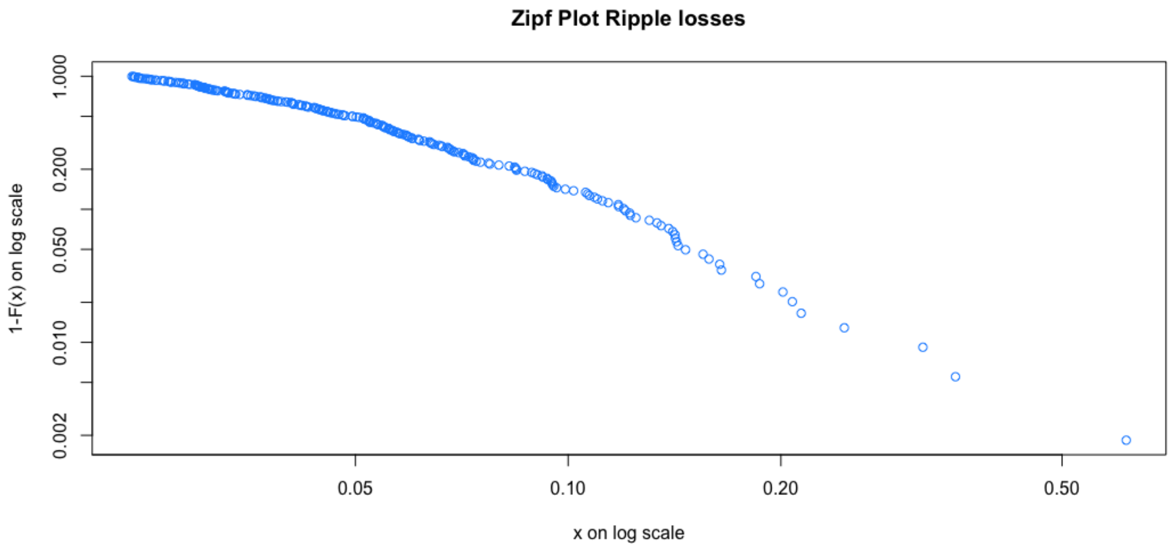

Figure 21: A Zipf plot for Ripple losses above threshold $\beta \approx 0.02$

and, as for Litecoin, the tail of the Ripple losses appears to follow a Pareto Distribution.

When we compare the results to those by [21], who found values $\beta \approx 0.03$ and $\xi \approx 0.26$, we observe that the tail for the Ripple losses has also become heavier in the last 19 months.
Finally, for Ethereum, we obtain figure 22:

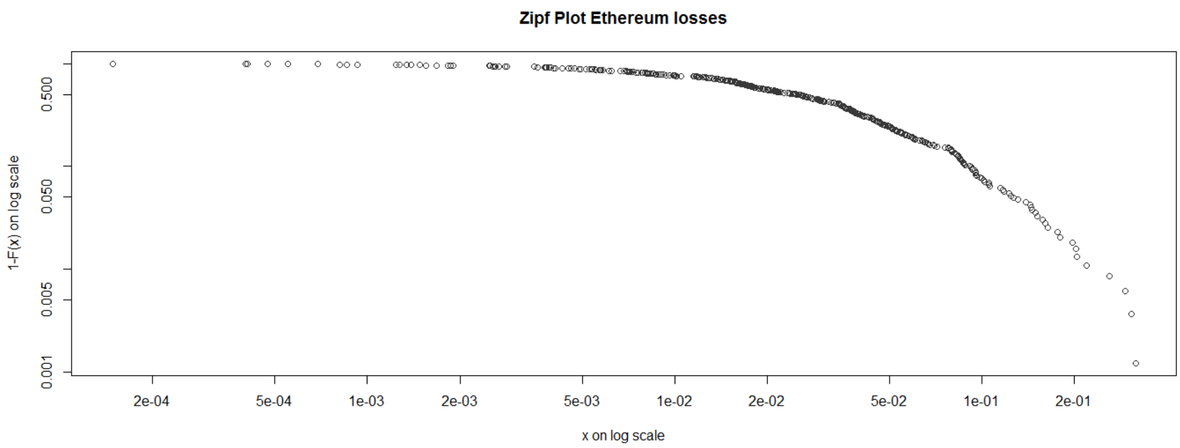

Figure 22: A Zipf plot for Ethereum

When we use R to obtain parameters for β and ξ, we obtain, using the maximum likelihood parameters, the values $\beta \approx 0.033$ (with a standard error of 0.003) and $\xi \approx 0.165$,(with a standard error of 0.066) indicating there is no Paretianity to speak of, i.e. no fat tail for the losses.

This is confirmed by a new Zipf plot in figure 23, zooming in at the far end of the tail for log losses greater than 0.03.

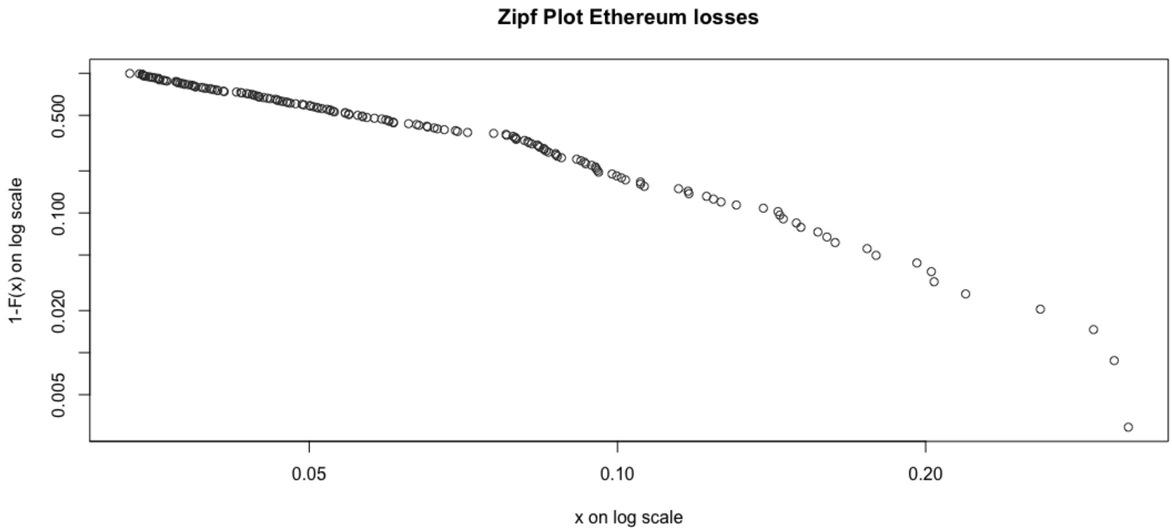

Figure 23: A Zipf plot for Ethereum losses above threshold $\beta \approx 0.03$

Unfortunately, the parameters obtained cannot be compared to previous research, as Ethereum was not included here.

To summarize, we have obtained the following parameters (and standard errors) for the Generalized Pareto Distribution for the log losses, using a threshold value of 0.01:

	β **(scale)**	ξ **(shape)**
Bitcoin	0.028 (0.001)	0.143 (0.075)
Litecoin	0.027 (0.002)	0.299 (0.075)
Ripple	0.024 (0.001)	0.300 (0.065)
Ethereum	0.033 (0.003)	0.165 (0.066)

Table 32: Losses: maximum likelihood estimates (with standard errors) for the Generalized Pareto Distribution

From table 32 we can conclude that the shape parameters for Litecoin and Ripple appear to be highest, and therefore the most extreme observations come from these two cryptocurrencies. For Litecoin and Ripple, the losses are also more extreme compared to two years ago, based on the results by [21].
The Bitcoin losses, however, appear to have become less extreme. For Ethereum, no comparison could be made.

3.5.2 Profits

For the profits, it is also interesting to examine the tail behavior of Bitcoin, Litecoin, Ripple and Ethereum. For Bitcoin, considering the slight negative skew obtained in section 3.1, and the fact that even the losses turned out not to be Paretian distributed, we expect the shape parameter for the profits to be even lower, and therefore the profits for Bitcoin are expected not to be Pareto distributed.
However, given the extreme value of Kurtosis for Ripple, and the fact that Ripple appeared to have the most significant positive skew, it is expected that the profits for Ripple will show some Paretianity. For the losses, we obtained $\xi \approx 0.3$ for Ripple, so for the profits we expect ξ to exceed this value at least.
We have, once again, used `R` to obtain maximum likelihood parameters for the Generalized Pareto Distribution, and obtained the following parameters (along with standard errors) for the four cryptocurrencies of interest in table 33:

	β (scale)	ξ (shape)
Bitcoin	0.026 (0.002)	0.068 (0.052)
Litecoin	0.031 (0.003)	0.356 (0.076)
Ripple	0.032 (0.003)	0.443 (0.078)
Ethereum	0.056 (0.005)	0.008 (0.060)

Table 33: Profits: maximum likelihood estimates for the Generalized Pareto Distribution

Figure 24 shows four Zipf plots corresponding to the tails of the profits for all four cryptocurrencies.

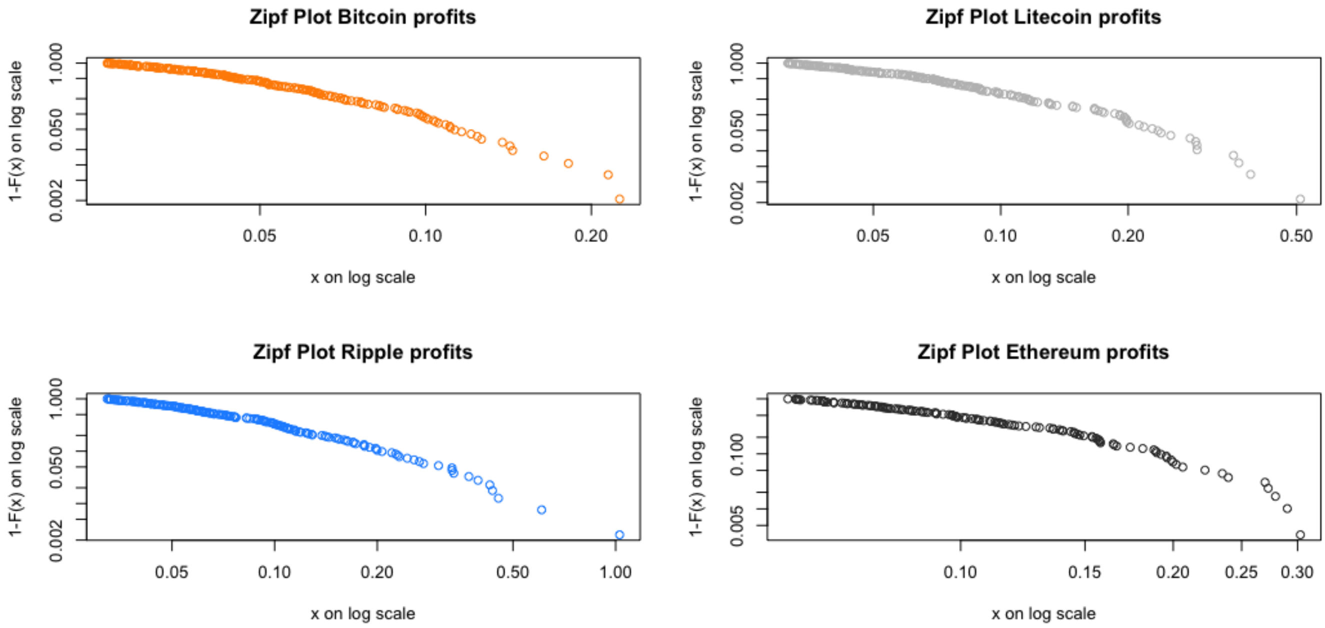

Figure 24: Zipf plots for profits above thresholds β provided in table 33

What stands out is that the shape parameters for Bitcoin and Ethereum are very low ($\xi \approx 0.07$ and $\xi \approx 0.01$ respectively). Based on these values, the tails are highly likely not Pareto distributed. For Litecoin and Ripple, however, the values that were obtained ($\xi \approx 0.36$ and $\xi \approx 0.44$ respectively), appear to be even higher than the values for the losses. Especially for Ripple, the shape parameter, combined with the standard error, is indicative of an extremely fat tail. This is consistent with earlier findings (the value for Kurtosis, but also the Value at Risk and Expected Shortfall values for the profits were indicative of extreme tail behavior).
For Litecoin, Paretianity is also suspected based on the values obtained combined with our findings from previous sections.

Unfortunately, the results obtained cannot be compared to previous research done by [21], as they only took into account the losses, not the profits.

Short conclusion For Litecoin and Ripple, Paretianity was expected based on the Descriptive Statistics obtained in section 3.1, and the Values at Risk and Expected Shortfall values obtained in section 3.3, and the values in table 32 and 33 confirm our expectations: for Ripple, fat tails were observed, especially for the profits but also for the losses. Also, for the Litecoin profits and losses, at least mildly fat tails appear to be present in the data.
It also stands out that, compared to the losses, the shape parameters for Bitcoin and Ethereum are lower for the profits, whereas for Litecoin and Ripple this appears to be the other way around.
Compared to previous research done by [21], the left tail for Bitcoin has become less fat, whereas

for Litecoin and Ripple, the tails have become fatter. For Ethereum, no comparison could be made. For the profits, no comparisons could be made either, unfortunately.

3.6 Maximum to Sum

In this section, Maximum to Sum ratios will be plotted against n to see whether our results from section 3.5 will be confirmed.
As explained in section 2.6, we are interested in the convergence of $R_n(p)$: we want $\lim_{n\to\infty} R_n(p) = 0$ to be able to assume that the $p-$th moment is finite.

- First, for Bitcoin, the following plot (25) is obtained, based on the absolute values of the log returns:

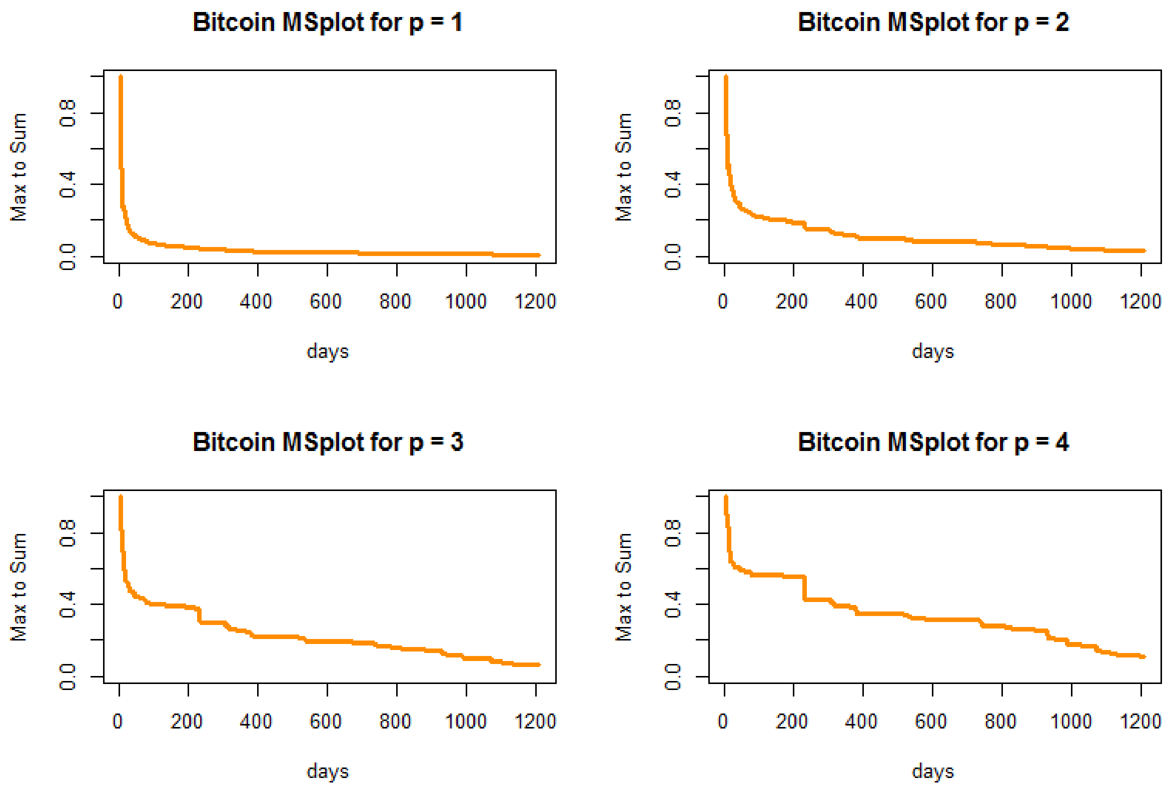

Figure 25: Maximum to Sum ratios for Bitcoin log returns, powers 1 to 4

From figure 25 we can conclude that the sequences appear to converge to 0 and therefore, the mean, variance, skewness and kurtosis appear to be finite.
This is in accordance with our findings from section 3.5, where we did not observe extremely heavy tails when testing for Paretianity.

- For Litecoin, there appears to be convergence for the first two moments $p = 1, 2$ (figure 26), and maybe for $p = 3$, although this is less clear from the (lower left) plot.

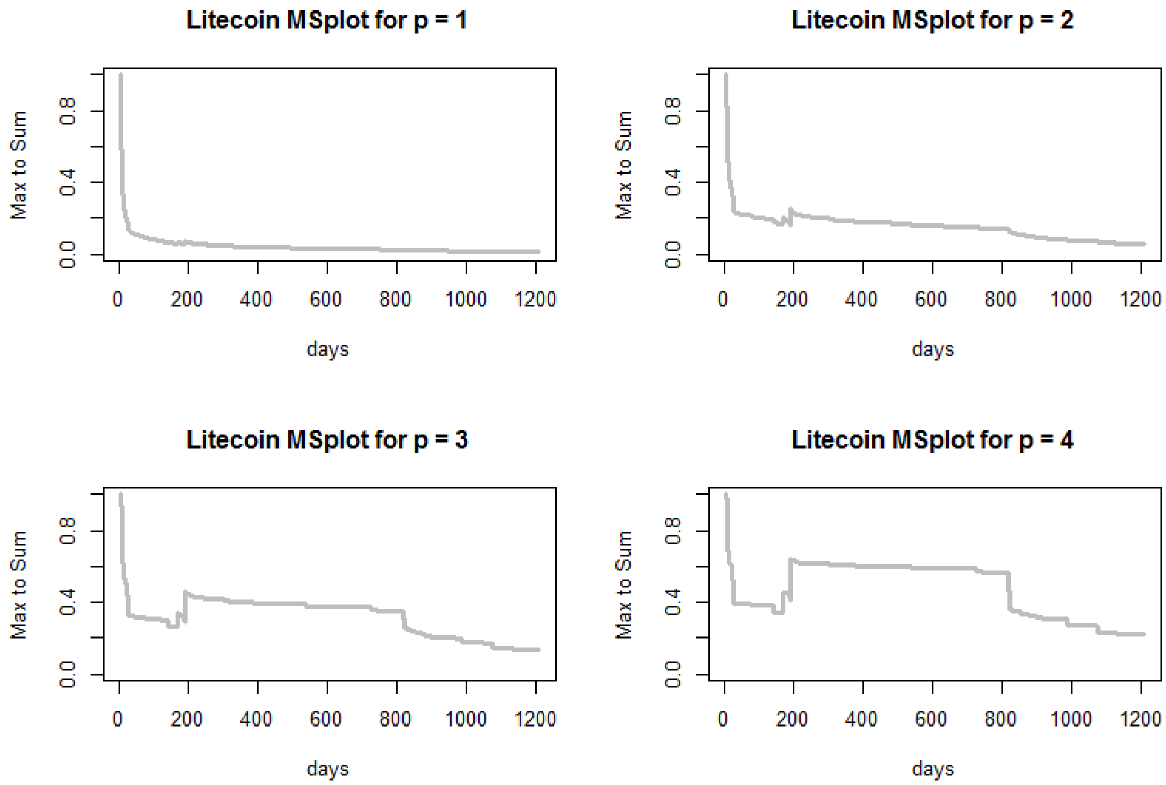

Figure 26: Maximum to Sum ratios for Litecoin log returns, powers 1 to 4

Based on the resulting plots for Litecoin in figure 26 and our findings from sections 3.1 and 3.5, it might be interesting to examine the losses and the profits seperately. Maximum to Sum plots for the losses and the profits can be found in figures 27 and 28 respectively:

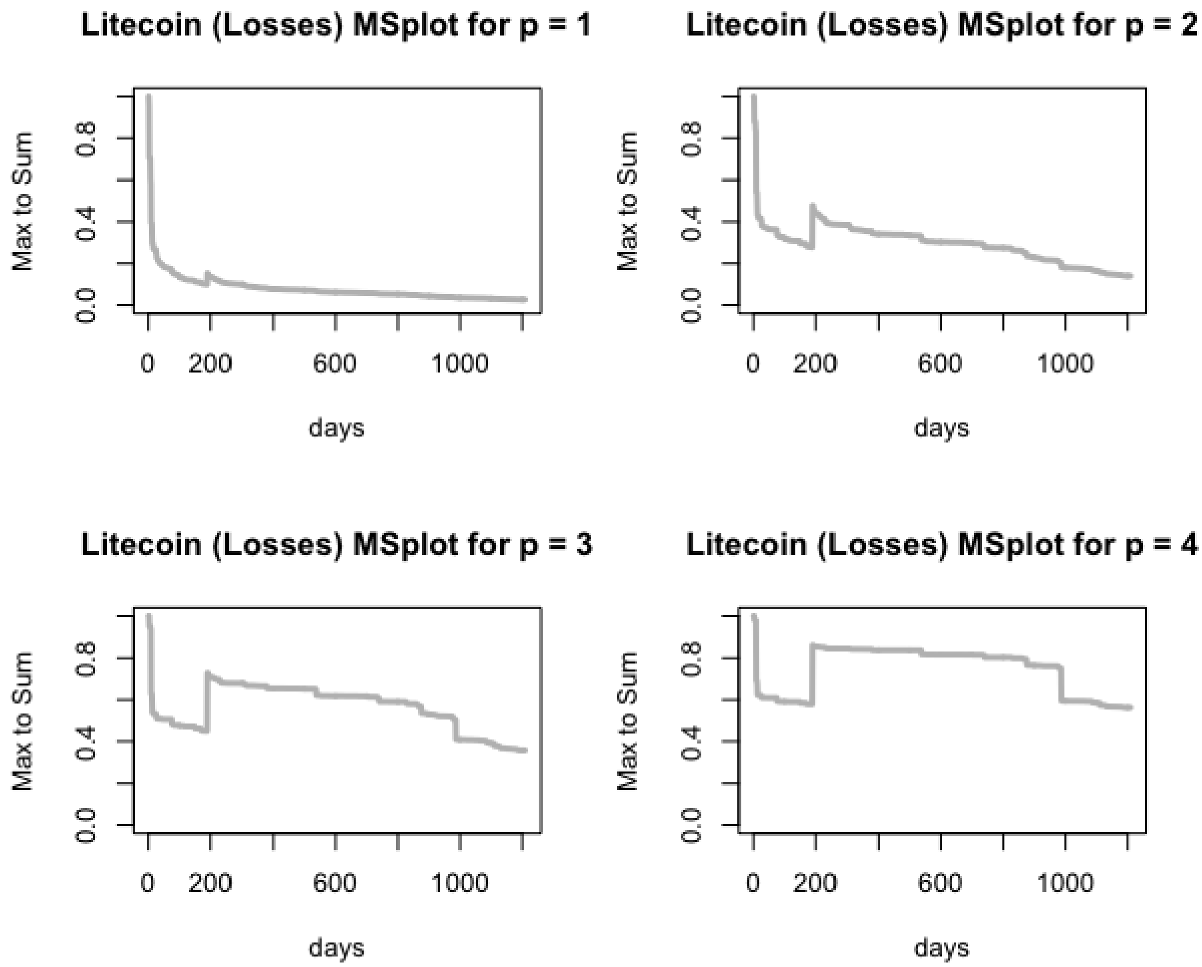

Figure 27: Maximum to Sum ratios for Litecoin losses, powers 1 to 4

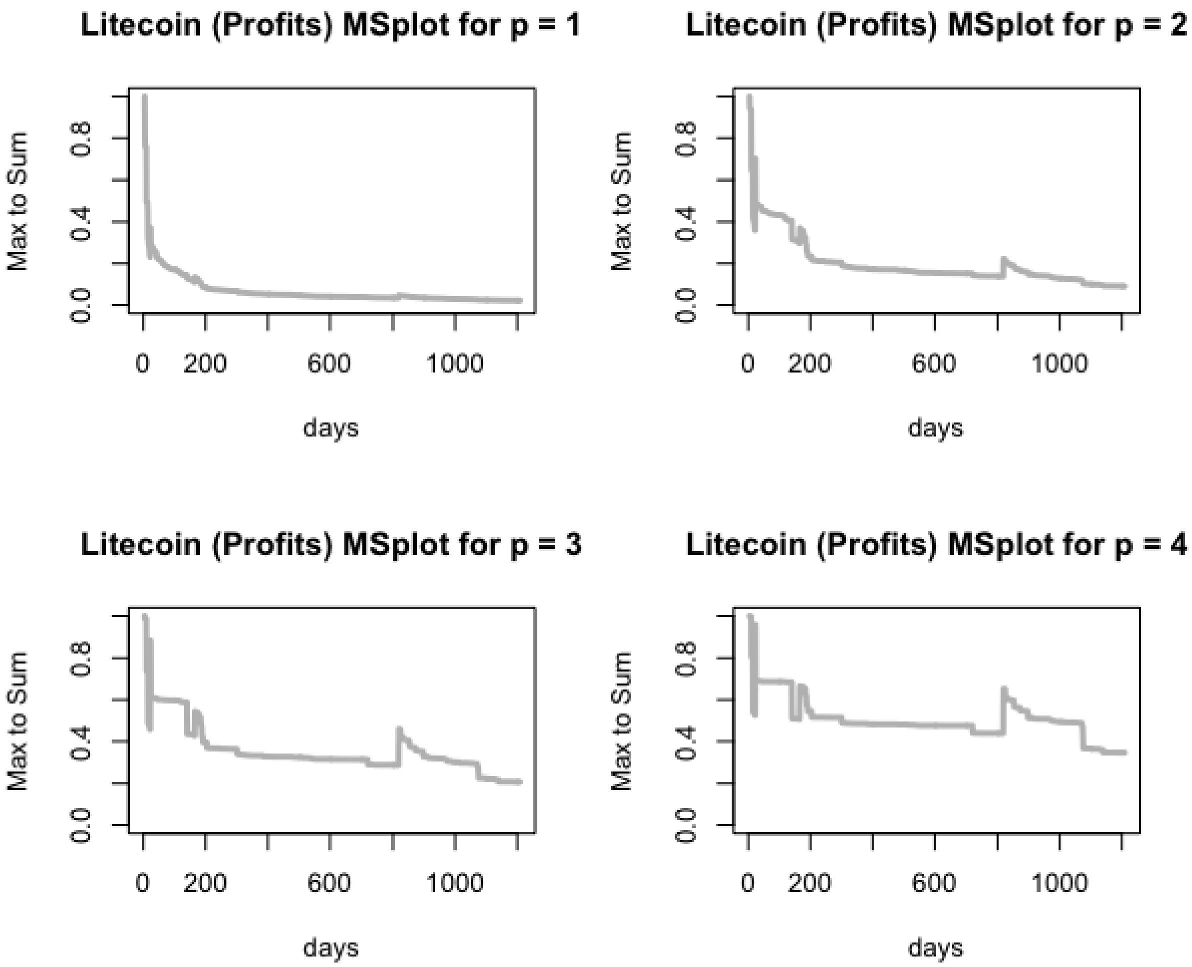

Figure 28: Maximum to Sum ratios for Litecoin profits, powers 1 to 4

By examining the losses and the profits seperately, it becomes especially clear that for $p = 3$ and $p = 4$ the sequences do not seem to converge: both the plots for the Litecoin losses as well as the ones for the Litecoin profits indicate that the data come from distributions with infinite third and fourth moment. This is consistent with the values we found for the skewness and kurtosis in table 2, equal to 0.71 and 13.964 respectively, which were relatively high compared to the ones for Bitcoin and Ethereum (but not Ripple), and is also consistent with our findings from section 3.5 where we assumed the Litecoin losses and profits appear to follow a Pareto Distribution.

- For Ripple (figure 29), we also observe that only the mean and variance (first and second moment) appear to be finite; for $p = 3$ and $p = 4$, the sequences do not converge to 0.

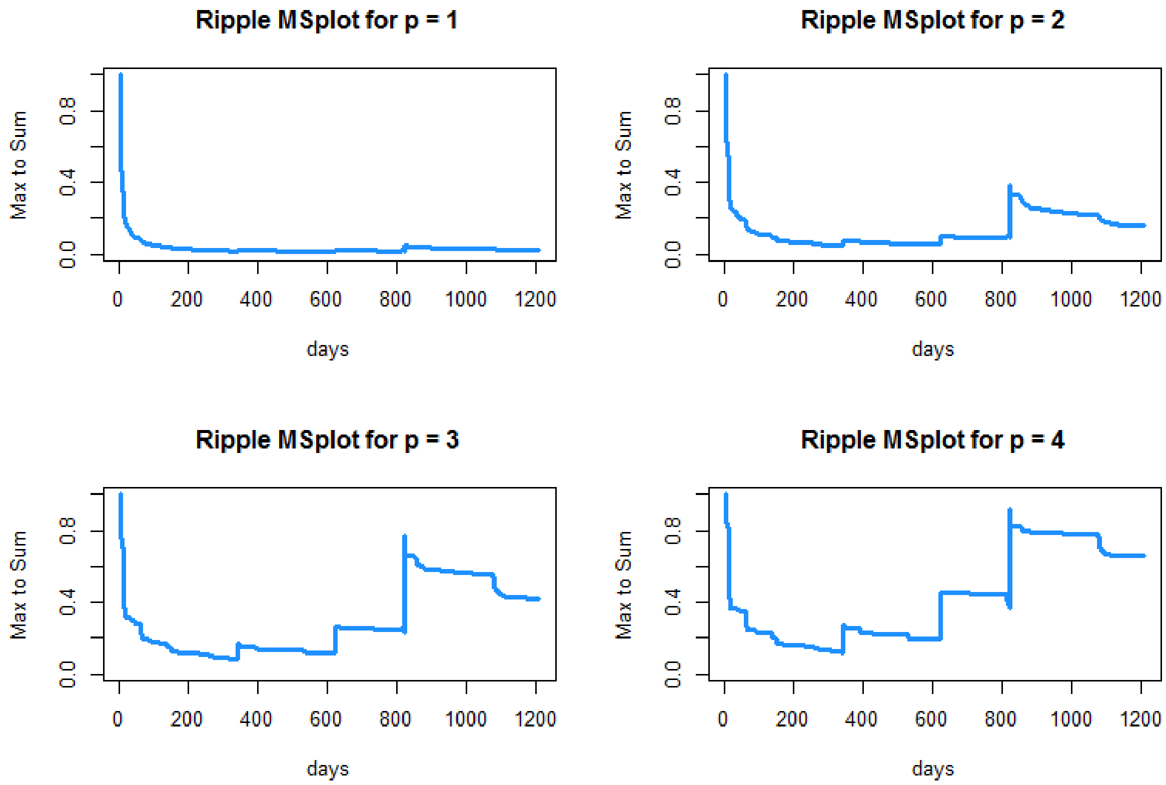

Figure 29: Maximum to Sum ratios for Ripple log returns, powers 1 to 4

In fact, based on the upper right plot of figure 29, even for $p = 2$, the convergence is not so clear. When we again examine the Ripple losses (figure 30) and Ripple profits (figure 31) seperately, it can be confirmed that both sides of the distribution produce extreme values and we are dealing with infinite third and fourth moments:

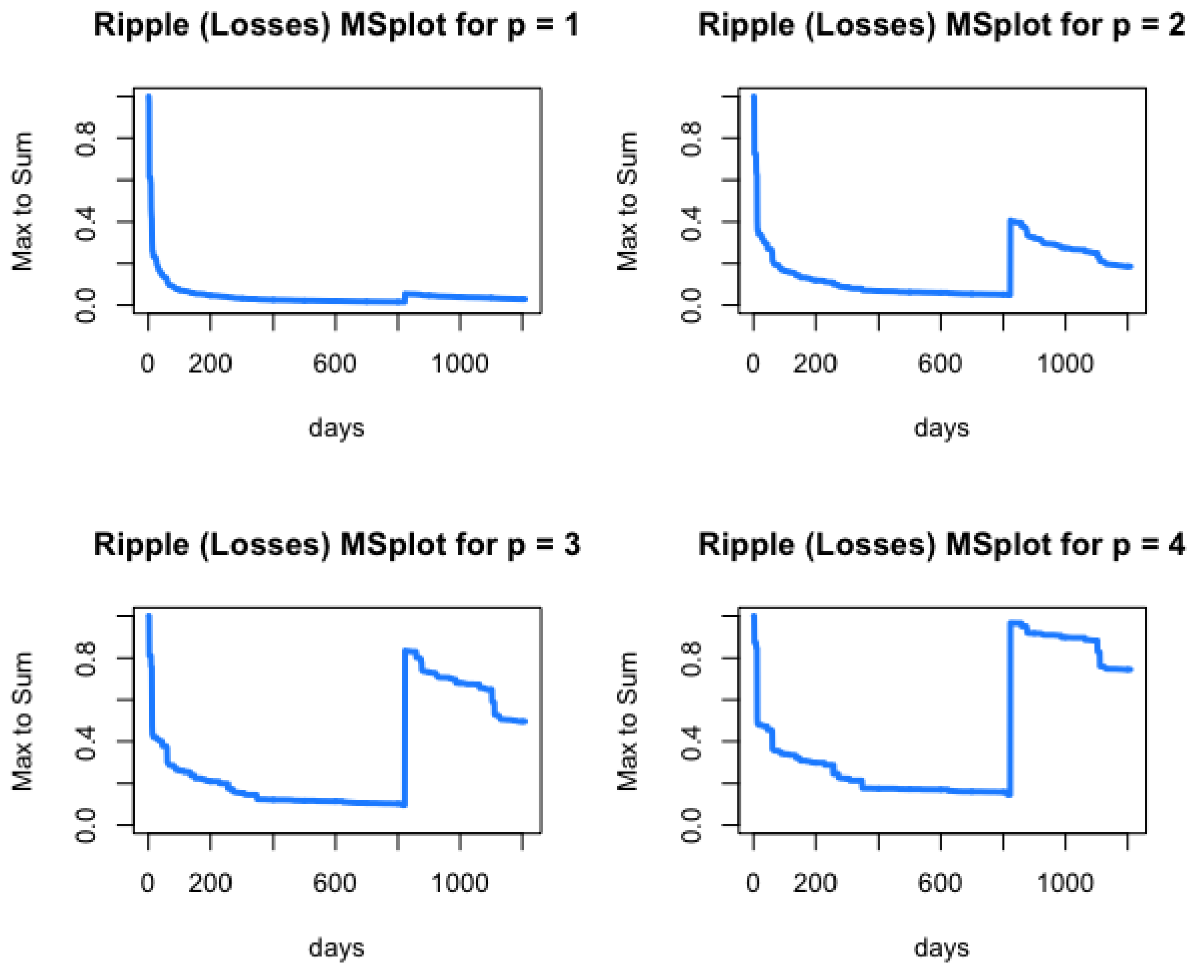

Figure 30: Maximum to Sum ratios for Ripple losses, powers 1 to 4

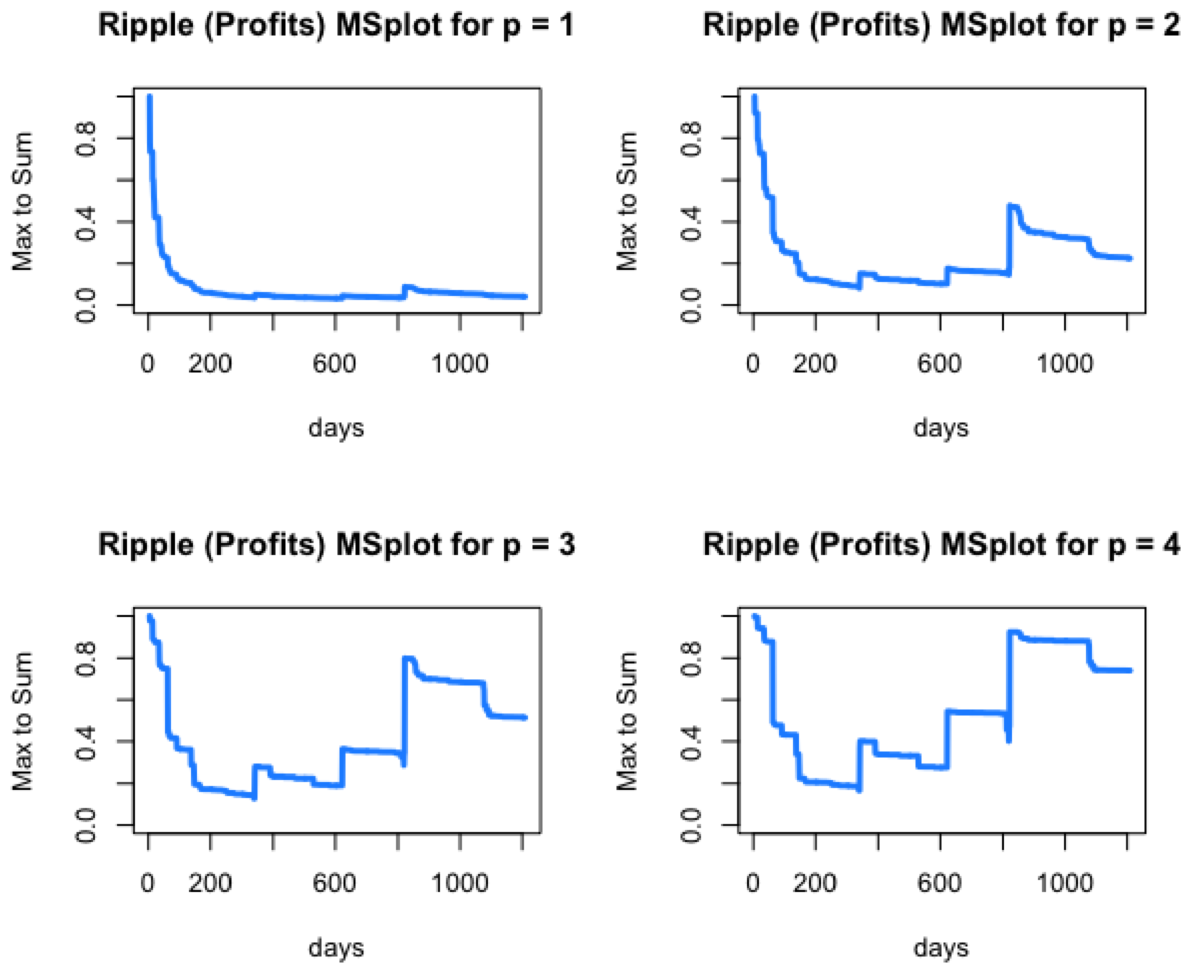

Figure 31: Maximum to Sum ratios for Ripple profits, powers 1 to 4

These results are consistent with earlier findings in section 3.1: the values we found for the skewness and kurtosis in table 2, equal to 3.132 and 41.637 respectively, which were extremely high compared to the values for the other three cryptocurrencies. Also, when we tested for Paretianity in section 3.5 we observed, especially for the profits, extremely fat tails for Ripple.

- Finally, Maximum to Sum plots are provided for Ethereum in (figure 32):

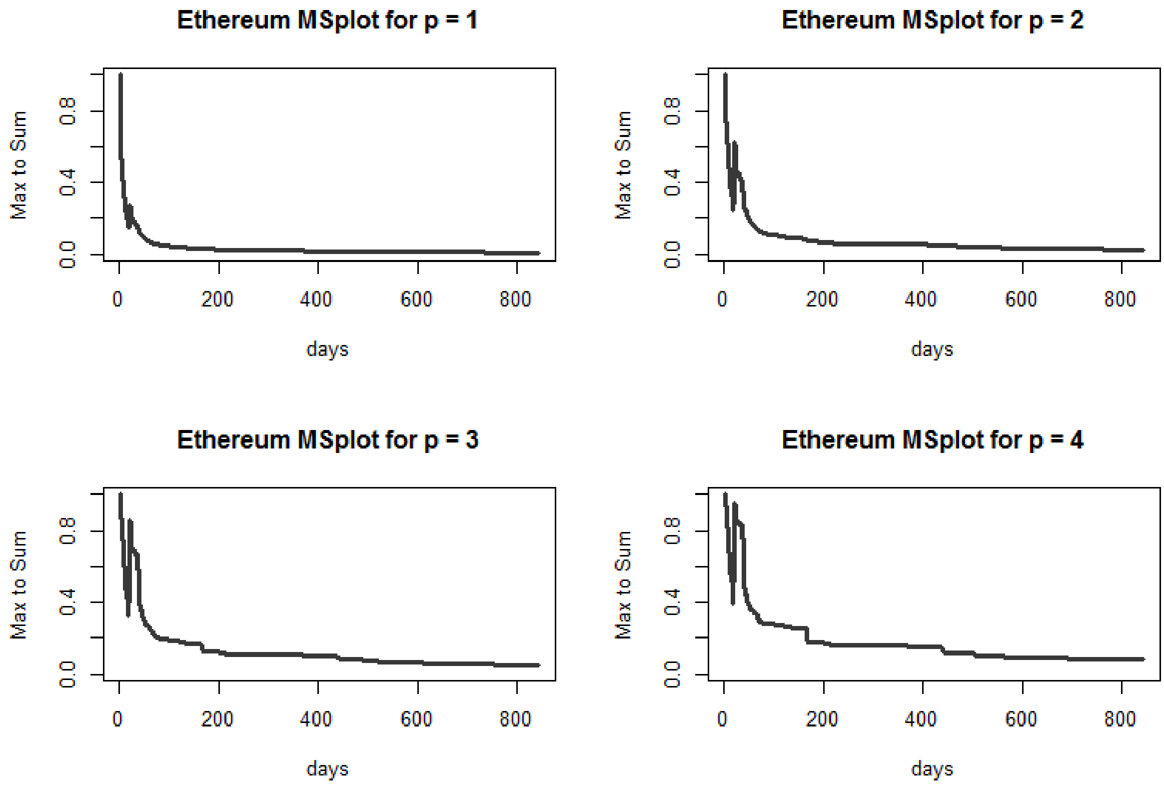

Figure 32: Maximum to Sum ratios for Ethereum log returns, powers 1 to 4

The plots in figure 32 confirm our expectations based on earlier findings: all moments appear to be finite, as there are no significant deviations from 0.
This is in accordance with the values we obtained in sections 3.1 for skewness and kurtosis, and also with the results we obtained when we tested the Ethereum losses and profits for Paretianity, as they did not appear to be Pareto distributed.

Short Conclusion

The plots in figures 25 to 32 are consistent with earlier findings: the distributions of the log returns for Bitcoin and Ethereum do not appear to have fat tails, but the ones for Litecoin and especially Ripple, do. Here, we see that the higher order moments ($p = 3, p = 4$) do not appear to converge. For Bitcoin and Ethereum, the sequences seem to converge for all p.
This is consistent with our findings in section 3.1, where we observed extreme skewness and kurtosis levels for Litecoin and Ripple, and in section 3.5 where we concluded that the Ripple tails appeared to follow a Pareto Distribution, and the ones for Litecoin appeared to as well.

4 Conclusion

With the emergence of Bitcoin in 2009, and more and more altcoins entering the market ever since, cryptocurrencies are very popular today. In this research paper, the behavior of the log returns of four cryptocurrencies, Bitcoin, Litecoin, Ripple and Ethereum, was investigated based on daily exchange rates from January 2015 until April 2018 for Bitcoin, Litecoin and Ethereum, and from January 2016 until April 2018 for Ethereum.
Here, a special emphasis was placed on the occurrence of extreme values.
Several research methods have been used in an attempt to answer the following research question:

How do the returns of Bitcoin and other altcoins behave over time, and what can we say about extreme values for losses and profits?

Comparisons have been made with respect to previous research conducted on this topic, to see how the behavior of the returns has changed since 2016, considering big changes have taken place in the cryptocurrency market since then (altcoins have taken up a significantly larger market share since then, and also, a huge market crash took place during the last few months of 2017, likely affecting investment strategies and therefore the behavior of returns).
Also, Ethereum was not included in most earlier studies, because of its relatively short existence, but it has been included in the current research, providing valuable information with respect to the topic, considering Ethereum ranks second place, right after Bitcoin, based on market capitalization.

Previous research conducted by Osterrieder et. al. [21] suggested that the returns of cryptocurrencies were extremely volatile, exhibiting heavier tail behavior than traditional fiat currencies.
In the current research, 19 months later, this still appears to be true.
Volatility results show an annual volatility of about 70% for Bitcoin, much higher than traditional stocks, although explicit comparisons between cryptocurrencies and stocks are not included in this research.
For Litecoin, Ripple and Ethereum, volatility levels were even higher.Especially Ethereum, the 'youngest' cryptocurrency of all four, turned out to have extreme annual volatility rates of around 130%.

Given the relatively high volatility rates, a large emphasis has been placed on the risks with respect to investments in cryptocurrencies, but of course, potential gains were also investigated.
Useful information with respect to the behavior of the returns was provided by, for example, summary statistics in section 3.1.
Here, the values of skewness and kurtosis provided the most information: Bitcoin appeared to be negatively skewed, but the other three cryptocurrencies showed positive skews, which was especially clear for Litecoin and Ripple. When not taking into account the most extreme values, Bitcoin also appeared to be positively skewed.
Here, it must also be kept in mind that negative skewness does not necessarily imply negative returns, as a significant area left of the mean can still be positive.
The kurtosis values made it clear that all four cryptocurrencies come from distributions with many extreme values compared to the Normal Distribution. For Litecoin and especially Ripple, the highest kurtosis values were obtained, indicating the possibility of extremely fat tails for these two cryptocurrencies. Combined with the positive skews, the expectation was that this was especially true for the right tails.
For Ripple, this was again confirmed by the analyses performed in section 3.3 taking into account the Value at Risk and Expected Shortfall. Here, Bitcoin turned out to be both the least risky and the least profitable, whereas Ripple, which always ranked first or second on the Expected Shortfall, could be considered the most risky and most profitable of all four cryptocurrencies.

For Bitcoin, the left tail has also become less fat compared to research by [21]. For Litecoin and Ripple, on the contrary, the left tails have become fatter.
This means that, when investing in cryptocurrencies, one can expect to lose a substantial portion of their investment during extreme events, when deciding to invest in Ripple or Litecoin (although it must be kept in mind that in general, the mean return is still positive). It can therefore be stated that, when

investing in Ripple or Litecoin, one is taking a relatively big risk compared to before and also compared to Bitcoin, which has, on the contrary, proven to have become less risky than it was before.
To make an explicit comparison, the losses for Bitcoin came down to 6.6% once every 20 days, with an expected loss of 10.4%, whereas for Ripple we obtained losses of 7.8% once every 20 days, with an expected loss of 13.7%.
Bitcoin can therefore be considered a 'safer' bet when it comes to making a decision with respect to what cryptocurrency to invest in, compared to for example Ripple. However, the expected loss is still substantially greater when investing in cryptocurrencies in general, compared to traditional stocks (although no explicit comparisons to any of those were made in this research).
For the profits, the difference between Bitcoin and Ripple was even greater, with gains of 6.3% and expected gains of 9.5% for Bitcoin as opposed to gains of 10.2% with expected gains of 21% for Ripple, based on the 0.95 percentile.
What this means is that, when deciding to invest in Bitcoin as opposed to Ripple, while this can be considered a safer bet, the expected gains are also substantially less. An investor that could afford to lose some money and is more interested in some quick, extreme gains, might therefore want to go for Ripple instead.

Consistent with the results aforementioned, when the tails were tested for Paretianity, the ones for Ripple and Litecoin appeared to be fat enough to be considered Pareto distributed. The gains for Ripple resulted in the highest value for the shape parameter of $\xi \approx 0.44$ with a standard deviation of approximately 0.08), which is consistent with the highest value for kurtosis that was obtained for Ripple, as well as the earlier results with respect to the Value at Risk.
This suggests that, out of the four cryptocurrencies under investigation, Ripple did exhibit tail behavior as heavy as initially expected, especially for the profits, together with Litecoin, which turned out to have the second fattest tail of all four.
For investors, this means that, based on the data from the past three years, extremely high profits are to be expected when investing in Ripple and Litecoin.
The losses turned out to be slightly less extreme but might still be considered Pareto distributed for Ripple and Litecoin.
Bitcoin and Ethereum appeared not to follow a Pareto Distribution, not for the losses nor the profits: these did not exhibit extreme tail behavior.

Another important aspect of this research was the fitting of distributions on the returns of Bitcoin, Litecoin, Ripple and Ethereum. As it turned out, the returns for all four cryptocurrencies appeared to be best modeled by the Generalized Hyperbolic Distribution. Visualizations appeared to be very useful here in order to provide some intuition with respect to the best fitting distributions, which were later confirmed by the Kolmogorov-Smirnov and Chi-Square statistics. The Kolmogorov-Smirnov statistic did, however, tend to overestimate how good the fit really was; for the Chi-Square statistic, the results were somewhat less convincing, although they were still statistically significant. Also, when more than five bins were used, the empirical and theoretical distributions were not similar according to the Chi-Square statistic.

Apart from the fact that Generalized Hyperbolic Distributions appeared to fit the data best for all four cryptocurrencies, it became clear from the plots that the areas where the distributions sometimes did not make a good fit, were indeed the tails, as was expected.
This is in accordance with the results found when comparing the theoretical and empirical distributions with respect to the Value at Risk and Expected Shortfall for the losses: at the far ends of the tails (considering quantile 0.01), the theoretical distribution for Litecoin appeared to overestimate the risks at the far end of the tail, whereas the one for Ethereum appeared to grossly underestimate the risks by more than 3%. For Bitcoin and Ripple, however, the Generalized Hyperbolic Distribution appeared to fit the data quite accurately, even at the tails. For the gains, a different conclusion was made: the Generalized Hyperbolic Distribution appeared to overestimate the gains for Bitcoin and Ethereum, whereas the gains for Ripple were in fact, underestimated (which is again consistent with the extremely fat right tail for Ripple).

To summarize, the behavior of the log returns appeared to be best described by Generalized Hyperbolic Distributions. The tails, however, needed to be considered seperately, as the distributions did not

always appear to fit the tails so well for all four cryptocurrencies.
When the tails were tested for Paretianity, the ones for Bitcoin and Ethereum turned out to be less heavy than expected, but for Ripple and Litecoin, the left tails were heavier than expected based on previous research, and the right tails turned out to be even heavier.

For the investor, there is no straight answer when it comes to deciding on the 'best' cryptocurrency to invest in. This really depends on your risk profile: all results considered, Bitcoin might be seen as a relatively 'safe' bet compared to the other three cryptocurrencies, being the least volatile of all four and exhibiting the smallest values when it comes to the Value at Risk and the Expected Shortfall for the left tail, representing the losses, while on the long run still showing dramatic increases in value, as was clear from the time series plots.
However, an investor that could afford to lose some money on occasion, and is interested in some short-term, extreme gains, might decide to invest in Ripple, Litecoin, or even Ethereum, which have also shown tremendous gains in the past.
A word of advice, however, is to never invest more than you can afford to lose, in any of the cryptocurrencies, as nothing is certain and extreme losses are not uncommon.

Shortcomings, possible explanations and/or solutions for the results just described, will be addressed in section 5, the Discussion section, along with future recommendations for further research.

5 Discussion

As summarized in section 4, the log returns are best described by Generalized Hyperbolic Distributions, apart from the far end of the tails.
The tails for Bitcoin and Ethereum did, however, not appear to be Pareto distributed: they were less extreme than expected. For Litecoin and Ripple, Paretianity is suspected to be present in the data, especially for the right tails: the profits.
In this section, shortcomings of the methods used will be addressed. When possible, solutions or ideas for future research are discussed.

First of all, results obtained in the current research, were often compared to results from previous research, out of interest to see how things have changed since approximately two years ago.
Given the fact that a lot has happened in the past years, like an extreme crash in the market and a shift in market share where altcoins have taken up a much larger share, as explained in the Introduction (section 1), we were interested to see how this influenced the results. However, we cannot be certain that differences that were obtained, can be attributed to these specific events. Market forces are generally very complex and therefore observed changes could be attributed to many other factors as well.

Other general comments on the current research can be made with respect to the comparisons between Bitcoin, Litecoin, Ripple and Ethereum.
First of all, for Ethereum, less data was used, simply because less data is available due to its emergence only late 2015.
Comparisons could therefore have been easier to make if the exact same time intervals were used (so if we, for example, used data for Bitcoin, Litecoin and Ripple starting from 2016 as well).
Also, it would have been interesting to make comparisons with regular fiat currencies, like the EURO, commodities or equities, when making statements with respect to, for example, volatility.
Doing so would have given an even better impression of the volatility of cryptocurrencies in general, as opposed to only comparing them to each other.
Finally, the four cryptocurrencies under investigation are all considered high-caps (given their positions with respect to market capitalization, which are 1,2,3 and 5 for Bitcoin, Ethereum, Ripple and Litecoin, respectively (on 18 June 2018; note that the positions frequently change).
It would also have been interesting to include one of the cryptocurrencies that are considered to be low-caps (with a market capitalization of, for example, below 500 million USD).
For the current research, however, cryptocurrencies were selected based on the amount of data available, and it is not so easy to select a low-cap cryptocurrency that is currently considered to be promising with respect to its technical applications, for which there is also sufficient data available to perform useful analyses.

In section 3.2, volatility was analyzed for all four cryptocurrencies. Here, the daily log returns were considered, but it would also have been interesting to analyze weekly, or even monthly volatility rates. This would eliminate day-to-day fluctuations and therefore give a more general impression of the long term volatilities. In future research, this could be taken into account.
On the other hand, it would also be interesting to consider shorter time intervals instead of longer ones. In the current research, data was used based on the daily exchange rates obtained from [3]. One of the main advantages when it comes to cryptocurrency analyses, is, however, the fact that a lot of data is available, and even exchange rates with time intervals of as short as one minute, could be obtained if wanted. It would then be interesting to 'zoom in' to a specific period in time, for example to the beginning of the large market crash in December 2017, and analyse the volatility and the risk during that period in time using time intervals of only minutes.
Such analyses could provide useful information with respect to the risks in future crises, or for example for companies where purchases can be made using cryptocurrencies.
Also, one of the results that was obtained, was that the empirical and theoretical distributions differed somewhat with respect to the Value at Risk and Expected Shortfall: at the beginnings of the tails, the theoretical distributions, which were Generalized Hyperbolic Distributions, (slightly) underestimated the losses and also the profits, but at the end of the tails, both the losses as well as the risks were mostly

overestimated (apart from the ones for Bitcoin and Ethereum, where it was the other way around for the losses). It would be interesting to further investigate the reason as to why this phenomenon occurs. Perhaps if a certain percentage of the most extreme values were to be eliminated from the data, such that the outliers would not influence the theoretical distribution so much, a different result would arise. Then again, it is, after all, the extreme values we're most interested in, and therefore, for now, the decision was made not to exclude any data.

In section 3.4, theoretical distributions were fitted to the daily log returns.
The goodness of fit tests that were used to determine how well each of the distributions described the behavior of the returns, were the Kolmogorov-Smirnov and the Chi-Square statistics.
There are of course, many other ways to determine whether a certain distribution fits the data well. The authors of [4], for example, used, apart from the Kolmogorov-Smirnov statistic, several information criteria, such as the Akaike information criterion and the Bayesian information criterion. In future research, distributions could be fitted more thoroughly by making use of criteria such as these.

Also, what stood out was that for the Chi-Square test, when making use of 5 bins, the results were consistent with the Kolmogorov-Smirnov statistics: the Generalized Hyperbolic Distributions were the best fits. However, when 10 bins were used, the results were surprising: for Ripple and Ethereum, the Normal Inverse Gaussian suddenly appeared to fit the data better than the Generalized Hyperbolic Distribution.
It would be interesting to further look into that, as it might have something to do with the difference in nature between the Kolmogorov-Smirnov and the Chi-Square tests. Whereas the Kolmogorov-Smirnov test only looks at the maximal distance D between the empirical and theoretical distribution, the Chi-squared test takes into account all data points, by comparing the number of data points per bin.
Finally, in future research it would be interesting to perform the same analyses that were used in this research, but instead of making use of the entire dataset, consisting of data for over three years, this could be handled differently. For example, a distinction could be made between so-called 'bull' markets (characterized by optimism and expectations of continuously strong results) and 'bear' markets (characterized by falling prices and pessimism), when the market is into a phase of almost constant increase or decrease, respectively. It would be interesting to see how volatilities and risk measures differ when distinguishing between these two types of markets.

Also, one might wonder how relevant data from three years ago really is today, anyway.
Perhaps it would be better to look at more recent data, or maybe put more weight on recent data than on data further in the past, by using for example exponentially decaying weights.

Finally, and perhaps the most interesting suggestion for future research, would be to investigate how to diversify between the four cryptocurrencies investigated, and consider the effects diversification would have on for example the Value at Risk and the Expected Shortfall. After all, one of the main advantages of risk management are the gains that can be obtained via diversification. Correlations between the four cryptocurrencies as well as copulas could be taken into account here as two measures of dependence between the cryptocurrencies. It would be interesting to investigate what linear combination of the cryptocurrencies might result in the least risky portfolio, or the most profitable one.
All in all, a lot of topics were investigated, but there are still many improvements and ideas for future research concerning the returns of cryptocurrencies.

References

[1] Bandorff-Nielsen, O. (1977). Exponentially decreasing distributions for the logarithm of particle size. *Proceedings of the Royal Society of London Series A Mathematical and Physical Sciences, 353.*

[2] Bitcoin Project (2018).
URL= `https://bitcoin.org/en` .

[3] CoinMarketCap (2018). Cryptocurrency market capitalizations.
URL= `https://coinmarketcap.com/`.

[4] Chan, S., Chu, J., Nadarajah, S. and Osterrieder, J. (2017). A Statistical Analysis of Cryptocurrencies. *Journal of Risk and Financial Management,10.*

[5] Cirillo, P. (2013). Are your data really Pareto distributed? *Physica A: Statistical Mechanics and its Applications,392(23):5947-5962.*

[6] Clauset, A., Shalizi, C.R. and Newman, M.E.J. (2009). Power-law distributions in empirical data. *JSIAM Review, 51(4).*

[7] Embrechts, P., Kluppelberg, C. and Mikosch, T. (1996). *Modelling Extremal Events,6.* Springer-Verlag.

[8] Ethereum Project (2018).
URL=`https://ethereum.org`

[9] Gosset, W.S. (1908). The probable error of a mean. *Biometrika,6.*

[10] Guo, L. and Li, X.J. (2017). Risk Analysis of Cryptocurrency as an Alternative Asset Class. *Applied Quantitative Finance, Statistics and Computing.*

[11] Hull, J. (2015). *Risk Management and Financial Institutions, 4.* Wiley.

[12] Kinateder, H. (2015. What drives tail risk in aggregate European equity markets? *Journal of Risk Finance,16*

[13] Kinateder, H. (2016). Basel II versus III- A comparative assessment of minimum capital requirements for internal model approaches. *Journal of Risk,18*

[14] Kleinbaum, D.G., Klein, M. (2012). *Survival analysis: A Self-learning text.*

[15] Kolmogorov, A. (1933). Sulla determinazione empirica di una legge di distribuzione. *Giorno dell'Instituto Italiano degli Attuari,4.*

[16] Laplace, P. (1774). Mémoire sur la probabilité des causes par les évènements *Mémoires de l'Academie Royale des Sciences Presentés par Divers Savan,6.*

[17] Litecoin Project (2018). Litecoin-Litecoin Wiki.
URL= `https://litecoin.info/Litecoin`

[18] Mandelbrot, B.B. (1963). The variation of certain speculative prices. *Journal of Business, 36.*

[19] Nakamoto, S. (2009). Bitcoin: A Peer-to-Peer Electronic Cash System.

[20] Noron, R.M. (1984). The Double Exponential Distribution: Using Calculus to Find a Maximum Likelihood Estimator. *American Statistical Association, 38(2).*

[21] Osterrieder, J., Lorenz, J. and Strika, M. (2016). Bitcoin and Cryptocurrencies- not for the faint-hearted. *Advanced Risk & Portfolio Management Paper.*

[22] Osterrieder, J. and Lorenz, J. (2017). A Statistical Risk Assessment of Bitcoin and its Extreme Tail Behavior. *Annals of Financial Economics,12.*

[23] Pagan, A. (1996). The econometrics of financial markets. *Journal of Risk and Financial Management,3.*

[24] Pareto, V. (1896). La courbe de la répartition de la richesse. Reprinted in *Rivista di Politica Economica, 87 (1997).*

[25] Ripple (2018). Welcome to Ripple.
URL=`https://ripple.com`

[26] Snedecor, G. and Cochran, W.G. (1989). Statistical Methods, *8, Iowa State University Press.*

[27] Szczepanski, M. (Members' Research Service, 2014). Bitcoin: Market, Economics and Regulation. *European Parliamentary Research Service.*

[28] Wilk, M.B. and Gnanadesikan, R. (1968). Probability plotting methods for the analysis of data. *Biometrika,55.*

6 Appendix

6.1 Datasets

Data that was used, was retrieved from [3], more specifically:

Bitcoin:
https://coinmarketcap.com/currencies/bitcoin/historical-data/
For Bitcoin, data was obtained from January 1, 2015 until April 26, 2018.

Litecoin:
https://coinmarketcap.com/currencies/litecoin/historical-data/
For Litecoin, data was obtained from January 1, 2015 until April 26, 2018.

Ripple:
https://coinmarketcap.com/currencies/ripple/historical-data/
For Ripple, data was obtained from January 1, 2015 until April 26, 2018.

Ethereum:
https://coinmarketcap.com/currencies/ethereum/historical-data/
For Ethereum, data was obtained from January 1, 2015 until April 26, 2018.

From these, the daily closing prices were used as daily exchange rates.

6.2 Plots

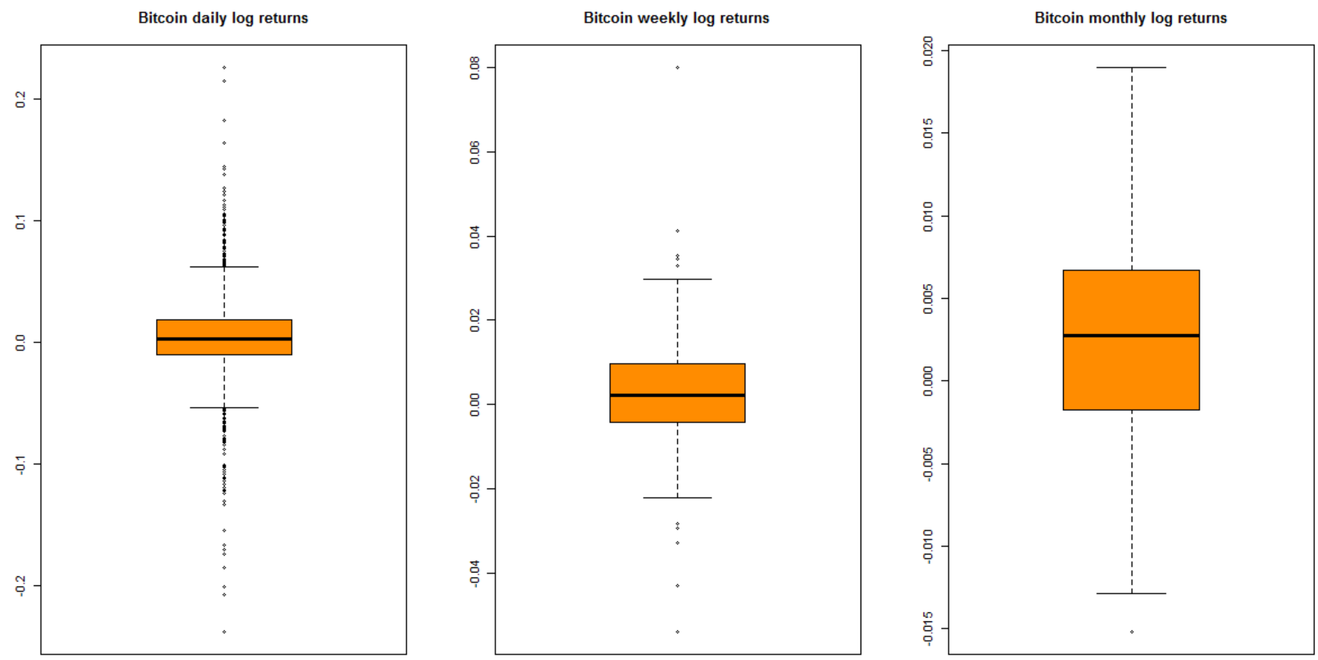

Figure 33: Boxplots based on log returns for Bitcoin

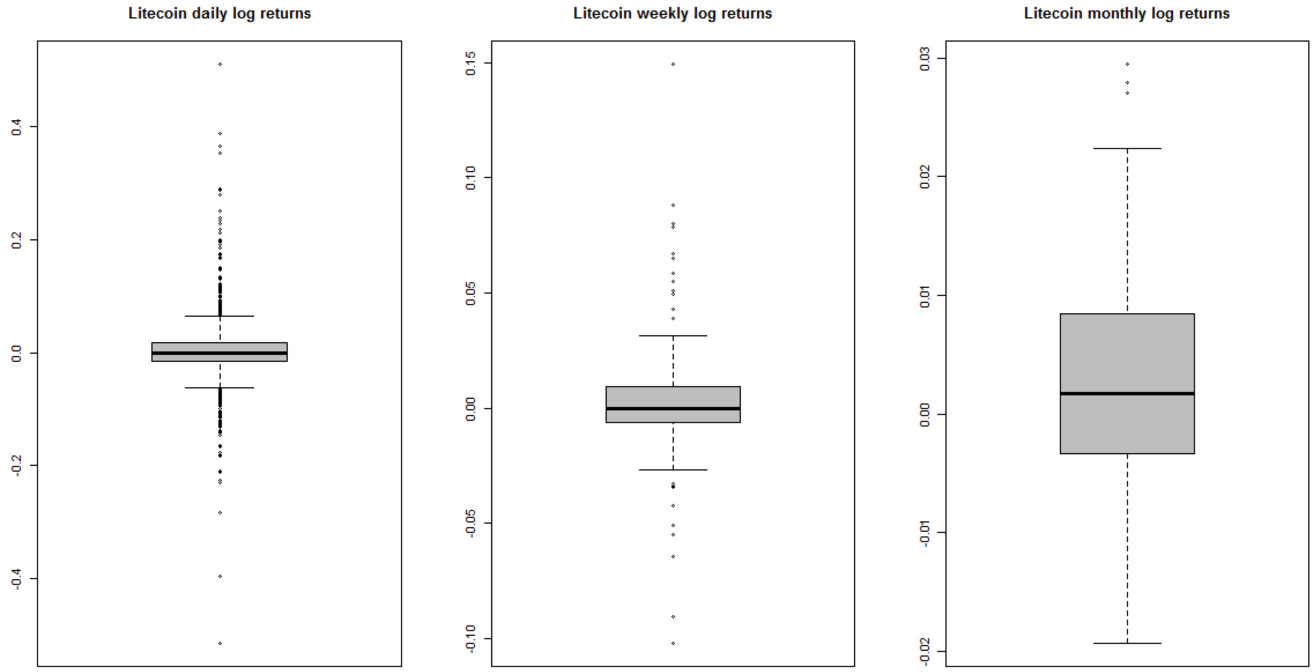

Figure 34: Boxplots based on log returns for Litecoin

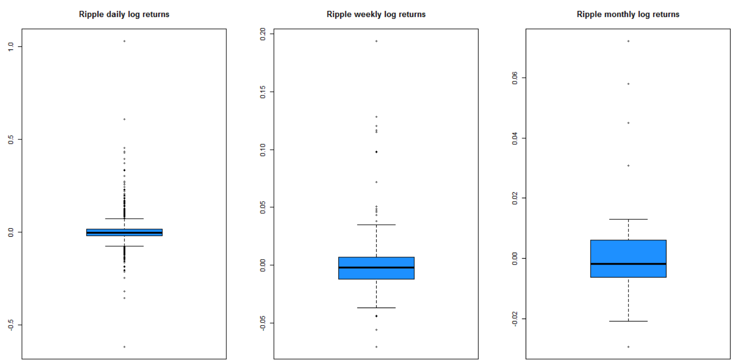

Figure 35: Boxplots based on log returns for Ripple

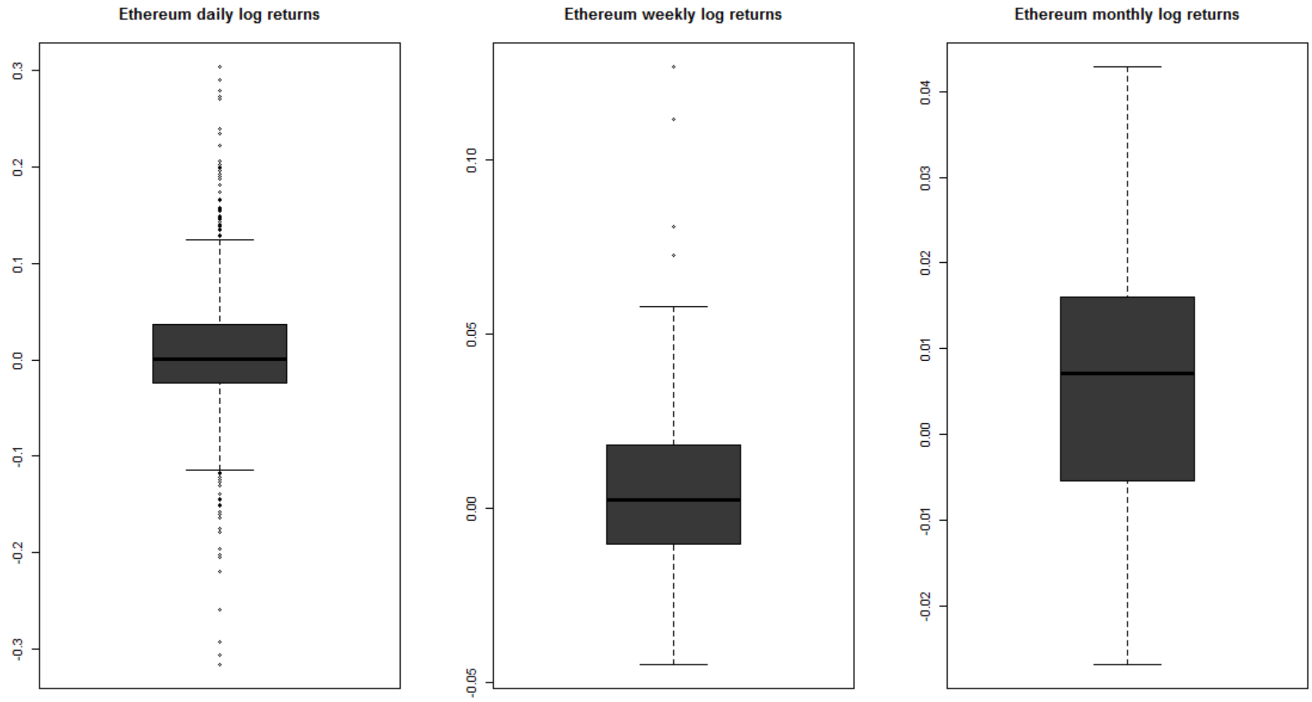

Figure 36: Boxplots based on log returns for Ethereum

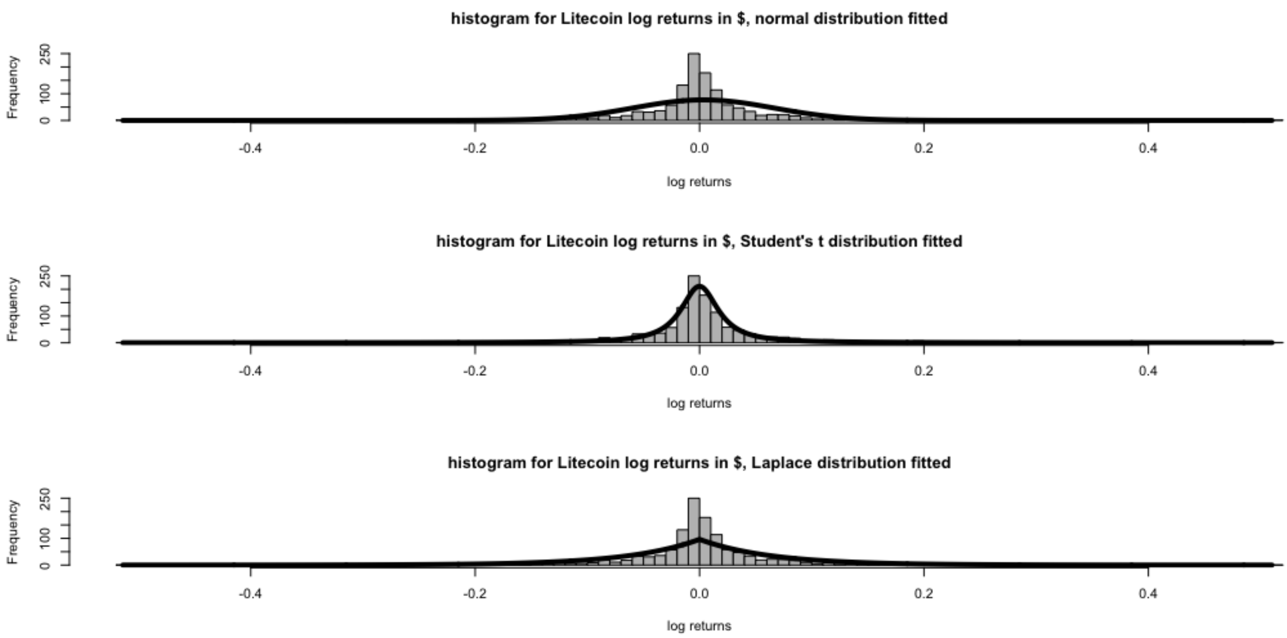

Figure 37: Fitting of distributions on Litecoin log returns

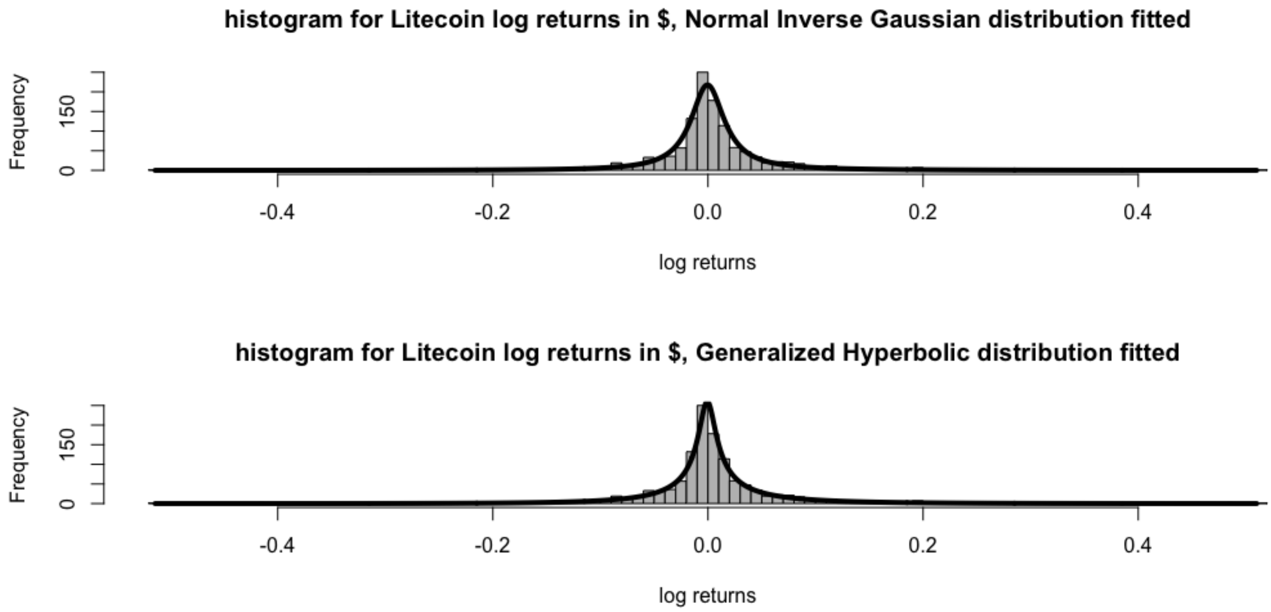

Figure 38: Fitting of distributions on Litecoin log returns

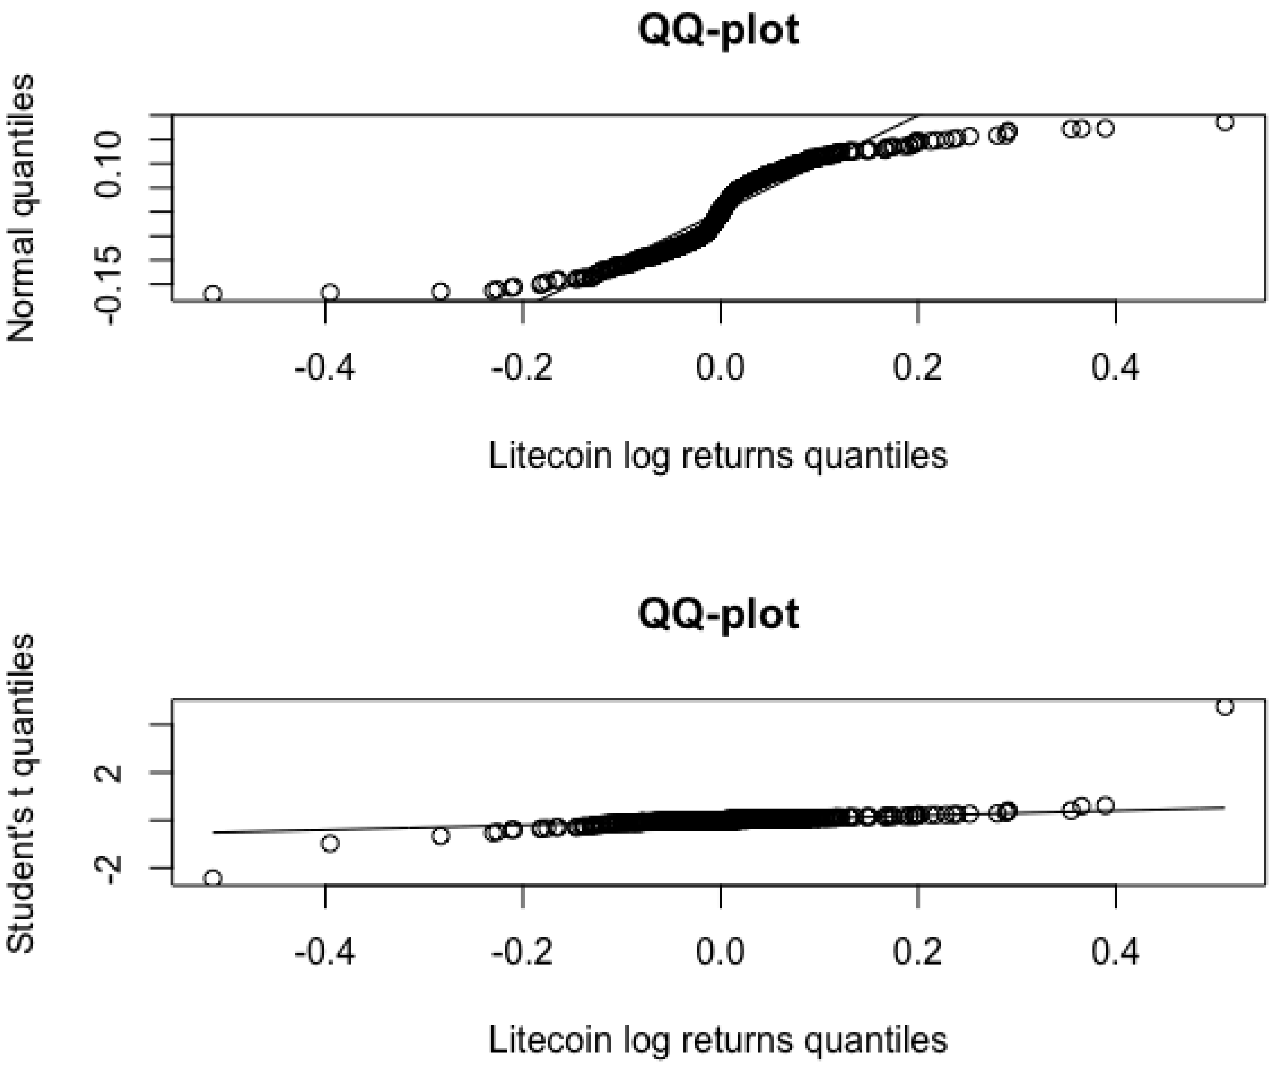

Figure 39: Corresponding QQ-plots for distribution fits on Litecoin log returns

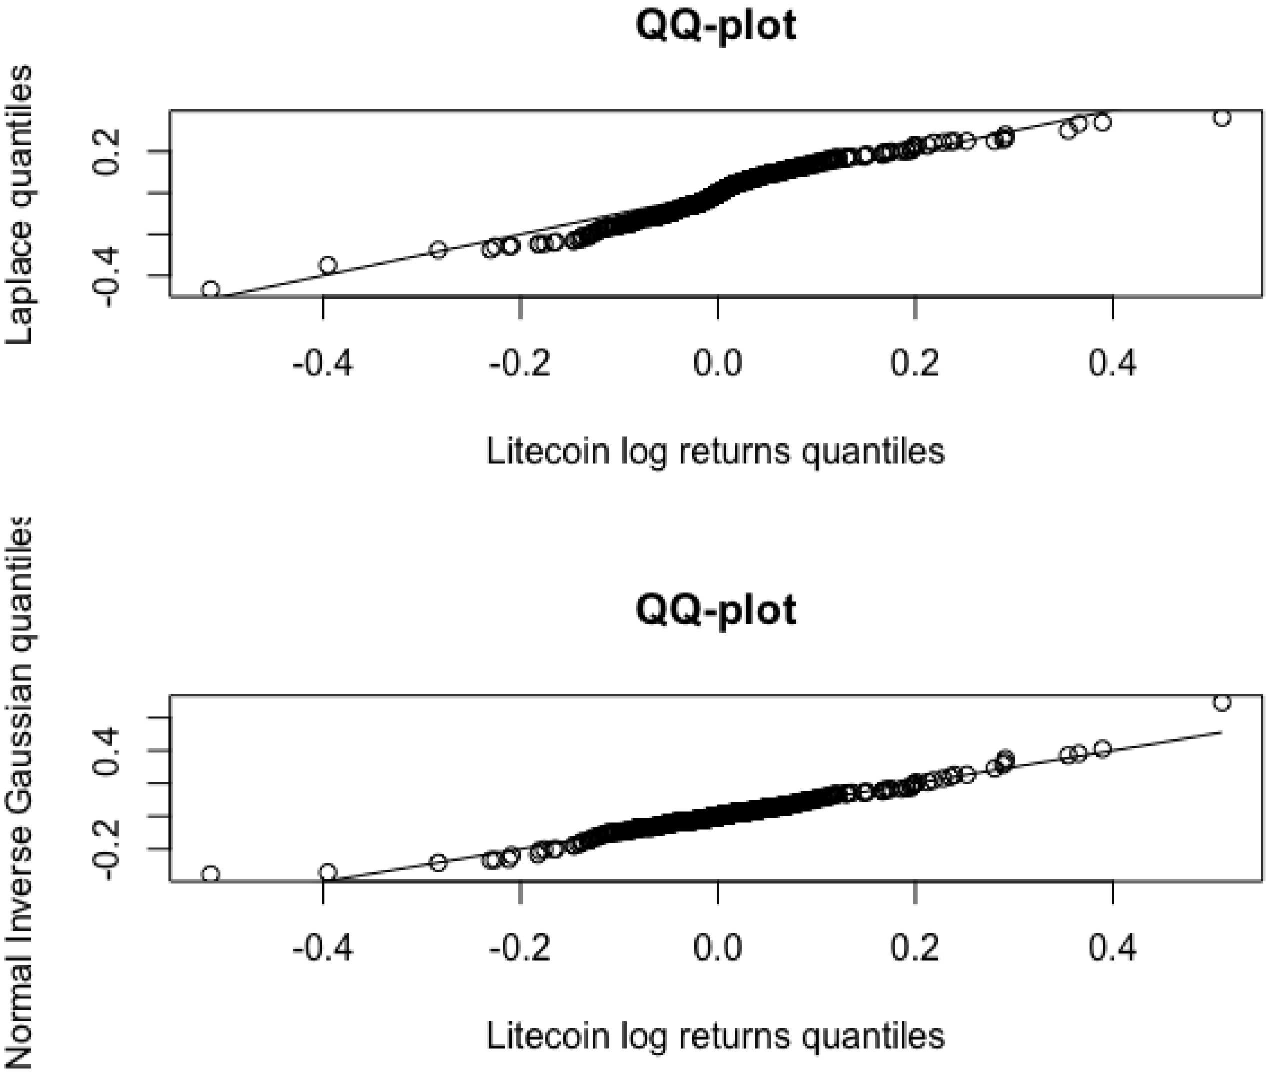

Figure 40: Corresponding QQ-plots for distribution fits on Litecoin log returns

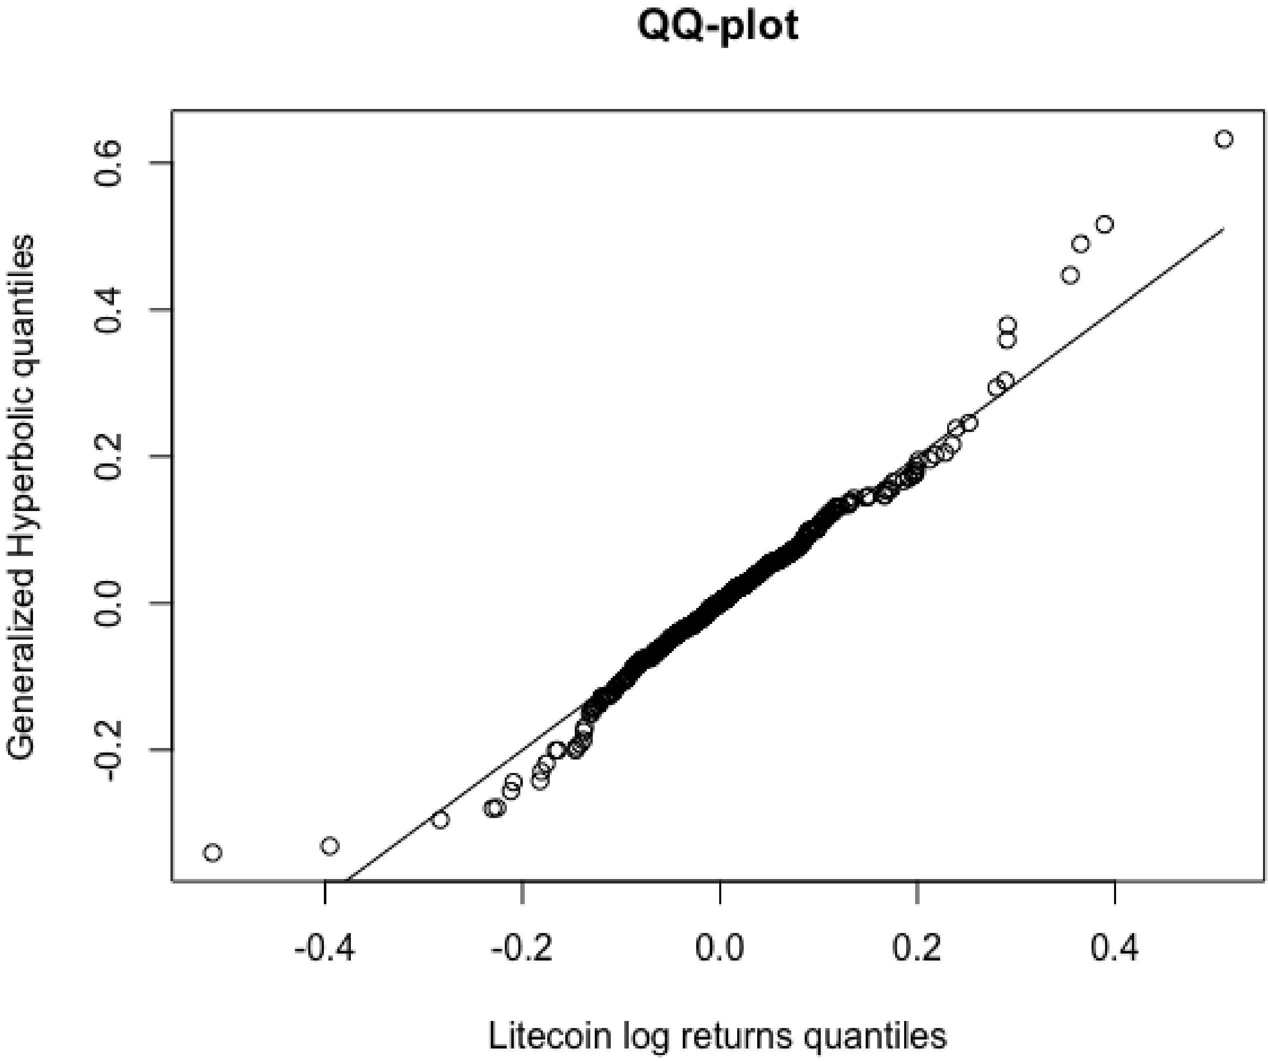

Figure 41: Corresponding QQ-plots for distribution fit on Litecoin log returns

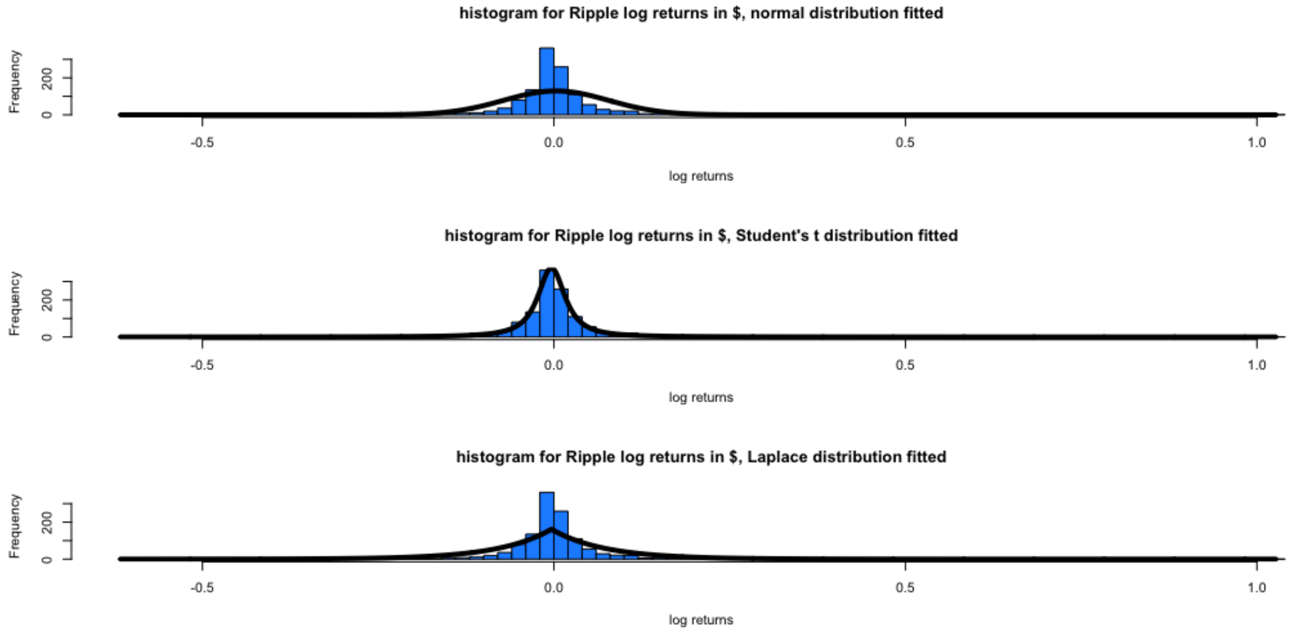

Figure 42: Fitting of distributions on Ripple log returns

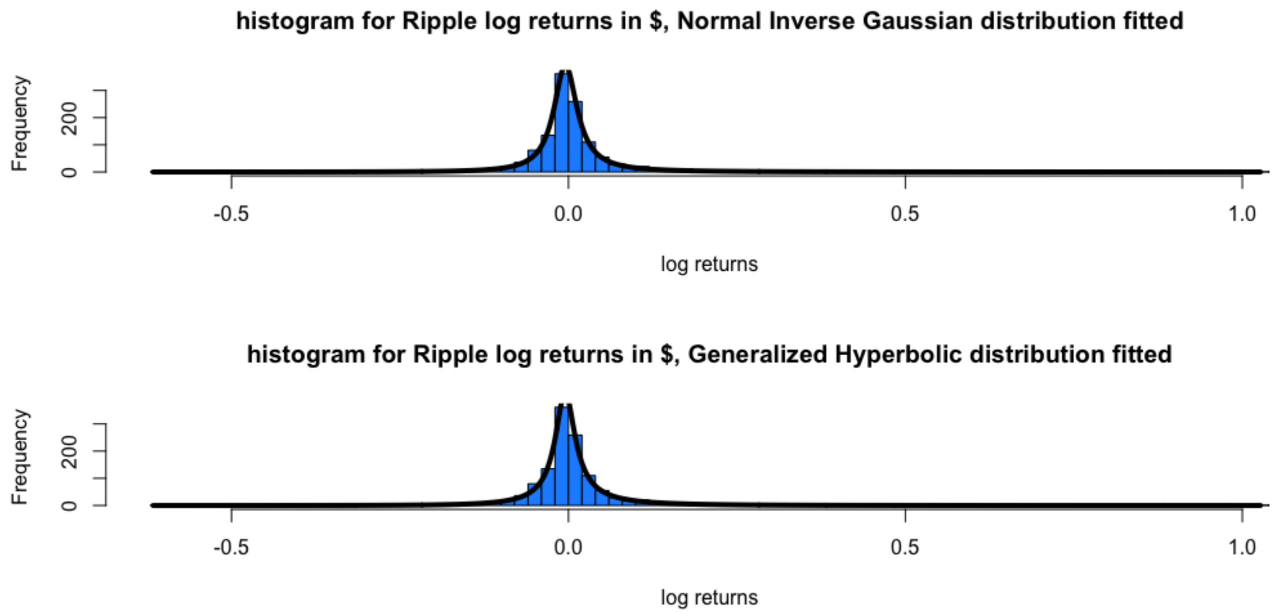

Figure 43: Fitting of distributions on Ripple log returns

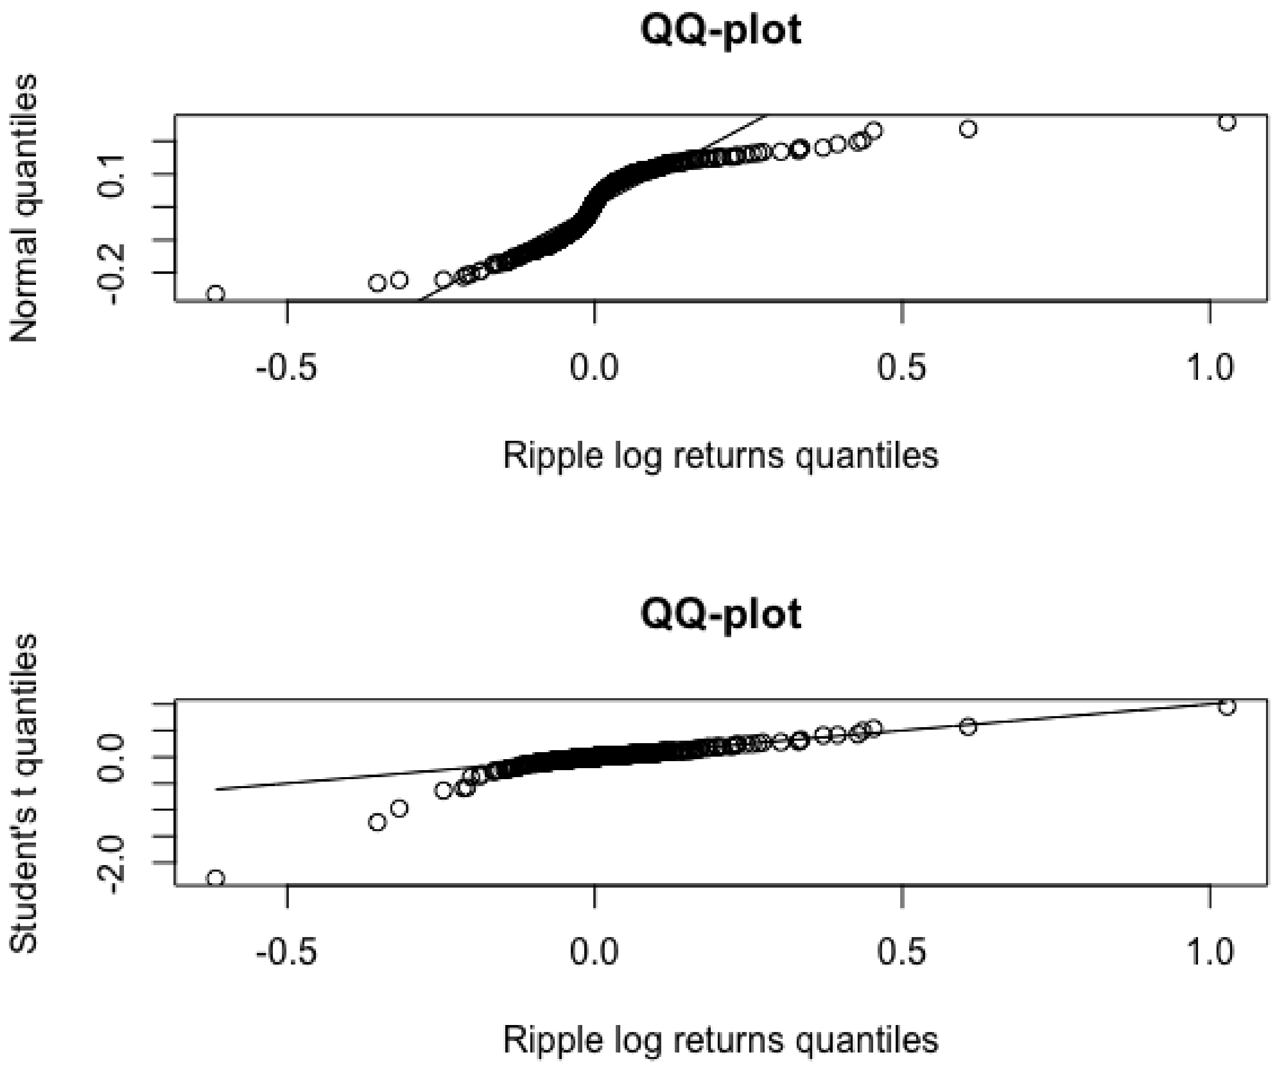

Figure 44: Corresponding QQ-plots for distribution fits on Ripple log returns

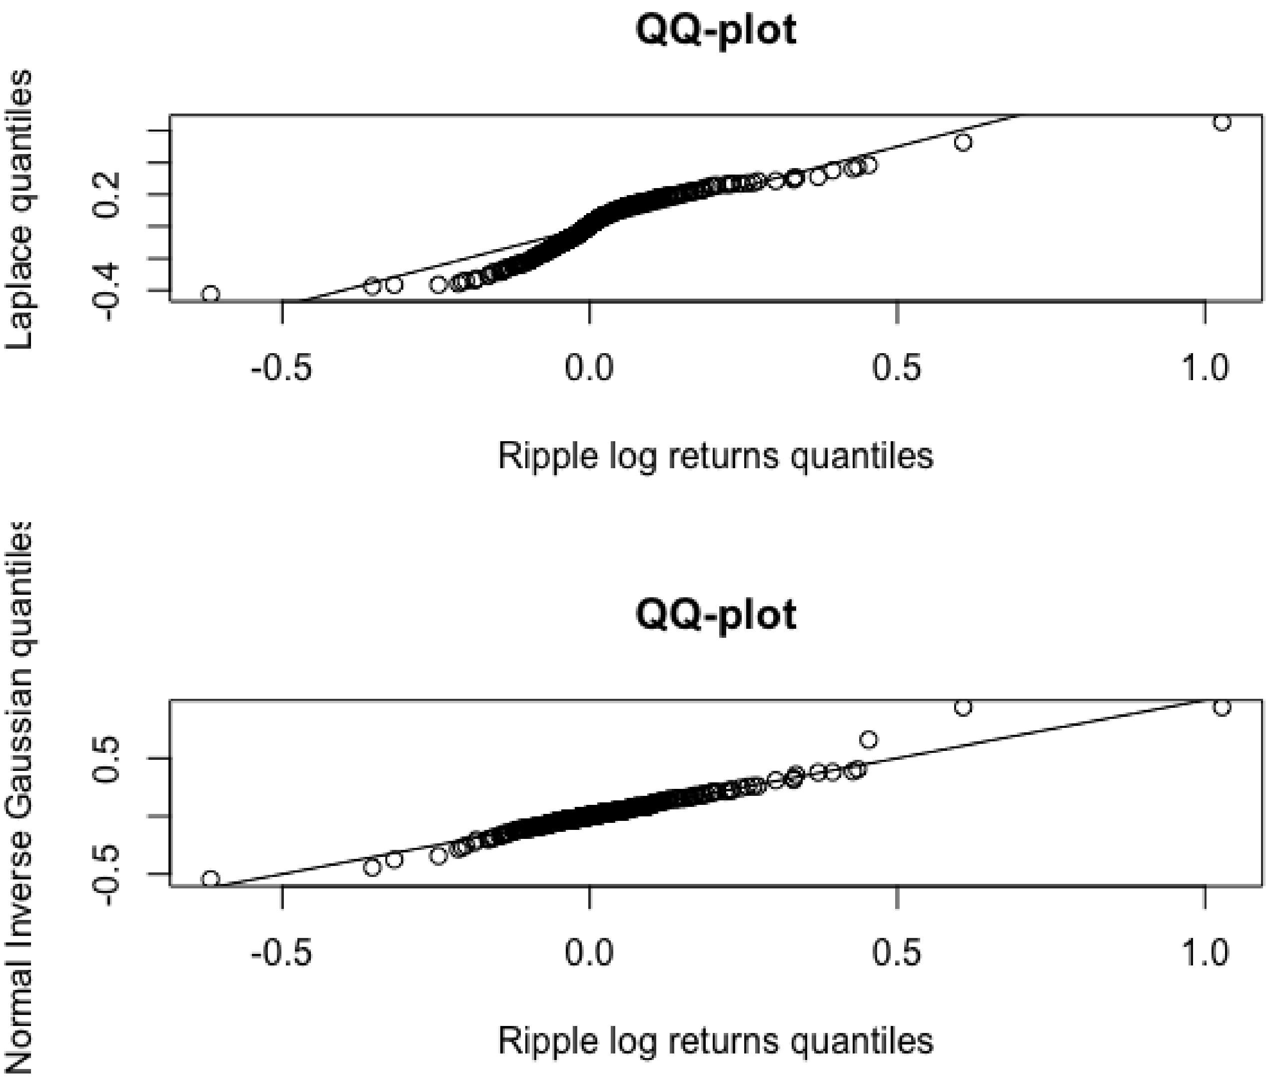

Figure 45: Corresponding QQ-plots for distribution fits on Ripple log returns

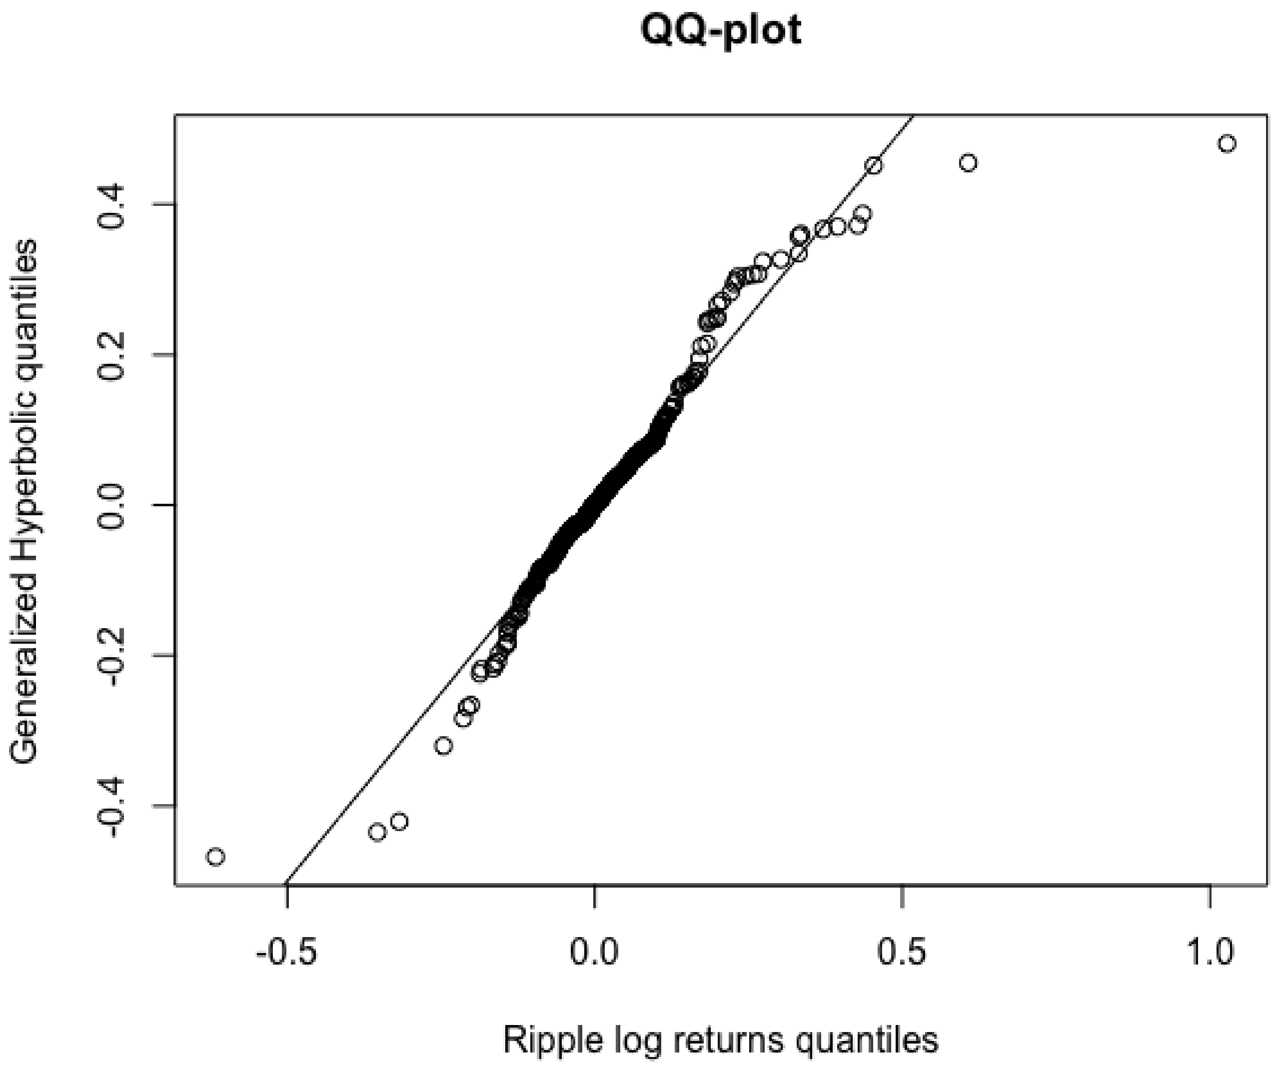

Figure 46: Corresponding QQ-plots for distribution fit on Ripple log returns

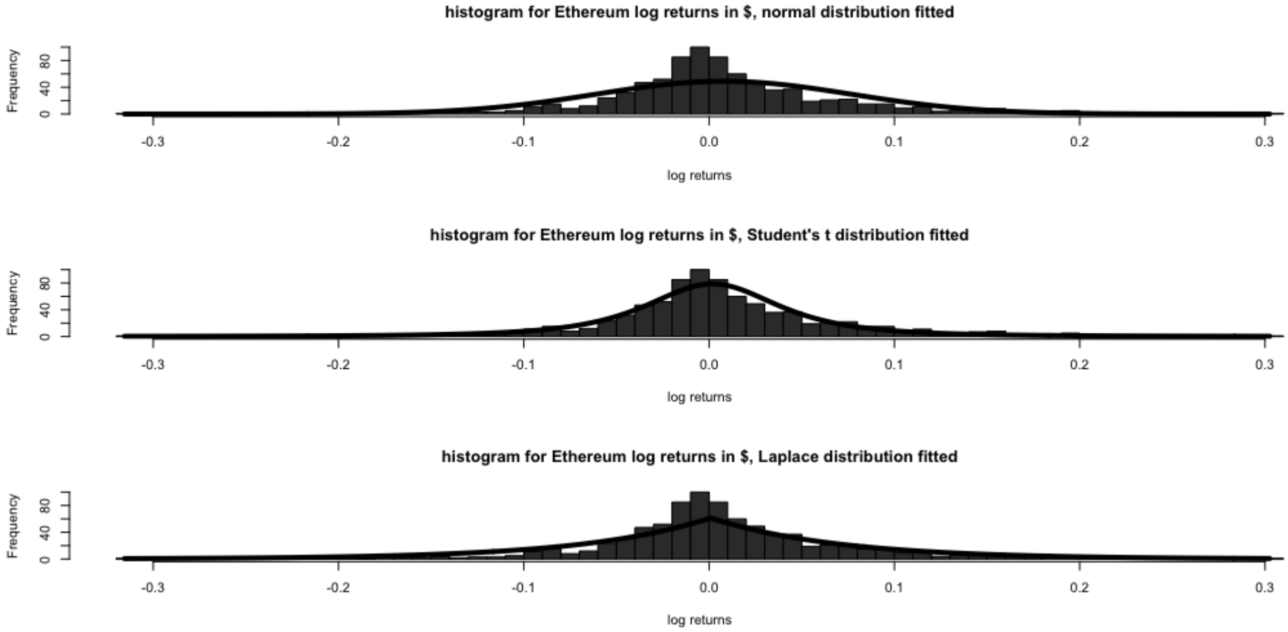

Figure 47: Fitting of distributions on Ethereum log returns

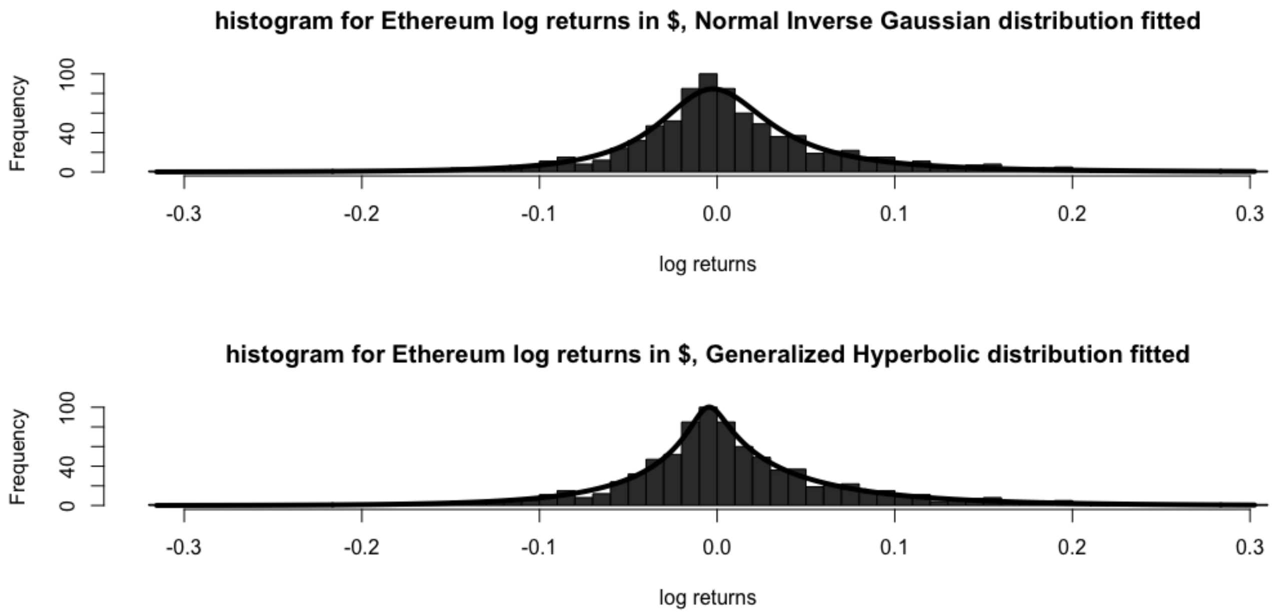

Figure 48: Fitting of distributions on Ethereum log returns

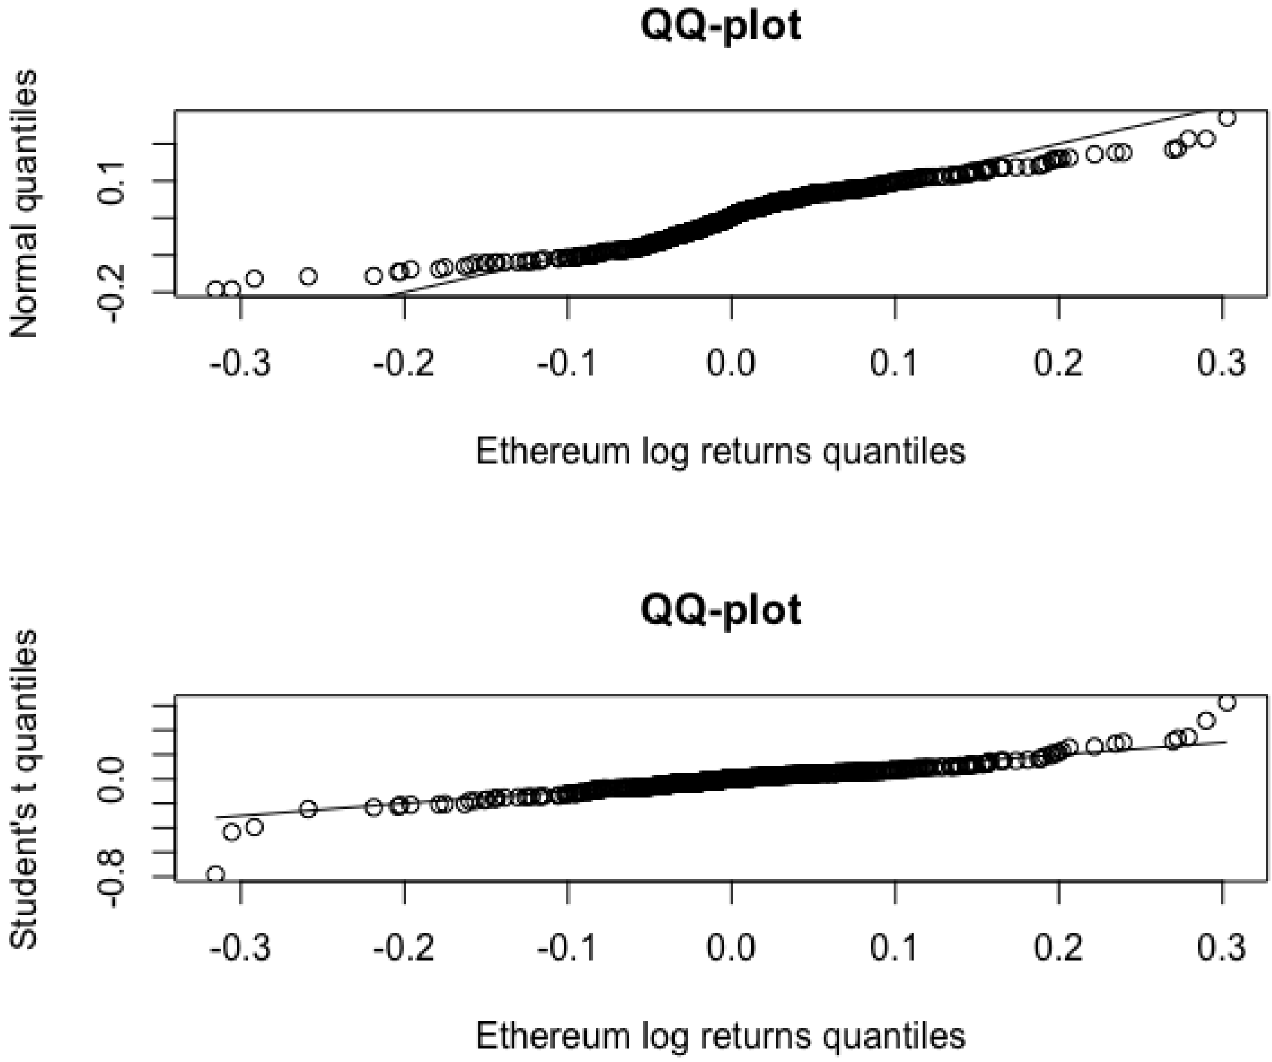

Figure 49: Corresponding QQ-plots for distribution fits on Ethereum log returns

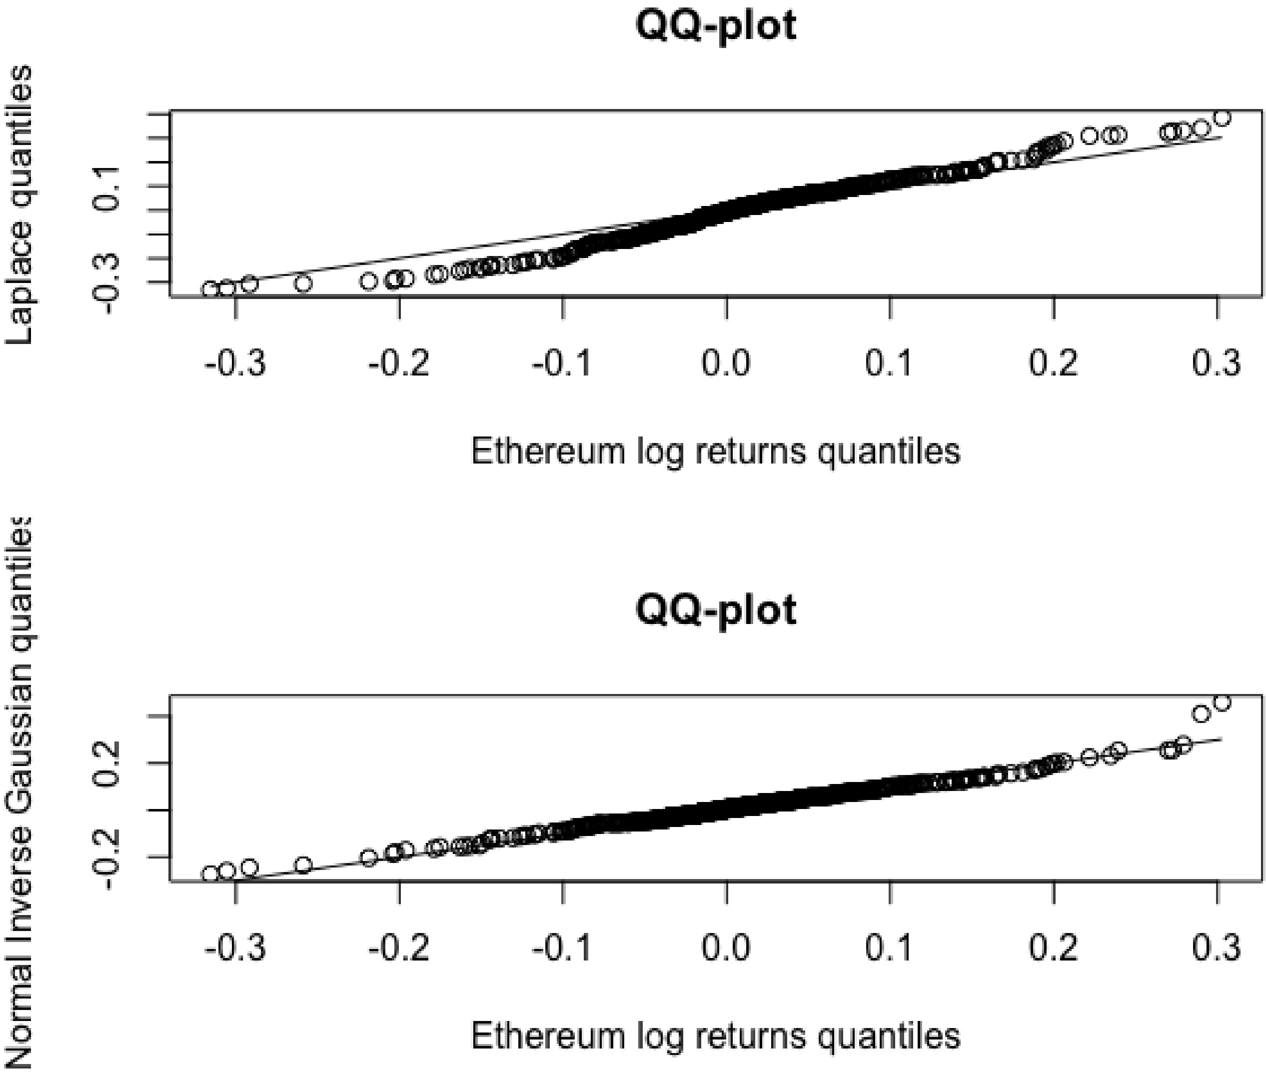

Figure 50: Corresponding QQ-plots for distribution fits on Ethereum log returns

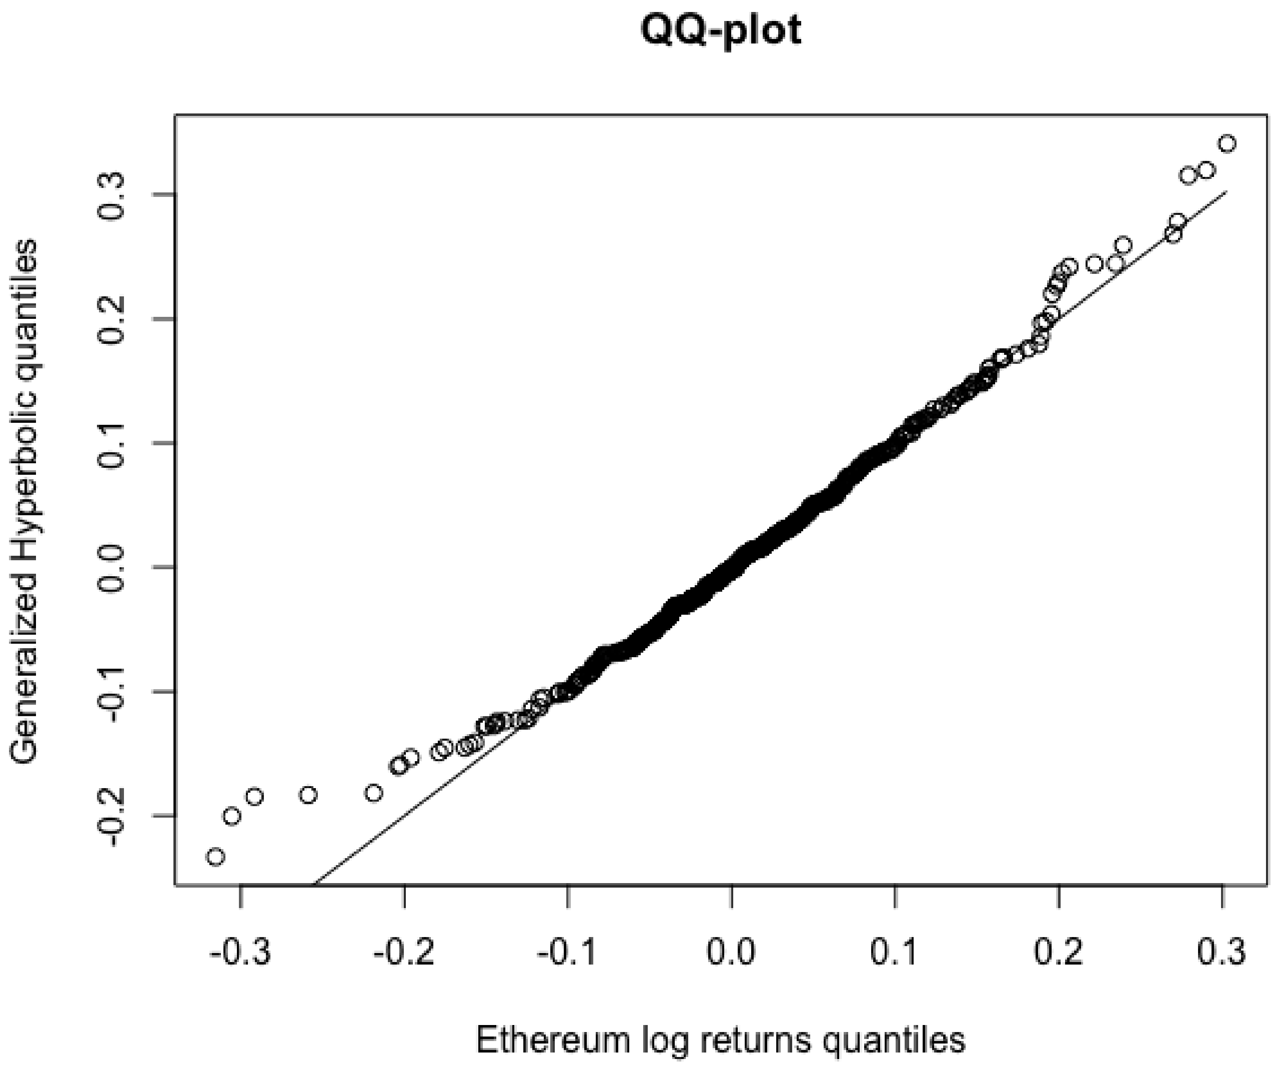

Figure 51: Corresponding QQ-plots for distribution fit on Ethereum log returns

6.3 Source code

Source code can be provided upon request.