LUND UNIVERSITY
School of Economics and Management

Volatility and Risk – FIGARCH Modelling of Cryptocurrencies

M.Sc. Thesis

Lund School of Economics and Management (LUSEM)
Lund University

Langkun Chen
Julian Ulmer

Supervisor:
Hans Byström

28 May 2021

Abstract

With increasing investor and thus media attention, academia has realized the need for more in-depth investigations of cryptocurrencies asset class as well as their volatility implications. We rely on these previous works for our model choice, to capture volatility differences between six cryptocurrencies and one fiat currency, namely the Euro, by using a FIGARCH(1,d,1) model. By including different proxies for monetary policy in the conditional variance equation, we want to investigate whether monetary policy significantly influences the volatility of the abovementioned assets. Deploying the value at risk, the expected shortfall as well as the stressed expected shortfall, using volatility weighted historical simulation, does not only allow us to incorporate current market conditions into the risk measures but also gives another way to compare the riskiness of the different assets. We conclude, that monetary policy does not significantly influence conditional volatility, due to the non-frequent changes in such. Nevertheless, we find meaningful parameter estimates within our model to explain the modelled series and the differences regarding persistence and memory of the different cryptocurrencies. Lastly, cryptocurrencies display a much higher market risk under the two standard risk measures and are even worse off under extremely stressed market conditions.

Keywords: *Cryptocurrencies, Conditional volatility, FIGARCH model, Monetary policy, Stressed expected shortfall*

Outline

Abbreviations

ADA	Cardano
ADF Test	Augmented Dicky-Fuller test
ARCH	Autoregressive conditional heteroscedasticity
ARMA	Autoregressive moving average
AWHS	Age weighted historical simulation
BIC	Bayesian information criteria
BHS	Basic historical simulation
BNB	Binance Coin
BTC	Bitcoin
ECB	European Central Bank
ES	Expected shortfall
ETH	Ethereum
EUR	Euro
FED	Federal Reserve System
FIGARCH	Fractionally integrated generalized autoregressive conditional heteroscedasticity
GARCH	Generalized autoregressive conditional heteroscedasticity
LM	Lagrange multiplier
LTC	Litecoin
MLE	Maximum likelihood estimation
OMO	Open market operations
USD	US dollar
VaR	Value at risk
VWHS	Volatility weighted historical simulation
XRP	Ripple

1. Introduction

Cryptocurrencies have been attracting more and more attentions from the media, governments, regulators and investors since the proposal of the first decentralized cryptocurrency, Bitcoin, in 2008 and its release in 2009 for public trading. Noticeably, there have been eye-catching rises in the price of Bitcoin since 2017: it increased by over 1300% throughout the whole year of 2017 alone. Its price witnessed a dramatic sky-rocketing since last October, from 10,619.45 USD on October 1st 2020 to 58,918.83 USD on March 31st 2021.[1] One could foresee an unignorable possibility where Bitcoins and other cryptocurrencies become even more relevant in the coming years.

The core technology behind Bitcoin and many other cryptocurrencies is named blockchain. It is a growing list of records, namely "blocks", linked with cryptography. Specifically, it could be regarded as a public ledger, with all transactions stored in this list of blocks (Zheng, Xie, Dai, Chen & Wang, 2017). Since the data in any block cannot be altered retroactively without altering all its subsequent blocks after being recorded, the blockchain is in itself resistant to data modification. This allows for user security and ledger consistency. In addition, blockchain technology has other key characteristics of decentralization, persistency, anonymity and auditability; blockchain thus can vastly save costs and improve efficiency with these given features (Phillip, Chan and Peiris, 2019). With blockchain, payments are allowed to be finished without any bank or other types of intermediary, which makes it possible for the technology to be used in various financial services, including digital assets. The proposal to use blockchain to build something as Bitcoin was first put up back in 2008 and later implemented in 2009. Several other cryptocurrencies, though not all, whose data we are going to use in this paper, are also designed with the blockchain technology.

Drastic changes in prices of cryptocurrencies and the potential large pay-offs that could be realized with those changes do not come without a huge and well-known volatility and thus uncertainty. Amidst the debates on the essential properties of cryptocurrencies, whether it functions more like a hedge asset or a conventional fiat currency, previous researches have been able to find evidence for similarities of cryptocurrencies and also its uniqueness in its place in investment and asset management. In the meantime, cryptocurrencies have been starting to play a growing role with a variety of aims in portfolios. This, together with the potential of an introduction of derivatives with cryptocurrencies as an underlying asset in the future, is making

[1] All data used in this paragraph is from Coinmarketcap (n.d.).

it ever more relevant to better understand the volatility characteristics of cryptocurrencies. Modelling cryptocurrencies' volatility also casts light on risk measurement of cryptocurrencies.

In this paper, we aim to model the conditional variance of several cryptocurrencies and one fiat currency, the Euro. By doing so, we look for evidence of the disparities and also similarities among our selected cryptocurrencies themselves as well as between them and the one fiat currency, in order to contribute to the understanding of the nature of cryptocurrencies at least to some degree. Here we look into adopting a time varying volatility model that also incorporates structural breaks in variances, as according to a number of previous literatures (Caporale and Zekokh, 2019; Katsiampa, 2019; Shen, Urquhart, and Wang, 2020), such an approach would tend to perform better in the modelling of cryptocurrencies' volatility. We would also like to include an explanatory variable as a proxy to mostly reoccurring types of monetary policies in our model, in order to see whether the changes in monetary policies impact cryptocurrencies' volatility. The estimated volatility by said model will finally be used in a further estimation of our chosen risk measure for cryptocurrencies, namely the stressed expected shortfall. We will proceed to compare this more advanced risk measure to more basic ones like the value at risk and the basic expected shortfall.

Some researchers have investigated risk measure estimations of cryptocurrencies after their volatility modelling. Chu, Chan, Nadarajah and Osterrieder (2017) estimated the value at risk. Caporale and Zekokh (2019) estimated both the value at risk and the expected shortfall and tried to pick the best volatility models based on the accuracy of these estimations. Klein, Thu and Walther (2018) calculated the average return of their constructed constantly rebalanced minimum variance portfolio consisting of one cryptocurrency and one market index under stressed times and argued that this mean could be seen as a sort of expected shortfall or conditional value at risk. According to the incoming new Basel regulations, the expected shortfall is set to take a more central stage than it has before. Thus, we hold the opinion that estimation of the expected shortfall, and on top of that, the stressed expected shortfall, could be of great interest in the very near future.

To conclude, one of our contributions to existing work is the finding regarding the similar trend in cryptocurrencies and the Euro, although they differ significantly in range and therefore in risk. This is supported by multiple findings from our estimated risk measures. For one the relative difference between VaR and ES is higher for cryptocurrencies, indicating more extreme tail events in the loss distribution. Furthermore, cryptocurrencies suffer significantly more under stressed market conditions, driving their stressed ES estimate to a much higher multiple than the one of the Euro. We also find that cryptocurrencies with a lower market capitalization display a higher volatility. A minor point is that we test the inclusion of a monetary policy proxy, which unfortunately is found to be insignificant.

The rest of this paper is organized as follows: section 2 is a summary of previous literature and a description of the theoretical framework where this paper lies; section 3 describes the methodology; section 4 describes the data used and also provides the summary statistics of the data; the description of our model is dedicated to section 5; section 6 presents our empirical results and finally in section 7 we arrive at the conclusion of the paper.

2. Literature Review and Theoretical Framework

In the following section we will discuss the previous literature on cryptocurrencies and their respective volatility implications. We start off with the papers covering the ongoing discussion about the classification of cryptocurrencies. Subsequently, we take a look at the papers regarding different approaches capturing the volatility, before we conclude this section with papers that elaborate on the use of risk measures to quantify risk in cryptocurrencies.

2.1 Asset Class of Cryptocurrencies

Since Bitcoin was introduced as the first publicly traded cryptocurrency in 2009, it has mostly been the researchers' consensus that cryptocurrencies are extremely volatile compared to traditional currencies or assets (Baur and Dimpfl, 2018; Caporale and Zekokh, 2019; Chu et al., 2017; Fang, Bouri, Gupta and Roubaud, 2019; Phillip et al., 2019). Following this insight, one of the main topics in economic literature is the debate under which asset class the highly volatile cryptocurrencies should be classified. According to Andrada-Félix, Fernandez-Perez and Sosvilla-Rivero (2020) cryptocurrencies are supposed to display huge similarities to normal currencies, as they function as a medium of exchange, a store of value and a unit of account. Many researchers state that most people treat Bitcoin as a risky investment vehicle instead of a medium of exchange, thereby classifying it as an asset (Dyhrberg, 2016; Fang et al., 2019; Liu and Tsyvinski, 2018). Furthermore, cryptocurrencies are not backed by a government and rely solely on their cryptographic integrity to ensure the security of transaction, whereas fiat currencies are backed by central banks and a regulatory framework, differentiating themselves from currencies even further (Andrada-Félix et al., 2020). Chu et al. (2017) support this line of argumentation with the fact that cryptocurrencies are digital assets relying on cryptography to secure extraction of new units. As the abovementioned extraction is limited for Bitcoin, Binance Coin, Ripple and many others, they derive value from that resource scarcity similar to commodities like gold or oil (Dyhrberg, 2016). According to the Commodity Futures Trading Commission (CFTC), virtual money is a declared commodity (Klein et al., 2018). Another argument supporting this classification of cryptocurrencies as a commodity is the finding of Fang et al. (2019) who state that Bitcoin serves as a hedge when global economic policy uncertainty is high. This result is in line with similar findings of Dyhrberg (2016) and Cheikh, Zaied and Chevallier (2020). In contrast to these results, Klein et al. (2018) find that Bitcoin does not possess the same hedging abilities as gold since it has a positive coupling effect.[2] Finally, it can be said that there remains no clear answer to the question which asset class

[2] A positive coupling effect means that the asset moves in the same direction as the market. In the case of Klein et al. (2018), Bitcoin prices decrease shock-like during stressed market conditions.

cryptocurrencies belong to. According to Dyhrberg (2016), cryptocurrencies are a special case of assets, displaying properties of multiple different asset classes. Therefore, Dyhrberg (2016) puts them somewhere between a currency and a commodity.

Although cryptocurrencies embody a highly secure architecture, the decentralized and unregulated markets may cause unpredictable price implications (Klein et al., 2018). Explaining these price and return implications has been one of the main goals of researchers, to help investors predict returns and thus yield abnormal returns. Liu and Tsyvinski (2018) find that cryptocurrencies have an extremely low exposure to traditional assets. Instead, they find that returns of cryptocurrencies are mainly predicted by investor attention and momentum. Positive investor attention is proxied by the number of google searches and twitter posts for the cryptocurrency in question, while negative investor attention is proxied by search results for "Bitcoin hack" (Liu and Tsyvinski, 2018). Nguyen, Nguyen, Nguyen and Pham (2019) try to explain cryptocurrency returns using monetary policy in monetary tightening (China) and monetary easing regimes (US), by deploying the Open Market Operations (OMO) target rate for the responsible central banks. Their result of a simple regression with multiple control variables, shows that only the Chinese OMO target rate significantly influences four cryptocurrencies (Liu and Tsyvinski, 2018), while the OMO rate of the US has no significance influence on the return of cryptocurrencies. They refer to the fact that China is the most active market for cryptocurrency trading.

2.2 Volatility Models

When modelling volatility of different cryptocurrencies, one of the most common procedures is to rely on a GARCH-type model. To ensure that GARCH modelling is possible, one common approach is to deploy the test for an ARCH effect (LM) suggested by Engle (1982). Katsiampa (2019), Cheikh et al. (2020), Dyhrberg (2016), Klein et al. (2018) and many others follow that procedure. Chu et al. (2017) fit 12 different GARCH-type models on log returns of different cryptocurrencies, and choose the best model using five goodness of fit estimates (AIC, BIC, CAIC, AICc and HQC).[3] Under normal innovations an IGARCH(1,1) provides the best fit for five of their seven tested cryptocurrencies (Chu et al., 2017). Exceptions are Ripple, for which a GARCH(1,1) yields better results, and Dogecoin, for which Chu et al. (2017) recommend a GJRGARCH(1,1).[4] Katsiampa (2017) runs six GARCH-type models on log returns of Bitcoin and finds an AR-CGARCH to be the optimal model, as it allows for a short-term and a long-

[3] Akaike information criterion (AIC), Bayesian information criterion (BIC), consistent Akaike information criterion (CAIC), corrected Akaike information criterion (AICc) and Hannan–Quinn criterion (HQC).
4 Both are again obtained under normal innovations.

term component, when describing conditional volatility. Fang et al. (2019) use a GARCH-MIDAS model, which also splits up the volatility into short- and long-term components. They use this model to show the impact of global economic policy uncertainty (GEPU) on the long-run volatilities. They obtain their explanatory variable using a purchasing power parity adjusted and GDP weighted average of national economic policy uncertainty indices. Phillip et al. (2019) state that cryptocurrencies display oscillatory long run autocorrelation behavior whereas in fiat currencies this long run autocorrelation is slowly decaying. To account for that and the rather frequent jumps in cryptocurrencies, Phillip et al. (2019) deploy a stochastic volatility model with a Gegenbauer filter, which adjusts the volatility measure for long run autocorrelation and the Buffered Autoregressive model to account for jumps. They conclude that cryptocurrencies with a higher transaction speed display oscillatory characteristics and a lower liquidity risk during such transactions and are therefore more valid as a medium of exchange.

Shen et al. (2020) use intraday data with a time interval of five minutes to run 18 heterogenous autoregressive (HAR) models. They find that HAR models that include jumps have a higher explanatory power in sample and a better out of sample forecasting ability. After the inclusion of structural breaks, which increases the forecasting ability even further, Shen et al. (2020) find the HARQ-F-J to be the best model. Mensi, Al-Yahyaee and Kang (2019) run four GARCH models that all are adjusted to take structural breaks as well as long memory into account. After an out of sample analysis using the mean squared error and the mean absolute error, Mensi et al. (2019) find that the FIGARCH with structural break dummy variables has the best fit. Katisampa (2019) investigates whether certain events change the structural properties of a conditional volatility series and are thus responsible for structural breaks. According to Caporale and Zekokh (2019), standard GARCH models can yield biased results if the volatility series displays structural breaks, which is very likely to occur for cryptocurrencies. Katisampa (2019) deploys an asymmetric multivariate GARCH model and checks for structural breaks using the Bai-Perron test and the Chow-test for both mean and variance equations. The author finds one structural break in the variance equation of each the Bitcoin and the Litecoin volatility series. Katsiampa (2019) relates the structural break in the Bitcoin volatility series to the shutdown of China-based cryptocurrency exchanges. Caporale and Zekokh (2019) tackle the problem of structural breaks in the volatility series by deploying a Markov-Switching GARCH model, where parameters can change over time to account for possible structural breaks. Rapach and Strauss (2008), who check for structural breaks in exchange rate volatility, do so by using the iterative cumulative sum of squares (ICSS) algorithm developed by Inclan and Tiao (1994). To account for possible observed structural breaks, Rapach and Strauss (2008) fit their GARCH model on the obtained respective subsamples that are separated by their detected structural breaks.

According to Cheikh et al. (2020), equities display a leverage effect which causes negative news to have a higher impact on volatility than positive ones. Many researchers also look at asymmetric effects for cryptocurrencies and thus adjust their model accordingly. Baur and Dimpfl (2018) take a behavioral focused approach, investigating asymmetric effects in cryptocurrencies. They do so, using a threshold GARCH (TGARCH) model in combination with an asymmetric response measure to capture potential asymmetric effects. Their result that a majority of cryptocurrencies exhibit an inverted asymmetric effect (meaning that negative shocks increase the volatility by less than positive shocks do), is explained by herding behavior of uninformed investors buying, afraid of not being part of upcoming profits. Cheikh et al. (2020) find an inverted asymmetric effect, using a smooth transition GARCH instead of the usually used threshold GARCH model. Katsiampa (2019) supports this finding with the Asymmetric Diagonal BEKK model, which considers co-movements as well as structural breaks. By using an APARCH and a FIAPARCH model to take asymmetry as well as long memory into account, Klein et al. (2018) find an inverse leverage effect for Bitcoin. Only Dyhrberg (2016) finds no asymmetric impact on volatility, contradicting the other researchers. Klein et al. (2018) explain this difference with the fact that they use a t-distribution, which naturally accounts for heavy tails in the return distribution. Baur, Dimpfl and Kuck (2018) attempt to replicate the work by Dyhrberg (2016), during which they realize that Dyhrberg (2016) has made quite a few mistakes regarding different assumptions about her GARCH modelling (stationarity, spurious correlation, multicollinearity, etc.). This outlines the importance of further work on this topic, as only with enough relevant and recent research a concluding hypothesis can be made. Nevertheless, the consensus of the majority of researchers is that cryptocurrencies display an inverse leverage effect.

A slightly different approach when investigating volatility of cryptocurrencies in comparison to fiat currencies is taken by Andrada-Félix et al. (2020). Whereas most papers are modelling conditional volatility with the goal to find the best volatility model, Andrada-Félix et al. (2020) study the connectedness of volatilities between multiple fiat and cryptocurrencies. They thereby also check for volatility spillovers and the drivers of such volatility, using a Time-Varying Parameter Vector Autoregressive (TVP-VAR) approach.

2.3 Risk Measures for Cryptocurrencies

A quite common approach when investigating cryptocurrencies is to additionally incorporate risk measures, as the conditional volatility implications can be incorporated into the risk measures, for example by using volatility weighted historical simulation (VWHS). Chu et al. (2017) compute parametric value at risk estimates for eight different underlying distributional assumptions (for innovations). Furthermore, they deploy the Christoffersen test to check whether the obtained value at risk estimates are acceptable. One day ahead estimates of value at risk and expected shortfall are calculated by Caporale and Zekokh (2019) using a rolling window. Mensi et al. (2019) find that the FIGARCH model including dummy variables for structural breaks is the best model to quantify value at risk. With the upcoming Basel III regulations in the beginning of 2023 and the associated implications for risk measures, researchers like Fang et al. (2019) suggest using the (stressed) expected shortfall measure instead of the value of risk estimates. Instead of using one of the common parametric or non-parametric methods to estimate the value at risk or expected shortfall, Klein et al. (2018) opt to use a rather unusual approach. They start with constructing a constantly rebalanced minimum-variance portfolio of a market index and a cryptocurrency, and then calculate the mean of the returns of this portfolio at distressed times. Then Klein et al. (2018) argue that this mean can be seen as a kind of expected shortfall or conditional value at risk estimate.

3. Methodology

The methodology section gives insight about our used models and their respective econometric properties. This is crucial in order to verify our model choice and understand the later discussed empirical results.

3.1 Data Characteristics

Oftentimes price data is the easiest to access in various data sources, but to do statistical inference about return variance it is crucial to transform price data into return data. This can be done either looking at the relative price change from time $t-1$ to time t or by using log returns, and the latter is used in this paper. To get a better understanding of the underlying data series, we will compute the standard descriptive statistics, which display measures as mean, median, standard deviation, skewness and kurtosis for each return series. To check whether the return series follow a normal distribution, we deploy the Jarque-Bera test for normality.

3.1.1 Log Returns

The abovementioned transformation of prices into returns will be done using transformation with the natural logarithm, according to the following equation:

$$r_t = \ln\left(\frac{p_t}{p_{t-1}}\right) \tag{1}$$

Where:

r_t	is the return in period t
p_t	is the price in period t
p_{t-1}	is the price in period $(t-1)$

3.1.2 Jarque-Bera Test

The Jarque-Bera test (1980), which we deploy to understand the underlying distribution of the return series, is a goodness of fit test of whether the sample skewness and kurtosis of the sample data matches those of a normal distribution. The test statistic for the Jarque-Bera test is given by:

$$\mathrm{JB} = \frac{T}{6} \times \left(S^2 + \frac{1}{4}(K-3)^2\right) \sim \chi^2 \tag{2}$$

Where:

T	is the number of observations
S	is the sample skewness
K	is the sample kurtosis

A test statistic that is higher than the critical value leads to a rejection of the null hypothesis of normality.

3.2 Conditions for GARCH Modelling

After familiarizing us with the characteristics of the return series for each of our cryptocurrencies, we get back to the main focus of this paper, which is modelling the volatility implications. To explain the differences in volatility amongst each asset and the influence of monetary policy on the conditional variance, we are to model a conditional volatility series. To justify such GARCH modelling, there must be a present ARCH effect as well as a stationary underlying return series. To check for the ARCH effect, we first run the Ljung-Box test on the squared residuals of our return series and then verify the results with the LM test for an ARCH effect. To test our time series for a unit root / stationarity we will use the Augmented Dickey-Fuller test.

3.2.1 Ljung-Box Test

We first give an insight into the general Ljung-Box test, as this is crucial in understanding the Ljung-Box test when it is performed on the squared residuals. The Ljung-Box test (1978) proposed by Ljung and Box was originally a diagnostic tool testing the lack of fit in a time series model. After fitting an ARMA(p,q) model to the time series, the test is performed on the residuals by examining m autocorrelations of the residuals. The test statistic for the Ljung-Box test is given by:

$$Q_m = T(T+2)\sum_{k=1}^{m}\frac{\hat{\rho}_k^2}{T-k} \sim \chi^2 \tag{3}$$

Where:

$\hat{\rho}_k^2$	is the estimated autocorrelation of the series at lag k
T	is the number of observations

The underlying model could be concluded to not exhibit significant lack of fit if the autocorrelations are very small and the null hypothesis, that the residuals from the ARMA model have no signification autocorrelation, is not rejected. On the other hand, the test could be applied to revealing evidence for potential existence of residual autocorrelation in a time series with the null being rejected instead. With that being said, the Ljung-Box test does not directly tackle other serial dependencies than autocorrelation.

Nevertheless, by testing squared residuals with the Ljung-Box test, we could also identify autoregressive conditional heteroscedasticity (ARCH) effects in the series (McLeod and Li, 1983). We rewrite the abovementioned Ljung-Box test to the specific test that tackles the detection of an ARCH effect using the squared residuals. The equation stays the same besides that the test is now performed on the squared residuals, precisely by examining m autocorrelations of the squared residuals.[5] This procedure should yield similar results to Engle's Lagrange multiplier test (1982), which is described below.

3.2.2 LM Test for ARCH Effect

The test for ARCH effect, based on the Lagrange multiplier, proposed by Engle (1982), is the standard test for detecting autoregressive conditional heteroscedasticity in time series. We use this test to verify our results of the above explained Ljung-Box test. The following explanation of the test is inspired by Sjölander (2011). If an ARCH(q) model with a return r_t that is equal to the conditional expectation μ plus a stochastic error η_t is written as:

$$r_t = \mu + \eta_t \tag{4}$$

$$\eta_t = \sigma_t \varepsilon_t \tag{5}$$

With:

$$\varepsilon_t \sim N(0,1) \qquad\qquad \eta_t \sim N(0, \sigma_t^2)$$

And the conditional variance, depending on the past values of squared errors:

$$\sigma_t^2 = \omega + \alpha_1 \eta_{t-1}^2 + \alpha_2 \eta_{t-2}^2 + \cdots + \alpha_q \eta_{t-q}^2 \tag{6}$$

Where:

ω is the intercept / constant term

α_i is the weight of the lagged error term at $(t-i)$

After regressing r_t with the specification as in equation (4) with OLS and obtaining the residuals $\hat{\eta}_t^2$, we use the estimated residuals in the following regression model:

$$\hat{\eta}_t^2 = \omega + \alpha_1 \hat{\eta}_{t-1}^2 + \alpha_2 \hat{\eta}_{t-2}^2 + \cdots + +\alpha_q \hat{\eta}_{t-q}^2 + \epsilon_t \tag{7}$$

With:

$$H_0\colon \alpha_1 = \cdots = \alpha_q = 0 \qquad\qquad H_1\colon \exists \alpha_i \neq 0 \quad for\ i = 1,2,\ldots,q$$

Then the test statistic is given by:

$$(T-q)R^2 \sim \chi^2 \tag{8}$$

[5] For computational details refer to the paper of McLeod and Li (1983).

Where:

T	is the number of observations
R^2	is the R-square obtained from regressing r_t

A rejection of the null hypothesis suggests existence of ARCH effects and thus provides justification for the usage of GARCH-type models on the underlying time series.

3.2.3 Augmented Dickey-Fuller Test

The augmented Dickey-Fuller (ADF) test looks for the presence of a unit root in a time series, thus we use it to check for stationarity in our seven return series. The exact alternative hypothesis for the ADF test can vary with the test specification (for us, the one that tests against stationarity is used).

It is based on the regular Dickey-Fuller test (Fuller, 1976; Dickey and Fuller, 1979). When the latter is no longer to be valid if the condition of the error term being white noise is not satisfied, which would be the case if there exists autocorrelation in the first difference of the dependent variable Δr_t, we can 'augment' the test using p lags of the dependent variable and thereby the augmented version for the test is given by:

$$\Delta r_t = \psi r_{t-1} + \sum_{i=1}^{p} \alpha_i \, \Delta r_{t-i} + u_t \tag{9}$$

And the test statistic as:

$$DF_\tau = \frac{\hat{\psi}}{SE(\hat{\psi})} \tag{10}$$

With:

$$H_0: \psi = 0 \qquad\qquad H_1: \psi < 0$$

A test statistic that is less than (i.e., more negative) the corresponding critical value leads to a rejection of the null hypothesis of a unit root; specifically, in our case, this then provides evidence for stationarity in the time series.

3.3 Concept of Conditional Variance

After verifying the conditions to run a GARCH model we go into detail of what is implied by such a conditional time-varying volatility series. The concept of time varying volatility is crucial to understand how a GARCH model works, as it incorporates current market conditions and thus volatility clustering into the variance equation. We furthermore explain the basic GARCH model and the implied coefficients, as it is the basis of understanding the later deployed more complex FIGARCH model.

3.3.1 Time-Varying Volatility

Measuring the unconditional volatility of a data series with a standard volatility measure like the standard deviation or variance yields a constant number for the volatility of a sample, as it only depends on an estimated set of persistent coefficients. Accordingly, the volatility is equal for each observation in the given sample making it impossible to differentiate between periods of low or high volatility, which would display more realistic market conditions. Allowing for time-varying volatility can solve this problem, as the volatility for each observation can vary now. Accordingly, current market conditions as well as volatility clustering are taken into consideration when modelling time-varying volatility.

To understand the different models which can be used to display such time-varying volatility, it is important to be familiar with the concept of conditional heteroscedasticity. Under homoscedasticity the error terms all have the same variance, whereas this variance differs from observation to observation under heteroscedastic errors. Conditional heteroscedasticity is not predictable, since volatility is non-constant compared to previous periods. In comparison to that unconditional heteroscedasticity is predictable (for example cyclicality).

3.3.2 GARCH

The Autoregressive Conditional Heteroscedasticity (ARCH) model, for which Engle and Granger won the Nobel Prize in Economics in 2003, would be one of many suitable options to display time-varying volatility. As we later use a more advanced model which is based on the Generalized Autoregressive Conditional Heteroscedasticity (GARCH) model, we think it is crucial to understand the underlying principles of such a model. With a GARCH model it is possible to model the volatility series with the help of the volatility observed in the previous periods and the error term (unexpected return shocks) observed in previous periods. How many periods one looks back depends on the model specification, more precisely on the order of the model. The difference between an ARCH and a GARCH process is that the GARCH process includes lagged values of the variance, whereas the ARCH process only looks at lagged values of the error term. Assuming the following return series (simple mean model) with return r_t, mean μ_t and unexpected return shock (error term) η_t, it is possible to model the volatility of the returns using a GARCH(1,1) process:

$$r_t = \mu_t + \eta_t \tag{11}$$

$$\eta_t = \sigma_t \varepsilon_t \tag{12}$$

With:

$$\varepsilon_t \sim N(0,1) \qquad \eta_t \sim N(0, \sigma_t^2)$$

GARCH(1,1):

$$\sigma_t^2 = \omega + \alpha\eta_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{13}$$

Where:

ω is the intercept / constant term

α is the weight of the lagged error term at $(t-1)$

β is the weight of the lagged variance at $(t-1)$

To obtain the abovementioned estimates of a GARCH model the maximum likelihood estimation (MLE) method is used.[6] Per definition, all coefficients on the right-hand side of the GARCH process are bigger or equal to zero to ensure a positive variance. Furthermore, the sum of the alpha and beta coefficients needs to be smaller than one to ensure a stationary process.[7] The higher the beta value in a GARCH model compared to the alpha value, the smoother the variance series, as past observations have a higher perseverance within the process. In the simple GARCH model the alpha parameter represents the ARCH effect while the beta parameter represents the GARCH effect.

3.4 Structural Breaks

Multiple previous researchers state that including structural breaks, will increase the goodness of fit of a GARCH model in an in-sample model and the forecast accuracy in an out of sample forecast. As we model multiple conditional volatility series it is crucial for us to check for structural breaks in the volatility series and not the return or price series. We do this on the sample volatility of each return series, which is computed over a rolling window of 60 observation. By using the Bai-Perron test we obtain the breakpoints and the according break dates for each volatility series.

3.4.1 Bai-Perron Test

The Bai-Perron test, first proposed in a seminal paper by Bai and Perron in 1998 is an endogenous technique that could detect multiple structural breaks in longitudinal data, that is, it allows for the data to learn itself where the breakpoints lie. A later paper in 2003 by the same authors dealt with the computational difficulties faced by the test by using a dynamic programming algorithm based on the Bellman principle.[8]

[6] Refer to APPENDIX A for a brief introduction into maximum likelihood estimation.

[7] An IGARCH(1,1) process is a special case of a GARCH(1,1) process where the sum of the alpha and beta coefficients are equal to one. Thus, an IGARCH model, displays a non-stationary volatility series as shocks are persistent.

[8] The Bellman equation, named after Richard E. Bellman, is the necessary condition for dynamic programming to reach mathematical optimization.

Bai and Perron proposed three types of tests for detecting structural breaks. The first two tests are tests of structural stability versus a fixed or an unknown number of breaks while the third test is a sequential test that determines on l and $l+1$ structural changes. One preferred strategy (Jouni and Boutahar, 2005) in determining the numbers of structural breaks is to first check for whether or not there are any structural breaks in the series using one of the first two tests, and then check for the actual number by repetitively deploying the sequential test (starting from one breakpoints) until the null hypothesis of l structural breaks is no longer rejected. This is, as concluded by Bai and Perron (2003), a method that leads to more satisfactory results and is recommended for empirical applications. This is also the approach taken by Katsiampa (2019) in search for structural breaks in both mean and variance equations of various cryptocurrencies.

The actual equations involved in the Bai-Perron test, which one can look for in the two abovementioned previous works by these authors, are quite complex and most possibly beyond the scope of this paper. Yet fortunately there have been rather simpler implementation approaches in various computational programs.

3.5 Risk Measures

After modelling the conditional volatility of each subsample using the FIGARCH model, we want to end our discussion with the investigation of different risk measures. After defining losses and rescaling them, we compute the value at risk using a non-parametric approach (VWHS). After that, we also compute the expected shortfall and the stressed expected shortfall, which will be used under the upcoming Basel III regulation. Comparing these risk measures yields insights into the risk of cryptocurrencies and the implications of the stressed expected shortfall.

According to Artzner, Delbaen, Eber and Heath (1999), risk is always related to uncertainty in future values, whether they are due to changes in market conditions or unforeseeable events. In line with Szegö's (2002) definition of a risk measure, a mapping f is used from the space of random variables X to a non-negative real number R.

$$f: X \rightarrow R$$

As a result, a scalar risk measure is obtained, which is desirable since it ensures comparability within one risk measure and thus can help to improve investment decisions. For any risk measure $f: X \rightarrow R$ to be classified as an acceptable risk measure, the following conditions need to be fulfilled (Szegö, 2002):

Positive homogeneity:

$$f(kx) = kf(x) \qquad \forall\, k > 0$$

Subadditivity:

$$f(x + y) \leq f(x) + f(y)$$

For a risk measure to be coherent the following conditions need to be fulfilled in addition to the previous two:
Monotonicity:

$$x \leq y \Rightarrow f(x) \leq f(y)$$

Translation invariance:

$$f(x - m) = f(x) - m \qquad \forall\, constant\ m$$

Due to time and space limitations and the fact that these insights are well known, we will not go into the economic significance of these axioms, but instead refer to the papers of Artzner et al. (1999), Frittelli and Gianin (2002) and Szegö (2002) who not only explain abovementioned axioms in detail, but also give reasoning why it is desirable to use a coherent risk measure.

3.5.1 Value at Risk (VaR)

To understand the later used risk measure, the stressed expected shortfall, it is crucial to be familiar with the concept of value at risk as it lays the foundation for the expected shortfall calculation. The VaR measure at a confidence level α, gives the α quantile of a prespecified or an estimated sample loss distribution. Economically the VaR estimate can be interpreted as a risk measure which states, that with a probability of α the portfolio in question will lose more than the VaR_α.

Normally losses for any asset are computed from the price series. In general losses equal profit with a flipped sign and can therefore be calculated in the following way:

$$l_t = -(p_t - p_{t-1}) \tag{14}$$

Where:

l_t	is the loss at time t
p_t	is the price / value of the asset at time t
p_{t-1}	is the price / value of the asset at time $t - 1$

Nevertheless, in our case we have to compute loss scenarios from the return series, since only then are we able to rescale the losses with the volatility series obtained from the return series.

This is in practice done by using the following equation and assuming a portfolio of size 10000 USD:

$$l_t = -10000 * \frac{r_t}{100} \tag{15}$$

Where:

r_t is the return in period t

Looking at an example of continuous normally distributed losses $L \sim N(0,1)$, the $VaR_{0.99}$ is the following:

Graph 3.1. VaR Illustration

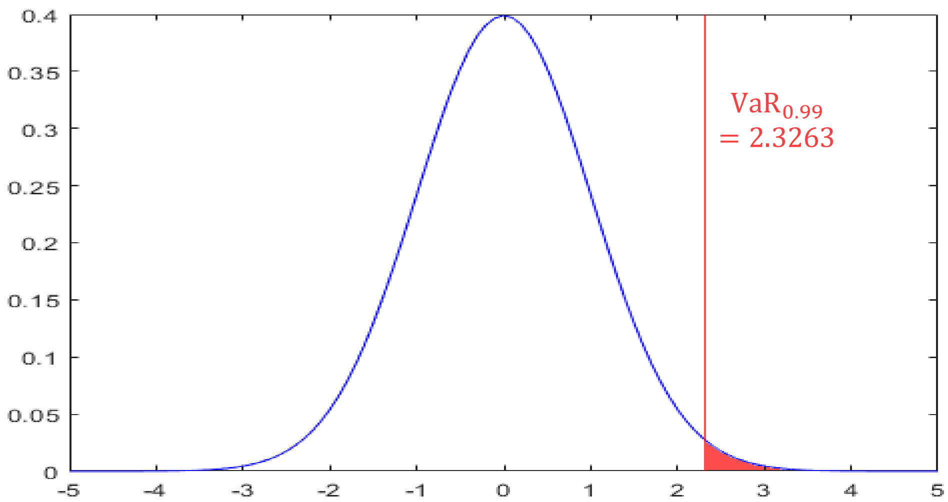

The value at risk estimate equals the smallest loss which still fulfills the requirement that the probability of losses being larger than that loss are smaller or equal than $1 - \alpha$. Putting that definition into an equation, one can compute the VaR in the following way:

$$VaR_\alpha = min\{l: \Pr(L > l) \leq 1 - \alpha\} \tag{16}$$

This definition of the value at risk nests the equation for a continuous loss distribution as the desired probability $1 - \alpha$ can always be obtained when the underlying loss distribution is a continuous one:

$$\Pr(L > VaR_\alpha) = 1 - \alpha \tag{17}$$

When Value at Risk or ES is computed using a non-parametric approach, the estimation relies on the empirical loss distribution, where different methods like the basic historical simulation (BHS), age weighted historical simulation (AWHS) and volatility weighted historical simulation (VWHS) can be used. As the later used GARCH-type model will yield a conditional volatility series, the VWHS will be deployed when computing the risk measure. While deploying the VWHS we scale losses by volatility and then apply a percentile function to obtain the according risk measure in question.

Scaled losses:

$$l_1^* = \frac{\sigma_{T+1}}{\sigma_1} l_1 \tag{18}$$

$$\vdots$$

$$l_t^* = \frac{\sigma_{T+1}}{\sigma_t} l_t$$

Where:

l_t is the original loss at time t

l_t^* is the rescaled loss at time t

σ_t is the volatility of the original losses at time t

σ_{T+1} is the one-day-ahead forecasted volatility/conditional volatility from the GARCH model

As time-varying volatility is accounted for under VWHS the risk measure estimates will take the then-current market conditions into account. Time-varying volatility can also be incorporated under a parametric approach, by replacing the sample volatility with the conditional volatility. Using a parametric approach requires an assumption about the loss distribution, as the estimate is then computed using parameters of the underlying distributional assumption.

Although VaR is not a coherent risk measure,[9] it was used under the Basel I regulation with a 99% confidence level and a 10-day horizon and stressed VaR was used with a 10-day horizon under the Basel II.5 regulation (Hull, 2018).

3.5.2 Expected Shortfall (ES)

The expected shortfall for continuous loss distributions can be interpreted as the expected loss of a portfolio conditional that the loss is larger than the value at risk estimate.

$$ES_\alpha = E[L | L > VaR_\alpha]^{10} \tag{19}$$

Following this interpretation, it is straightforward that the ES estimate must equal the average of all losses larger than the value at risk in question. Using the same loss distribution $L \sim N(0,1)$

[9] VaR does not fulfill the subadditivity requirement as it is unable to catch very high losses with a very low probability, which can be captured when combining two assets with such characteristics to a portfolio, thereby making the portfolio VaR higher than the added VaR of both assets.

[10] Holds only for continuous loss distributions.

as in the example used to display the $VaR_{0.99}$ (Graph 3.1), the $ES_{0.99}$ estimate equals the integrated values to the right of the $VaR_{0.99}$ estimate (red area in Graph 3.1).

Taking the discrete case into account as well as the continuous one, the definition of the ES changes to the average value at risk VaR_x for all $\alpha \leq x \leq 1$. Displaying this analytically yields following equation:

$$ES_\alpha = \frac{1}{1-\alpha}\int_{x=\alpha}^{x=1} VaR_x \,\mathrm{d}x \tag{20}$$

In comparison to the value at risk, expected shortfall is a risk measure which fulfills all four coherency requirements, making it a more desirable risk measure.[11]

3.5.3 Stressed ES

The stressed expected shortfall is a risk measure which in principle is just the expected shortfall for a consecutive period of extreme stress. This means that the ES is calculated for the 250 worst trading days displaying stressed market conditions and thus yielding the highest ES estimate for the portfolio in question (Hull, 2018). In comparison to the VaR and the current ES we do not use a rolling window of 250 observations, but we look at the one most stressed consecutive period of 250 observations.[12] Furthermore, the stressed ES is a risk measure for which back testing is not possible,[13] as it uses extreme data which is statistically not expected to occur again in similar fashion in the future (Hull, 2018).

Using a risk measure to quantify the risk of cryptocurrencies has been done before; value at risk estimates have been computed by Chu et al. (2017) and Ardia, Bluteau and Rüede (2019). In addition to the value at risk Caporale and Zekokh (2019) also computed ES estimates. With the upcoming Basel III regulation, the risk measure used will be swapped to a stressed ES estimate with a 97.5% confidence level and liquidity reflecting time horizons (Basel Committee on Banking Supervision, 2014).[14] Thus we will compute all our risk measures on a confidence level of 97.5%. Furthermore, we choose to compute the stressed ES as it will give more valuable insight compared to the other estimates once the Basel III regulation is implemented on the 1st January 2023. Chu et al. (2017) follow that line of argumentation and suggest using the expected shortfall instead of the value at risk estimates used by themselves.

---

[11] ES fulfills the subadditivity requirement as it takes the whole right tail of a loss distribution into account, thereby capturing very high losses with an extremely small probability. By overcoming this shortfall of value at risk, the expected shortfall estimate is able to display diversification benefits in all cases.

[12] We decide the most stressed period simply by looking at the amount of risk quantified by the ES estimates.

[13] VaR and ES are usually back tested to evaluate the estimates, using for example the Kupiec test or the Basel traffic light test for VaR and the Acerbi-Szekely test for ES.

[14] The liquidity adjusted varying time horizon will not be included here since they do not change the main results obtained from our model.

4. Data

Since cryptocurrencies are a rather new asset class, with new ICOs (initial coin offerings) happening frequently (Adhami, Giudici and Martinazzi, 2018), it is important to ensure relevancy and maturity for the underlying data. Therefore, we follow the example of Andrada-Félix et al. (2020), who choose cryptocurrencies with high market capitalizations. Furthermore, we extend the selection criteria of Katsiampa (2019) who requires cryptocurrencies to exist for at least 2 years by imposing the requirement that cryptocurrencies are traded for at least three years. This is done to ensure enough observations to make significant inference about the return volatility. Applying abovementioned selection criteria in our shortlist leaves us with six cryptocurrencies (Bitcoin, Ethereum, Binance Coin, Cardano, Ripple and Litecoin)[15] that together make up a market capitalization of over 1.4 trillion $ according to Coinmarketcap (n.d.). Bitcoins closing prices are obtained from 2017-01-01 up until the 2021-03-31, thereby capturing the highest price movements of the asset. Closing price series are obtained for the same time frame for Ethereum, Ripple and Litecoin. The data series for Binance Coin ranges from 2017-07-26 to 2021-03-31 as Binance launched in July 2017. The closing prices for Cardano range from the 2017-10-02 to the 2021-03-31. All closing price data series are extracted from Coinmarketcap (n.d.) and are than transformed to return series using log returns. The USD to EUR spot exchange rates are obtained from Datastream for a time period starting from the 2017-01-01 until the 2021-03-31. As data for cryptocurrencies is available 365 days a year whereas exchange rate data is only available on trading days, we remove weekend data for all cryptocurrencies to ensure consistency within our model and increase comparability. This is done as it might be possible that weekend dynamics differ from weekday dynamics, thus including weekend data would result in differences in the volatility estimates. Klein et al. (2018) came across the same problem and opted to synchronize their data as well, as interpolation of weekend prices would have resulted in biased parameter estimates when estimating the model.

To proxy for the influence of monetary policy there are multiple options, which could serve as fitting explanatory variables for our conditional variance series. The reserve requirement, the open market operations (OMO) target rate, the discount rate as well as the interest rate on reserves are all instruments by central banks to influence monetary policy. Nguyen et al. (2019) choose the OMO target rate to explain returns of cryptocurrencies with different monetary regimes. Although they find significant results for the OMO rate of China, we will stick to a

[15] We exclude Tether (USDT) from our sample given the fact that is a stablecoin, mirroring the price of the U.S Dollar by holding equivalent USD cash reserves. As the price is always very close to one USD the volatility implications are not of importance for us.

more commonly used rate, namely the discount rate. The discount rate describes the interest rate that commercial banks have to pay to the central bank when taking a loan for a certain time window.[16] The discount rate is usually higher than the interbank rates, as borrowing from a central bank is seen as a signal of a bank being in distress. A high discount rate signals a rather tightening monetary policy as borrowing becomes more expensive, whereas a low discount rate signals monetary easing. The choice of the central bank is as important as the choice of an appropriate measure of monetary policy. As cryptocurrencies are a global phenomenon, it is crucial to incorporate the discount rate from a central bank which is relevant for the global economy. As the US FED is one of the most influential central banks is the world, we will explore the influence of the FED's discount rate on the seven different conditional volatility series. More specifically, we will use the primary credit as our exogenous proxy for monetary policy. We obtain this rate from the official website of the federal reserve (FED, n.d.).

4.1 Descriptive Statistics

We now document the major statistical properties of the time series in historical prices of the chosen assets (Bitcoin, Ethereum, Binance Coin, Cardano, Ripple, Litecoin and the Euro) and the corresponding returns for each of them.

APPENDIX B presents the historical price data (regard the EUR/USD exchange rate as a sort of price for the Euro in dollar terms) as well as the return data for each of the seven assets. All of the cryptocurrencies, except for Binance Coin, display a first high peak almost simultaneously around December 2017 to January 2018, while the second significant rise (from the late second half of 2020 to the end of our observation period) is witnessed by the five cryptocurrencies apart from Ripple.

In APPENDIX C the histograms for each of the return series is shown. From the return plots, we can tell that all seven assets have their means at around zero throughout our selected time period, which can also be verified with the histograms. Arguably, volatility clusters can be identified already for some of the cryptocurrencies (Ripple, Binance Coin and Cardano) as well as heteroscedasticity. We can see several peaks in almost all return series in cryptocurrencies but also periods of relative tranquility; the return series of the Euro is instead smoother than the cryptocurrencies', and also features a much smaller range. We are also able to tell that all series are somehow symmetrically distributed from the histograms, while featuring much higher peaks than the one of a normal distribution.

[16] Primary credit (usually overnight, available for sound financial institutions), secondary credit (usually overnight, for financial institutions that are not eligible for primary credit but comes at a higher rate than the primary credit) and seasonal credit.

Table 4.1 provides the summary statistics for the return data. The insights that we can obtain from the table basically agrees with our observations from the figures. All seven assets have their means and medians very close to zero. The selected cryptocurrencies all appear to have a much larger range in their returns than the Euro while at the same time having a higher standard deviation.

Table 4.1 Summary Statistics for Return Data

	r-BTC	r-ETH	r-XRP	r-LTC	r-BNB	r-ADA	r-EUR
Mean	0.0037	0.0049	0.0041	0.0034	0.0083	0.0042	0.0001
Min	-0.4647	-0.5507	-0.5504	-0.4490	-0.8183	-0.5037	-0.0175
Max	0.2251	0.3914	0.7508	0.5398	1.1973	0.8867	0.0175
Median	0.0031	0.0012	-0.0022	0.0000	0.0025	-0.0006	-0.0001
Std	0.0495	0.0675	0.0835	0.0710	0.0941	0.0908	0.0041
Kurtosis	9.7378	8.0469	16.7994	9.9519	36.9042	22.8017	1.3778
Skewness	-0.7369	0.0228	1.8639	0.8775	1.9128	2.4854	0.1072

Note: This table reports the summary statistics for each of the seven return series for our chosen cryptocurrencies and the Euro. The return series for Bitcoin, Ethereum, Ripple, Litecoin and the Euro ranges from 2017-01-03 to 2021-03-31, while Binance Coin from 2017-07-27 to 2021-03-31 and Cardano from 2017-10-03 to 2021-03-31. The descriptive statistics included in the table are **(1) mean** (the mean of each of the return series from the entire sample period), **(2) min** (the minimum of each of the return series), **(3) max** (the maximum of each of the return series), **(4) median** (the median of each of the return), **(5) std** (the sample standard deviation of each of the return series), **(6) kurtosis** (the calculated excess kurtosis for each series) and **(7) skewness** (the calculated skewness for each series).

All cryptocurrencies have an excess kurtosis that is much higher than the Euro. While Bitcoin is found to be the only asset to be negatively skewed, the other assets all have a positive skewness (though for Ethereum and the Euro, the degree of skewness is quite minimal). Table 4.2 presents the results from the Jarque-Bera test (1980) for normality on these return series. The null hypothesis of normality is rejected for all return series at a one percent significance level, which serves as the reason for choosing a t-distribution in the returns rather than a normal one. The rejection is also in line with our observation of the histograms from above.

Table 4.2 Results for the Jarque-Bera Test

	r-BTC	r-ETH	r-XRP	r-LTC	r-BNB	r-ADA	r-EUR
JB-stat	4429.5***	2955.9***	1353.1***	4663.8***	5447.7***	2046.6***	88.207***
P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table reports the results of the Jarque-Bera test on each on the seven return series. The null hypothesis for this test is normality. */**/*** indicates statistical significance level at the 10/5/1 percent level. The two statistics presented in the table are **(1) JB-stat** (the test statistics for the Jarque-Bera test on each of the seven return series from the entire sample period) and **(2) P-value** (the corresponding p-values for each of the test-statistics). Clearly the null hypothesis is rejected at the 1% significance level for each of the return series in our sample period.

Table 4.3 and 4.4 are the results for the Ljung-Box test (1978) on the squares of the residuals and Engle's Lagrange multiplier test (1982) for the ARCH effects for the seven return series up until five lags.[17] The test results from both tests are almost identical, with the null hypotheses of no significant ARCH effect being rejected for all but the first lag of Binance Coin at a five percent significance level. Nonetheless, the null hypothesis is still rejected at a ten percent significance level for that lag of Binance Coin. The test results display evidence for the existence of the ARCH effect in the return series, thus providing justification of using a GARCH-type model.

Table 4.3 Results for the Ljung-Box Test

		BTC	ETH	XRP	LTC	BNB	ADA	EUR
Lag 1	t-stat	4.857**	9.627***	32.825***	46.509***	2.824*	39.894***	20.071***
	P-value	0.028	0.002	0.000	0.000	0.093	0.000	0.000
Lag 2	t-stat	7.444**	14.055***	53.113***	50.977***	16.905***	89.446***	32.380***
	P-value	0.024	0.001	0.000	0.000	0.000	0.000	0.000
Lag 3	t-stat	13.192***	23.085***	55.715***	51.654***	21.373***	89.653***	41.189***
	P-value	0.004	0.000	0.000	0.000	0.000	0.000	0.000
Lag 4	t-stat	14.170***	23.251***	56.172***	58.937***	35.558***	89.679***	51.155***
	P-value	0.007	0.000	0.000	0.000	0.000	0.000	0.000
Lag 5	t-stat	22.819***	28.837***	82.195***	61.485***	41.182***	89.679***	76.479***
	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table reports the results of the Ljung-Box test on the squared residuals of the seven return series. We allow for up until five lags in the test. Testing on the squared residuals means that the test now de facto tests for the ARCH effect in each series. */**/*** indicates statistical significance level at the 10/5/1 percent level. The two statistics presented in the table are **(1) t-stat** (the test statistics for the Ljung-Box test on the squared residuals of each of the seven return series from their entire sample period) and **(2) P-value** (the corresponding p-values for each of the test-statistics).

Table 4.4 Results for the Lagrange Multiplier ARCH Test

		BTC	ETH	XRP	LTC	BNB	ADA	EUR
Lag 1	t-stat	4.840**	9.604***	32.710***	46.348***	2.812*	39.766***	21.056***
	P-value	0.028	0.002	0.000	0.000	0.094	0.000	0.000
Lag 2	t-stat	7.005**	12.721***	45.360***	46.824***	16.170***	73.830***	27.436***
	P-value	0.030	0.002	0.000	0.000	0.000	0.000	0.000
Lag 3	t-stat	11.722***	19.707***	45.403***	46.914***	19.186***	78.071***	27.355***
	P-value	0.008	0.000	0.000	0.000	0.000	0.000	0.000
Lag 4	t-stat	11.806**	19.865***	45.359***	52.953***	29.182***	79.093***	35.094***
	P-value	0.019	0.001	0.000	0.000	0.000	0.000	0.000
Lag 5	t-stat	18.980***	24.0567***	69.076***	53.169***	31.908***	79.347***	53.608***
	P-value	0.002	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table reports the results of Engle's LM ARCH effect test on the seven return series. We allow for up until five lags in the test. */**/*** indicates statistical significance level at the 10/5/1 percent level. The two statistics presented in the table are **(1) t-stat** (the test statistics for the LM ARCH effect test on each of the seven return series from their entire sample period) and **(2) P-value** (the corresponding p-values for each of the test-statistics). Almost identical results are obtained from this test when compared to the results from the Ljung-Box test as presented in Table 3.

[17] We allow for up until five lags in the Ljung-Box test, ARCH LM test and the ADF test. Since we are using daily data with weekend excluded, five lags (thus technically a week) appears to be one reasonable option.

Lastly, we present results for the augmented Dickey-Fuller test (Fuller, 1976; Dickey and Fuller, 1979) up until five lags in Table 4.5, with the null hypotheses being of a unit root. The null hypotheses are rejected at a one percent significance level for all series, which shows evidence for the series' stationarity.

Table 4.5 Results for the Augmented Dickey-Fuller Test

		BTC	ETH	XRP	LTC	BNB	ADA	EUR
Lag 1	t-stat	-22.952***	-22.186***	-20.458***	-22.718***	-18.871***	-18.169***	-22.978***
	P-value	0.000	0.002	0.000	0.000	0.000	0.000	0.000
Lag 2	t-stat	-17.931***	-17.912***	-17.294***	-18.689***	-15.194***	-15.375***	-19.123***
	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Lag 3	t-stat	-15.256***	-14.868***	-14.328***	-15.134***	-12.455***	-13.787***	-17.239***
	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Lag 4	t-stat	-14.147***	-13.401***	-12.306***	-14.248***	-10.835***	-12.713***	-15.175***
	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Lag 5	t-stat	-12.605***	-12.533***	-10.525***	-13.263***	-10.082***	-12.090***	-14.864***
	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table reports the results of the augmented Dickey-Fuller test against a null hypothesis of a unit root on the seven return series. We presented the test results up until five lags augmented. */**/*** indicates statistical significance level at the 10/5/1 percent level. The two statistics presented in the table are **(1) t-stat** (the test statistics for the ADF test on each of the seven return series from their entire sample period with the given number of lags augmented) and **(2) P-value** (the corresponding p-values for each of the test-statistics). Strong evidence of series being stationary is detected by the test.

4.2 Structural Breaks

To test for structural breaks within the volatility series, there are two options that seem plausible: either fitting the GARCH (in our case, FIGARCH) model first, testing the conditional volatility series for structural breaks and then adjusting the GARCH for the structural changes, or otherwise creating a sample volatility series over a rolling window, checking this series for structural breaks and then fitting the GARCH model with the found structural breaks taken into account. We opt for the latter, since it seems like the more practical approach not having to run the GARCH model twice.

To create the rolling window of sample volatilities, we use a window length of 60 observations[18] with the date in question centered in the middle of the window. At the beginning and the end of the sample the window is truncated to the according size, so that the variance is only calculated for the elements inside of the window. The choice of which window size to use is important, as the larger the window the smoother the volatility series will be.

[18] This is close to a window of three months since we previously removed weekend data making the usual month about eight to ten days shorter.

To account for the obtained structural breaks there are multiple ways in which the GARCH model can be adjusted. Mensi et al. (2019) use structural break dummy variables, while Caporale and Zekokh (2019) use a Markov-Switching GARCH model, to account for these structural breaks by varying parameters. We follow the example of Rapach and Strauss (2008) and separate our series into subsamples, in which the breakpoints detect where the subsamples should start and end.[19] The advantage with this method is that the parameters can vary from subsample to subsample.

4.2.1 Bai-Perron Test

To successfully deploy the Bai-Perron test for multiple structural breaks, certain assumptions have to be made. We allow for a maximum of two structural breaks, considering that for each structural break we will obtain an additional subsample to fit our FIGARCH model on. Furthermore, we follow the usual rule of thumb to have at least 15% of the total observations in each subsample. This is done to ensure that each subsample has a sufficient size and thus enough relevant data. To decide whether to include none, one or two structural breaks we use the Bayesian information criteria (BIC), by selecting the case in which the obtained BIC value is the lowest.[20] APPENDIX D displays the output from the Bai-Perron test for each sample volatility series and thus represents the basis on which our choice of structural breaks is made. To graphically display the obtained structural breaks, the sample volatility series as well as the return series are modelled again including red dash-dotted lines separating the data series according to the structural breakpoints from the Bai-Perron test (see APPENDIX E).

First of all, it is interesting to note that each volatility series displays two structural breaks, indicating differences in parameters. When looking at the detected dates for our structural breaks we try to sort them and see whether some fall into the same timeframe. A similarity that all volatility series of our six cryptocurrencies share is a structural break in the beginning of 2018. They are located between the 2018-01-29 for Litecoin and the 2018-04-11 for Bitcoin and Cardano, so roughly in a timeframe of two and a half months. For all cryptocurrencies this is shortly after the price reversal of the first big rally of cryptocurrencies, which started in the summer of 2017. The only asset that does not display a structural break in the beginning of 2018 is our one fiat currency, the Euro. The first structural break in the volatility series of the Euro is on the 2018-09-21, quite at the beginning of a steady decrease in price until the

[19] We incorporate the date of the structural break into the former subsample and let the next subsample start on the day right after the structural break.

[20] Note that allowing for even more structural breaks would often yield to an even lower BIC, but this is unfeasible for our paper due to the time and space constraints that will come into play when analyzing all obtained coefficients for each subsample.

beginning of 2020. This is where the second structural break for the volatility series of the Euro occurs, specifically on the 2020-01-29. Litecoin and Ethereum's volatility series both have a structural break merely one day later, on the 2020-01-30. These structural breaks could coincide with the beginning of the Covid-19 pandemic in China. Ripple, Binance and Cardano have the second structural break in the late summer of 2020, namely between the 2020-08-11 for Ripple and the 2020-09-21 for Cardano. For the second structural break only one asset, namely Bitcoin, has its other structural break before 2018. The first breakpoint for Bitcoin is on the 2017-08-22, right before the first price rally of Bitcoin. With these two breakpoints the volatility series of Bitcoin is separated into three subsamples, before the price rally of 2017, during the price rally of 2017 and after.

5. Model

Since we have found significant ARCH effects in all seven return series, by using the Ljung-Box and the LM test, we have a strong reason to believe that a GARCH model is a great choice to model these series volatility. To further strengthen the use of such a GARCH model, we find that all of our obtained return series are stationary, which is favorable when deploying a GARCH model. These finding are in line with previous findings from cryptocurrency literature as many researchers use GARCH models to capture the time-varying volatility of these assets.

Since our paper does not aim to run multiple models and find the best fitting one by using multiple information criteria (as done by Chu et al., 2017), we rely on previous research to construct our model. As discussed in section 2 of this paper, different authors suggest different models, when it comes to best capture the volatility of cryptocurrencies. After carefully weighing arguments for different models, we choose to run a FIGARCH model as suggested by Mensi et al. (2019), since we believe its possibility to incorporate long memory effects in the conditional variance equation will help to model cryptocurrencies' volatility.

To incorporate our exogenous variable for the monetary policy of the FED into our FIGARCH model we came across two plausible options. Under the first we would include the explanatory variable in the return equation already, therefore the effect would show up in the error term (unexpected return shock), whereas under the second we include the explanatory variable directly in the conditional variance equation. The direct approach would leave us with an extra coefficient in our conditional volatility model, making it quite easy to interpret the effect of the monetary policy. Since we want to measure this effect on the variance of the cryptocurrencies and not necessarily on the returns, we choose to do the straightforward implementation of the exogenous variable into the FIGARCH equation.

As mentioned above the FIGARCH model will be deployed to each subsample of all seven return series, to account for the above detected structural breaks. In addition, the FIGARCH is run on the whole return series of each asset, as an anchor for us in modelling.

5.1 FIGARCH

The fractionally integrated generalized autoregressive conditional heteroscedasticity (FIGARCH) model was introduced by Baillie, Bollerslev and Mikkelsen in 1996. We closely follow their specification of the model in order to omit certain details. Interested readers can always refer to the original paper. Under the FIGARCH model the effect of the lagged squared innovations (unexpected return shocks) on the conditional variance decays with a slow hyperbolic rate. Therefore, the FIGARCH model is able to capture long-memory effects while still allowing the shocks to decay unlike the IGARCH model. To get a better understanding we take a look at the FIGARCH(p,d,q) model, in which p determines the number of AR lags (GARCH effect) and q determines the number of MA lags (ARCH effect). It is important to note that in the FIGARCH model, the coefficient of the ARCH effect is ϕ but not α as in the simple GARCH model. Nevertheless, it will still be interpreted in a similar way, namely as the ARCH parameter (following Mensi et al., 2019) and thus as the influence of the unexpected shocks. We will follow the general approach used in practice regarding the volatility modelling of cryptocurrencies and rely on first order models with only one lag, as they have proven to be a good representation of conditional variance processes (Baillie et al., 1996). Under the condition $0 < d < 1$, the process displays long memory for the conditional variance which nevertheless will die out over time (Mensi et al., 2019). The closer the fractional differencing parameter d goes to one, the higher the memory of the FIGARCH model (Ghalanos, 2020). Furthermore, with $d = 0$ the FIGARCH collapses to a simple GARCH and with $d = 1$ it converges to the IGARCH model. Modelling the return series as an ARMA(1,1) process:

$$r_t = \mu + \vartheta r_{t-1} + \eta_t + \zeta\eta_{t-1} \tag{21}$$

Allows us to already incorporate an AR and a MA effect in the mean equation, with η_t being the unexpected return shock (error term).

Given η_t is a discrete time real-valued stochastic process:

$$\eta_t \equiv \varepsilon_t\sigma_t \tag{22}$$

With:

$$\varepsilon_t \sim (0,1) \qquad\qquad \eta_t \sim (0, \sigma_t^2)$$

We are able model the conditional variance series as a FIGARCH(1,d,1) process, with σ_t representing the available information set at time $t - 1$, as a time varying function,
A FIGARCH(p,d,q) process can be displayed using the following equation:[21]

$$\phi(L)(1 - L)^d\eta_t^2 = \omega + [1 - \beta(L)]v_t \tag{23}$$

[21] We still show the general framework of the FIGARCH(p,d,q), since it is straightforward to obtain our used FIGARCH(1,d,1) model, by simply replacing p with one and q with one.

In which L is defined as the backshift operator, which lags the coefficients:

$$\phi(L) = [1 - \alpha(L) - \beta(L)](1 - L)^{-1} \tag{24}$$

$$\alpha(L) = \alpha_1 L + \alpha_2 L^2 + \cdots + \alpha_q L^q \tag{25}$$

$$\beta(L) = \beta_1 L + \beta_2 L^2 + \cdots + \beta_p L^p \tag{26}$$

The fractional differencing operator is the following:

$$(1 - L)^d = \sum_{k=0}^{\infty} \frac{\Gamma(d+1)}{\Gamma(\mathrm{k}+1)\Gamma(\mathrm{d}-\mathrm{k}+1)} L^k \tag{27}$$

In which $\Gamma(z)$ defines the gamma function:

$$\Gamma(\mathrm{z}) = \int_{x=0}^{x=\infty} x^{z-1} e^{-x} \, \mathrm{d}x \tag{28}$$

Rearranging equation (23) with:

$$v_t = \eta_t^2 - \sigma_t^2 \tag{29}$$

Yields the following expression for the FIGARCH(p,d,q):

$$[1 - \beta(L)]\sigma_t^2 = \omega + [1 - \beta(L) - \phi(L)(1 - L)^d]\eta_t^2 \tag{30}$$

The conditional variance of the stochastic process η_t is then given by:

$$\sigma_t^2 = \omega[1 - \beta(L)]^{-1} + \{1 - [1 - \beta(L)]^{-1}\phi(L)(1 - L)^d\}\eta_t^2 \tag{31}$$

To estimate the parameters of the underlying FIGARCH model, we again rely on the approach proposed by Baillie et al. (1996). To obtain maximum likelihood estimates for the FIGARCH(p,d,q) model the following log likelihood function must be optimized:

$$\log L(\theta; \eta_1, \eta_2, \ldots, \eta_T) = -0{,}5T log(2\pi) - 0{,}5 \sum_{t=1}^{T} [\log(\sigma_t^2) + \eta_t^2 \sigma_t^{-2}] \tag{32}$$

Where:

$$\theta' = (\omega, d, \beta_1, \ldots, \beta_p, \phi_1, \ldots, \phi_q) = (\omega, d, \beta_1, \phi_1) \tag{33}$$

Displays the starting values under which the log likelihood function is maximized. The second equality in equation (33) holds since we use a first order model with only one lag each. These parameters are namely $\omega, d, \beta_1, \phi_1$, respectively, a constant, the long-memory parameter, the GARCH effect and the ARCH effect. A problem that might occur with the abovementioned maximum likelihood estimation is that under the default starting values supplied by R-Studio the global optimum might not be reached. Nevertheless, we stick with that MLE as approaches to guarantee a global maximum are outside the scope of this thesis, given the limited time frame. To work around this problem of being stuck at a local maximum, we will vary the starting values using a loop function in R until we find the parameter starting value combination which gives us the highest log likelihood value.[22] We believe that under this approach we are able to model the dynamics to the best of our abilities, without an ad-hoc choice of parameters. To verify our obtained results, we manually compare the conditional volatility series from each subseries, with the sample volatility series as well as the conditional volatility series of the whole data set, to see whether the according parts resemble each other. In certain cases, we obtain implausible parameter estimates from starting values that result in log likelihoods, which are single outliers and are only above the second largest log likelihood value by an extremely small margin. If this happens, we manually adjust the optimal starting parameters to yield more plausible results, namely by choosing the starting values which result in the next largest log likelihood value.

[22] We vary the most important parameters, namely the long memory parameter, the GARCH parameter and the ARCH parameter for each data series / subseries, with steps of 0.2 between zero and one.

6. Empirical results

In our first estimation of the abovementioned model with the inclusion of the FED's primary credit rate, as a proxy for monetary policy, we have not found any significant exposure of the volatility series to monetary policy. Basically, this means that the gamma parameter of the following FIGARCH equation is insignificant:

$$\sigma_t^2 = \omega[1-\beta(L)]^{-1} + \{1-[1-\beta(L)]^{-1}\phi(L)(1-L)^d\}\eta_t^2 + \gamma x_t \tag{34}$$

To crosscheck our results, we decide to re-run the estimation of all our series, but change the explanatory variable to the interest rate of the European Central Bank (n.d.). We again find no significant impact of the monetary policy on our seven conditional variance equations. Our explanation for that is that the monetary policy proxies change not very frequently and therefore are too close to a constant to significantly influence the conditional volatility of our seven return series. As monetary policy rates, no matter which one, are not changed on a day-to-day basis, we believe that it is very unlikely to find an appropriate proxy, that properly displays the monetary policy decisions and is also changed frequently enough to have a significant influence on the conditional variance. Therefore, we decide to re-estimate our model excluding the explanatory variable in the FIGARCH equation, as this procedure will give us more precise parameter estimates that are not altered by an insignificant variable.

Through preliminary estimation we also notice that, a vast majority of the series would turn out to display a constant term in their variance equation that is not significantly different from zero. The inclusion of a constant term did in fact raise the conditional variance series of some of our chosen sub-periods on a whole by an amount that could be considered less than plausible when the very sub-period is supposed to display less volatility. Thus, in order to better both the conditional variance and the parameter estimations, we decide to also exclude the constant term from the variance equation. This is done in practice by setting the value of ω to zero manually before the estimation. The variance equation specification is thereby:

$$\sigma_t^2 = \{1-[1-\beta(L)]^{-1}\phi(L)(1-L)^d\}\eta_t^2 \tag{35}$$

The same phenomenon, however, is not observed on the constant μ of the mean equation, as some of our selected assets do appear to have a statistically significant non-zero mean. On top of that, the focus of this paper is on the volatility of our chosen assets, which, in essence, would not vary, whether or not we subtract a constant from all the returns of a sub-period alike. Thus, the constant term in the mean equation is kept.

6.1 FIGARCH Results

The FIGARCH model in general functions in a fine fashion after certain adjustments as mentioned previously in the estimation process when it comes to our case. It is able to yield an overall satisfying estimation of the conditional variances of our selected cryptocurrencies and the fiat currency Euro. We report the parameter estimates of the FIGARCH model consisting of the return equation as specified by equation (21) and the variance equation as specified by equation (35) in APPENDIX G, and the plotted conditional variance series accompanied by the corresponding return series in APPENDIX F.

We start with analyzing the results for Bitcoin, arguably the most sought-after cryptocurrency at the moment. The pattern of the conditional variance series modelled for Bitcoin is more or less in line with what we can expect when looking at the parameter estimates as reported in Table G.1. The long memory parameter estimates of the first two sub-periods for Bitcoin are close to each other; though period 2 of Bitcoin appears to possess a higher GARCH effect parameter, which in theory would mean the series should be smoother. The ARCH effect parameter is what plays a larger role when comparing these two periods; period 2 presents itself to be much more affected by shocks as it is estimated to have a larger and statistically significant ARCH effect parameter. Period 3 for Bitcoin, however, is estimated to have both a large GARCH effect term and long memory parameter, which suggests the variance series would be smoother, with the effects of shocks to be more persistent. Apart from the single huge spike in early 2020 (which one could argue to be an outlier and possibly related to the COVID-19 pandemic), this is to a large extent the case that we can see from the plot in APPENDIX F. It is also worth noticing that, the range of conditional variances modelled for Bitcoin is much smaller than the rest of the selected cryptocurrencies, but still significantly larger than that of the Euro.

According to the graph plotted with the estimated conditional variance, the Euro experienced two periods of relative tranquility in the first two sub-samples, also reflected in the high GARCH effect estimates. Still, the first period, as dominated by a gradual increase in the exchange rate, is seen to be slightly more affected by shocks, which provides evidence for a supposed higher ARCH effect term. The dynamics of the Euro is also in line with a sudden increase in the long memory term from sub-period 1 to sub-period 2. The last sub-period of the Euro, approximately taking place simultaneously with the start of the pandemic, is highlighted by its higher volatility; the conditional variance series is thus no longer as smooth as for the first two periods, alongside with a large decrease in the estimates of the GARCH effect parameter and the long memory parameter. Another interesting finding about the parameter estimates of the Euro is that, it seems not particularly straightforward to distinguish the Euro from cryptocurrencies by merely judging from these estimates, without knowing which is which first, as the parameters do not differ in an easily recognizable way. Still, one should notice that, the Euro is considerably much less volatile than all the cryptocurrencies that we have selected, for the range of its conditional variance is much narrower.

Comparing the conditional volatility graphs of the Euro and the Bitcoin, we find some quite interesting results. Although we have detected different structural breaks for both series, we still see some resemblance regarding the trend in their respective plots. Both series start off with a rather volatile period which is to calm down halfway through 2018. This following rather smooth series lasts until the beginning of the COVID crisis in January 2020, when a big negative return spike increases volatility tremendously. Obviously, there is still a difference between the two series as the smooth period of the Euro is a lot smoother than that of Bitcoin and the volatile one much less volatile than that of Bitcoin. This is due to the fact that, in spite of the resemblance in trend, the two series move in totally different ranges. The Euro is in general much less volatile than the Bitcoin, which is due to its much smaller changes in price and hence also in returns. Looking at the other cryptocurrencies, we see that the trend resemblance holds for all six volatility series, whereas the range of the volatility differs even more for other cryptocurrencies.

Ripple, Binance Coin and Cardano are the three cryptocurrencies to display a much wider range in their conditional variances, even compared to other cryptocurrencies, mainly from their first sub-periods. This comes as no real surprise as these three are from the cryptocurrencies with relatively smaller market capitalizations and had been of much smaller value per unit for a large part of the time since they first came into being, which means that during those times, even a "small" change in their values in absolute terms, would be seen as a large deviation in their returns (We could also see this from Table 4.1 as they are the three assets with the largest range in returns). This is in line with the findings of Cheikh et al. (2020), who state that altcoins with a relatively smaller market capitalization tend to fluctuate in a wider range than Bitcoin, which has by far the largest market capitalization.[23] The point mentioned above is also reflected by the large ARCH effect parameter estimates for the first sub-period of these three cryptocurrencies. Then these coins go on to experience a smoother period in their second sub-periods, as depicted by the smooth plots in APPENDIX F (smoother when compared to the first period though; they more or less just return to the "typical" levels of conditional variances for

[23] Bitcoin alone makes up 1.1 trillion USD of the total market capitalization of roughly 1.45 trillion USD of all six selected cryptocurrencies.

cryptocurrencies during this time). This finding is in line with the abovementioned trend of cryptocurrencies in general. After a relatively minor spike due to the COVID crisis and the following return to the previous pattern of smoothness, a second much larger spike is observed. This disruption that is the second structural break of the three coins in late 2020 can be explained by the increase in their prices starting earlier this year.

One could see, that almost every cryptocurrency that we have chosen has displayed said feature of increase in prices from the start of 2021, and question the reason why this nearly simultaneous structural break is not detected for the other assets. We assume that the reason lies in the fact that they are the "smaller coins" and less sought-after than the larger ones, which means that the large spikes that we see in their supposed third sub-periods are more important in determining their structural changes than for "larger coins". For the latter such spikes are sort of becoming the "norms" and thus these increases from earlier this year might happen not to be a deciding factor when looking for structural changes. This is, to be frank, also a limitation from only allowing for up till two breakpoints. The BIC could potentially get even lower with more breakpoints allowed for, and at the same time, it would be made possible for the potential break points, which are not yet recognized when only allowing for a smaller number of break points, to be determined to be places where structural changes happen.[24] These earlier mentioned changes in trend for Ripple, Binance Coin and Cardano are also shared by Litecoin, a cryptocurrency also smaller in scale when compared to Bitcoin, though especially in the first sub-period to a lesser extent. Still, we are glad to maintain that considering structural changes yields even better results for these smaller altcoins, for the somewhat obvious reason that their smaller scale in market capitalizations means that they are more susceptible to structural changes, and that in turn, considering structural changes when modelling volatility is suitable to these coins in nature.

As discussed previously, all seven assets suffer significant negative returns at the beginning of 2020 – the early stage of the global pandemic that we still find ourselves in. However, this is only translated into a structural change in Ethereum, Litecoin and the Euro while not for the other four assets. Ethereum, Litecoin and the Euro all have a decrease in their GARCH effect parameter estimates after this break point, which is to be expected since one could assume negative shocks such as a pandemic would induce more volatile behaviors from abovementioned assets. The long memory parameter estimates for these three assets are found to drop after this structural change as well – this suggests that even the effect of COVID did

[24] Though this would come along with the dilemma that more structural changes allowed for mean less estimation accuracy, more effort put into interpretations and more complexity.

not last exactly as long as some would imagine, as the shocks' impact on these assets' volatility turns out to be less persistent from then. By referring to the conditional variance plots, we could see that Bitcoin and Ethereum (the largest two cryptocurrencies we have used), and the Euro, have their conditional variance relatively significantly influenced at that point in comparison to the shocks within the same series, which justifies the existence of these structural changes in Ethereum and Litecoin (while for Bitcoin it is only determined to be a break point if more structural changes are allowed). In contrast, for the smaller coins that, without exception, also suffered losses during that time, their spikes in the conditional variance plots then appear to be less significant. This is generally in line with what we can see from the return graphs, in that the larger coins and the Euro see a large negative impact in returns in comparison to other negative returns of their own. For the smaller coins, although they are also pretty much influenced by the pandemic, as shown from the negative return spikes, there exist other high spikes that make these very spikes comparably "less important". This could be evidence that larger cryptocurrencies, given their more mature nature, are more correlated with fiat currencies, through shocks which could traditionally impact conventional assets, than smaller cryptocurrencies. Still, one could argue that for the cryptocurrencies smaller in volume, the time for the early stage of the pandemic lies within a sub-period that, for the most time (apart from said period of early 2020), has an extremely smooth conditional variance series. So, the parameters estimated for this period would naturally encourage the series to be smoother and with lower spikes, thus "dragging" the potential spikes downwards. Yet the reality might be more volatile than the modelled variances denote. One potential solution to this would be to allow for more break points in the Bai-Perron test. However, we should also take into consideration that the shocks brought about by the pandemic appear to be evanescent, which means that the sub-period that would be produced by the potential break point is most likely to be transient as well. Consequently, chances are that the parameter estimates would be less likely to be reliable when used to make references and hence would not necessarily improve our results in whole.

Looking back at the results in general we find a significant resemblance of the volatility graphs between all cryptocurrencies and the Euro, with a main difference in their respective range of returns and hence volatilities. Furthermore, we realize that market capitalization of cryptocurrencies plays a role in defining volatility of such assets, as "smaller coins" are influenced more heavily by relatively small price changes compared to "larger coins". We are also realistic with our obtained plots and parameters, and acknowledge certain shortcomings that might result in biased inference. The fact that we limit each series to two breakpoints most likely depresses shocks in the extremely smooth second period and thus might give a misleading view of the impact of the COVID crisis on the volatility.

6.2 Risk Measure Results

We compute our risk measures, namely the value at risk, the expected shortfall and the stressed expected shortfall, for all six cryptocurrencies and the one fiat currency at the 97.5% confidence level using the volatility weighted historical simulation (VWHS). Therefore, we should be able to incorporate the then-current market conditions, in form of the conditional volatility series, into the risk measures. In APPENDIX H we plot the different values for VaR and ES using a 250-day rolling window each, to measure the risk over the corresponding previously passed year. Concentrating on the displayed values we find that VaR and ES have almost a similar pattern within one asset, besides their differences in scale. This result is to be expected as the ES simply measures the right tail of a loss distribution (the losses larger than the VaR) and therefore is per definition at least as large as the VaR. One could argue that the ES estimate therefore does not add too much value to the VaR, as it could be obtained by simply upscaling the VaR estimate. This argument can be countered by the subadditivity axiom which is not fulfilled under the VaR, as mentioned in the methodology part of this paper, but is fulfilled by the ES. Thereby the ES has beneficial diversification properties, which the VaR lacks, and thus can be seen as an important addition under the scope of portfolio allocation and portfolio risk.

Starting with the VaR plot of Bitcoin, which can be found in APPENDIX H, the largest spike in the graph is found on the 2020-03-13 with a VaR estimate of 30.6. This spike coincides for one with the spike in the conditional volatility series at the same date and the negative spike in the return series on the day before.[25] This is no coincidence but expected, as a high negative return implies a high loss and the first date this loss can be displayed in the VaR measure is the day after as the VaR is computed over a rolling window of the 250 previous days. Furthermore, it is logical that this increase in VaR slowly decreases, as the losses are rescaled with volatility and the highest volatility estimate coincides with the highest VaR value. Moving along in time, the volatility decreases and therefore the weight assigned to the largest loss is lower. Comparing the general movement of the VaR plot we find an extremely high resemblance with the conditional volatility plot, which is caused by the VWHS. Furthermore, we can relate the other spikes in the VaR plot to negative return spikes in the return series. On an absolute scale the VaR estimates of Bitcoin vary between a minimum of 2.7 on the 2018-11-14 and the abovementioned maximum of 30.6. The whole comparison amongst variance, return and VaR

[25] Both, the conditional volatility graph and the return graph can be found in APPENDIX F.

could also be made with the ES plot as it very closely follows the trend of the VaR, with an upwards shift on the scale. Therefore, the absolute values of the ES vary between a minimum of 4.0 on the 2018-11-14 and a maximum of 57.4 on the 2020-03-13. We will later find that this maximum ES estimate actually represents the stressed ES estimate within our timeframe.

The same procedure could be done for all other cryptocurrencies, but we suspect that there is not much insight to be found in doing so. Comparing the VaR plots with the according conditional volatility graph and the return plot directly shows, that the relation we found for Bitcoin holds true for all other cryptocurrencies as well. This does not come as a surprise, as under the definition of the VWHS this result is exactly what we expect. Obviously, the same holds true for the ES plot, as they follow the same trend as mentioned above. What is of more interest to us is to compare the absolute values of the risk measures between the cryptocurrencies and our one fiat currency, the Euro. To do so we start with an in-depth analysis of the risk measures for the Euro in a similar fashion to how it was done for Bitcoin and then use the last day of our timeframe as an example to compare absolute values of the risk measures.

For the VaR plot of the Euro, the largest spike in the VaR amounts to 1.67 on the 2020-03-23 while the smallest VaR estimate equals 0.46 on the 2018-10-15. The big spike coincides with the largest spike in the conditional volatility series on the same day and also with the largest loss in the return series which occurred on the 2020-03-12, a few days before the highest VaR estimate. Again, this is no surprise as the VaR depends mainly on the amount of the loss, the volatility on the date of the loss and the one day ahead forecast, the conditional volatility for the day of which we calculate the VaR. Looking at the remaining trend of the graph we find again that it very closely resembles the trend of the conditional volatility plot. The actually interesting part of this interpretation is when we look at the range in which the risk measures move for the Euro. For all cryptocurrencies the respective lowest VaR estimate is still higher than the highest VaR estimate for the Euro. The same holds true for the ES shortfall estimates. This insight shows that the Euro in general is way less risky than the six cryptocurrencies in question. This result follows directly from the scale of the returns and therefore losses of the Euro. The returns of the Euro move way closer around the zero mean[26] which can be seen in the scale of its return plot in APPENDIX F compared to the other return plots of the six cryptocurrencies. Accordingly, also the volatility is much smaller in the Euro, but this does not influence the scale of the VaR as in the rescaling process the volatility is in the numerator as

[26] See the insignificant return intercept coefficients for EUR 1, EUR 2 and EUR 3 in Table G.7.

the one day forecast and in the denominator as the volatility on the day of the loss (see equation (18)), thus cancelling each other out to a certain extent.

In table 6.1 the three different risk measures are reported for the latest day in our timeframe, the 2021-03-31. This is done to quantify the risk measures for a certain day and thereby allow for easier interpretation of the absolute values.[27] We also include the averages for VaR and ES of each asset across our chosen period in the table.

Table 6.1 Risk Measures on the 2021-03-31 over a 250 day period

	BTC	ETH	XRP	LTC	BNB	ADA	EUR
VaR	7.22	8.75	9.74	10.28	15.25	13.84	0.78
ES	11.07	12.41	14.65	13.70	22.17	16.85	0.88
Average VaR	9.26	13.34	12.59	12.50	11.82	13.39	0.73
Average ES	15.32	18.95	18.27	16.90	17.82	20.86	0.89
Stressed ES	57.43	77.70	144.78	50.68	118.32	99.89	2.06

Note: This table reports the reports the risk measures for all seven loss series on the 2021-03-31. The displayed results are **(1) VaR** (the value at risk estimate for the 2021-03-31), **(2) ES** (the expected shortfall estimate for the 2021-03-31), **(3) Average VaR** (the average VaR estimate over the whole timeframe in which VaR estimates are computed), **(4) Average ES** (the average expected shortfall estimate over the whole timeframe in which ES estimates are computed) and **(5) Stressed ES** (the stressed expected shortfall estimate for the 2021-03-31).

First of all, it is important to note that the 2021-03-31 is a day when the risk measures are on a low point in absolute terms compared to the average values for all assets besides Cardano and Binance Coin. Therefore, the stressed ES will be a lot higher than the ES of the current point in time. We find that the VaR is numerically the smallest of the three risk measures and that the ES is strictly larger than the VaR. These results underline the findings we made when comparing the scale of the plots and are expected. Comparing the relative difference of average VaR to average ES, we find that that of the Euro is quite smaller than those of the six cryptocurrencies. This result underlines that the fiat currency has fewer extreme losses at the right tail of the loss distribution compared to the cryptocurrencies and is therefore less risky in extreme situations.

Furthermore, we are now able to incorporate the stressed ES, as modelling it would simply yield a constant as it does not change over time unless there exists another time period of 250 days that is more severely stressed than the current most stressed period. Looking at the obtained estimates for the stressed ES, we directly see that it is much larger than all the other estimates. For Ripple the stressed ES is roughly about ten times as high as the ES of the current period, whereas it is about four times the current expected shortfall of Litecoin. The other

[27] As a supplement, APPENDIX I displays the VaR and ES estimates at the beginning of each year throughout our sample.

cryptocurrencies lie in between these two multiples. For the Euro the stressed ES is a bit more than twice the current ES, showing that even under the most stressed market conditions the fiat currency is not nearly as risky as the cryptocurrencies. Looking at the absolute values of the stressed ES we see that the values for the cryptocurrencies lie in between 57.43 for Bitcoin and 144.78 for Ripple, the stressed ES estimate for the Euro is 2.06. Similar results are found when comparing the stressed ES with the average ES of each asset. These findings strengthen the above explained differences between cryptocurrencies and fiat currencies, as the cryptocurrencies display a much higher amount of market risk. This again can be explained by the scale of the underlying return series. Seeing the huge differences in the scale of the risk measures of cryptocurrencies and the fiat currency, we are able to draw some clear conclusions. For once the fiat currency is found to be less risky as the return series moves in a much smaller range than the return series of the cryptocurrencies. This gives reason to believe that cryptocurrencies should not be interpreted as a normal currency but rather as a risky investment as authors like Dyhrberg (2016), Fang et al. (2019) and Liu and Tsyvinski (2018) suggested previously as well. Furthermore, we find more extreme loss cases in the distribution of the cryptocurrencies compared to the one of the Euro. The implementation of the stressed ES manifests that gap between the Euro and the six cryptocurrencies again, as the cryptocurrencies suffer a lot more under extremely stressed market conditions compared to the fiat currency.

7. Conclusion

Generally speaking, the FIGARCH model (especially with its long memory parameter), together with the structural breaks, provides us with more flexibility when modelling volatility and satisfyingly fulfills our needs. It is able to yield parameter estimates and conditional variance series that are in line with what we have expected from the return series, and their behaviors prior to and after the supposed break points of our chosen cryptocurrencies and the Euro. From the modelled conditional variances, we are able to detect both similarities in trend as well as disparities in range of the volatilities. Although we did not find significant influence of monetary policy on the variance equation, we find that the market value plays a part in tracing the influence of price movements on volatility. However, we also recognize the dilemma resulting from only allowing for only a limited number of structural changes due to time limit. Nonetheless, we deem that our current setting is still one of the most feasible options that we could come across in the given timeframe.

Analyzing the assets in question using the VaR and ES yields that in general all cryptocurrencies display a much higher market risk than the Euro. The ES hereby follows the same trend as the VaR, as it captures the right tail of the loss distribution. We find that cryptocurrencies have a relative higher ES compared to the VaR, in contrast to the Euro. Both risk measures resemble the trend of the according conditional volatility model, which states the importance of accounting for the then-current market conditions when measuring risk. With these results in mind we can clearly differentiate between the fiat currency from the six cryptocurrencies, supporting the classification of cryptocurrencies as risky investment but not a currency. With the inclusion of the stressed expected shortfall estimate, these results are further strengthened. Cryptocurrencies suffer under stressed market conditions to a significantly larger extend than the Euro, as their stressed ES estimates differentiates the two even further. Under the upcoming Basel III regulation and the thereby included implementation of the stressed expected shortfall, cryptocurrencies will have an even harder stance regarding their quantified risk.

Regarding further work, there are certain improvements that could be made to our model, to improve the accuracy of the volatility. For one, we use daily price and therefore also daily returns in our paper, as the underlying data to model the volatility series. With cryptocurrencies and their extremely frequent trading pattern and price movements within a day, it would be feasible to use intraday data. Thereby, one would be able to capture volatility differences, not only from day to day, but also across different times throughout a day. This could be done using heterogeneous autoregressive models, as they are able to incorporate such intraday data.

Another problem we came across in our paper was the estimation of the FIGARCH model. Using the maximum likelihood approach has quite a few drawbacks as the global maximum cannot be reached with consistency[28], which means often the optimized values represent a local maximum and not the global one. To enhance the parameter significance and to ensure that the parameters which represent the global maximum are reached, a different estimation method like simulated annealing could be used. An aspect that is mostly neglected in this paper is the covariance and co-movements between different cryptocurrencies and fiat currencies, which could yield important insights on how the volatility series do influence each other. Including such in a future paper, might also bring about insights on what other explanatory variables influence the conditional volatility of cryptocurrencies.

[28] At least with the computational power that was at hand for us.

References

Adhami, S., Giudici, G., & Martinazzi, S. (2018). Why do businesses go crypto? An empirical analysis of initial coin offerings. *Journal of Economics and Business, 100*, 64-75.

Andrada-Félix, J., Fernandez-Perez, A., & Sosvilla-Rivero, S. (2020). Distant or close cousins: Connectedness between cryptocurrencies and traditional currencies volatilities. *Journal of International Financial Markets, Institutions and Money, 67*, 101219.

Ardia, D., Bluteau, K., & Rüede, M. (2019). Regime changes in Bitcoin GARCH volatility dynamics. *Finance Research Letters*, *29,* 266-271.

Artzner, P., Delbaen, F., Eber, J. M., & Heath, D. (1999). Coherent measures of risk. *Mathematical finance, 9(3)*, 203-228.

Bai, J., & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 47-78.

Bai, J., & Perron, P. (2003). Critical values for multiple structural change tests. *The Econometrics Journal, 6(1)*, 72-78.

Baillie, R. T., Bollerslev, T., & Mikkelsen, H. O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of econometrics, 74(1)*, 3-30.

Basel Committee on Banking Supervision. (2014). Fundamental review of the trading book: A revised market risk framework, Bank for International Settlements, Available online: https://www.bis.org/publ/bcbs265.pdf [Accessed: 7 April 2021]

Baur, D. G., & Dimpfl, T. (2018). Asymmetric volatility in cryptocurrencies. *Economics Letters*, *173*, 148-151.

Baur, D. G., Dimpfl, T., & Kuck, K. (2018). Bitcoin, gold and the US dollar–A replication and extension. *Finance Research Letters, 25*, 103-110.

Caporale, G. M., & Zekokh, T. (2019). Modelling volatility of cryptocurrencies using Markov-Switching GARCH models. *Research in International Business and Finance, 48*, 143-155.

Cheikh, N. B., Zaied, Y. B., & Chevallier, J. (2020). Asymmetric volatility in cryptocurrency markets: New evidence from smooth transition GARCH models. *Finance Research Letters, 35*, 101293.

Chu, J., Chan, S., Nadarajah, S., & Osterrieder, J. (2017). GARCH modelling of cryptocurrencies. *Journal of Risk and Financial Management, 10(4)*, 17.

Coinmarketcap. (n.d.). Available online: https://coinmarketcap.com/ [Accessed: 3 April 2021]

Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association, 74(366a)*, 427-431.

Dyhrberg, A. H. (2016). Bitcoin, gold and the dollar–A GARCH volatility analysis. *Finance Research Letters, 16*, 85-92.

European Central Bank. (n.d.). Available Online: https://www.ecb.Europa. eu/stats/policy_and_exchange_rates/key_ecb_interest_rates/html/index.en.html [Accessed: 4 May 2021]

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the econometric society*, 987-1007.

Fang, L., Bouri, E., Gupta, R., & Roubaud, D. (2019). Does global economic uncertainty matter for the volatility and hedging effectiveness of Bitcoin?. *International Review of Financial Analysis, 61*, 29-36.

FED. (n.d.). Federal Reserve, Available online: https://www.frbdiscountwindow.org /pages/discount-rates/historical-discount-rates [Accessed: 25 April 2021]

Frittelli, M., & Gianin, E. R. (2002). Putting order in risk measures. *Journal of Banking & Finance, 26(7)*, 1473-1486.

Fuller, W. A. (1976). *Introduction to statistical time series (Vol. 428)*. John Wiley & Sons.

Ghalanos, A. (2020). Introduction to the rugarch package. (Version 1.4-3). Available online: https://cran.r-project.org/web/packages/rugarch/vignettes/Introduction_to _the_rugarch_package.pdf [Accessed: 25 April 2021]

Hull, J. (2018). Risk management and financial institutions, John Wiley & Sons

Inclan, C., & Tiao, G. C. (1994). Use of cumulative sums of squares for retrospective detection of changes of variance. *Journal of the American Statistical Association*, *89(427),* 913-923.

Jarque, C. M., & Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics letters, 6(3),* 255-259.

Katsiampa, P. (2017). Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters, 158*, 3-6.

Katsiampa, P. (2019). An empirical investigation of volatility dynamics in the cryptocurrency market. *Research in International Business and Finance, 50*, 322-335.

Klein, T., Thu, H. P., & Walther, T. (2018). Bitcoin is not the New Gold–A comparison of volatility, correlation, and portfolio performance. *International Review of Financial Analysis, 59*, 105-116.

Liu, Y., & Tsyvinski, A. (2018). Risks and returns of cryptocurrency (No. w24877). *National Bureau of Economic Research.*

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

McLeod, A. I., & Li, W. K. (1983). Diagnostic checking ARMA time series models using squared-residual autocorrelations. *Journal of time series analysis, 4(4)*, 269-273.

Mensi, W., Al-Yahyaee, K. H., & Kang, S. H. (2019). Structural breaks and double long memory of cryptocurrency prices: A comparative analysis from Bitcoin and Ethereum. *Finance Research Letters, 29*, 222-230.

Myung, I. J. (2003). Tutorial on maximum likelihood estimation. *Journal of mathematical Psychology, 47(1)*, 90-100.

Nguyen, T. V. H., Nguyen, B. T., Nguyen, K. S., & Pham, H. (2019). Asymmetric monetary policy effects on cryptocurrency markets. *Research in International Business and Finance, 48*, 335-339.

Phillip, A., Chan, J., & Peiris, S. (2019). On long memory effects in the volatility measure of cryptocurrencies. *Finance Research Letters, 28*, 95-100.

Rapach, D. E., & Strauss, J. K. (2008). Structural breaks and GARCH models of exchange rate volatility. *Journal of Applied Econometrics, 23(1),* 65-90.

Shen, D., Urquhart, A., & Wang, P. (2020). Forecasting the volatility of Bitcoin: The importance of jumps and structural breaks. *European Financial Management, 26(5)*, 1294-1323.

Sjölander, P. (2011). A stationary unbiased finite sample ARCH-LM test procedure. *Applied Economics*, *43(8)*, 1019-1033.

Szegö, G. (2002). Measures of risk. *Journal of Banking & finance, 26(7)*, 1253-1272.

Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017, June). An overview of blockchain technology: Architecture, consensus, and future trends. *In 2017 IEEE international congress on big data (BigData congress)* (pp. 557-564). IEEE.

APPENDIX A: Maximum Likelihood Estimation

Under the maximum likelihood estimation, parameters of a certain underlying probability distribution are estimated by maximizing a likelihood function in such a way that the observed data is most likely to occur. It requires the number of observations to be large and their distribution to be known, as well as observations themselves to be independent and with the same probability distribution. Following our example of a simple mean model with normal distributed errors from section 2.5, we have to maximize the following log likelihood function to obtain the parameter estimates:

$$\max_{\mu,\omega,\alpha,\beta} \ln L(\mu,\omega,\alpha,\beta) = \sum_{t=1}^{T}\left(-\frac{1}{2}\ln(2\pi) - \frac{1}{2}\ln(\sigma_t^2) - \frac{\eta_t^2}{2\sigma_t^2}\right) \tag{36}$$

This can be done as the ln transformation of the likelihood function does not change the location of the maximum (maximizing coefficient values). For more insights into the maximum likelihood estimation method, underlying assumption and how the likelihood functions are derived from probability density functions we refer to the paper of Myung (2003).

APPENDIX B: Price and Return Series[29]

Bitcoin:

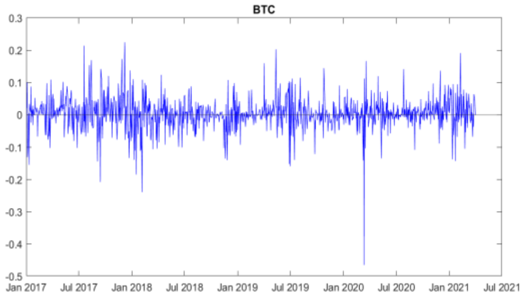

Ethereum:

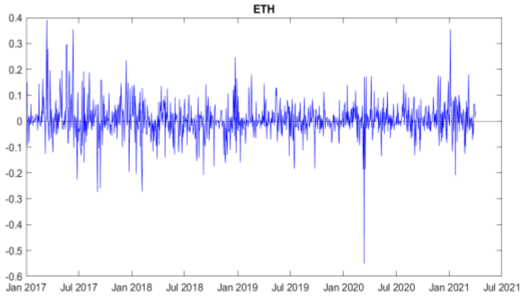

Ripple:

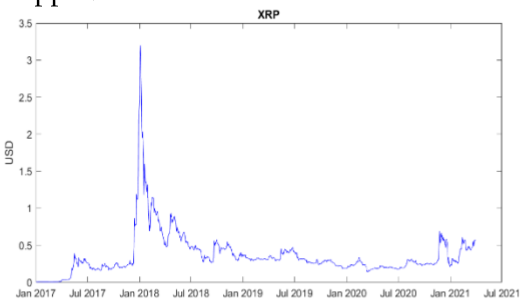

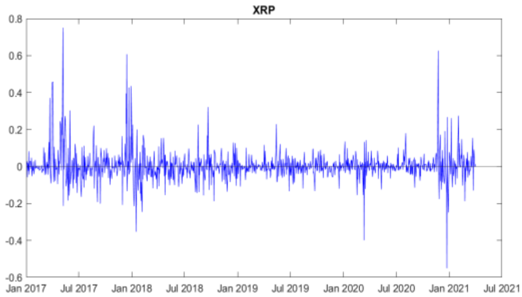

Litecoin:

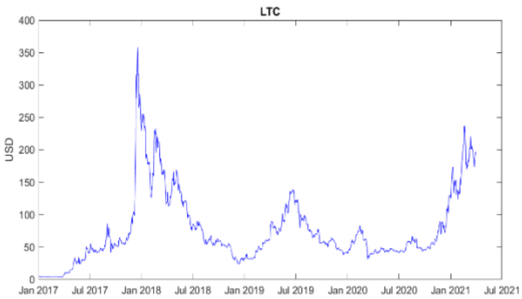

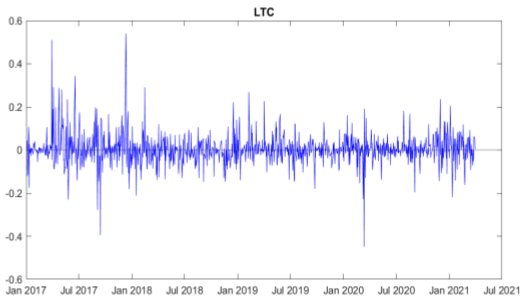

Binance Coin:

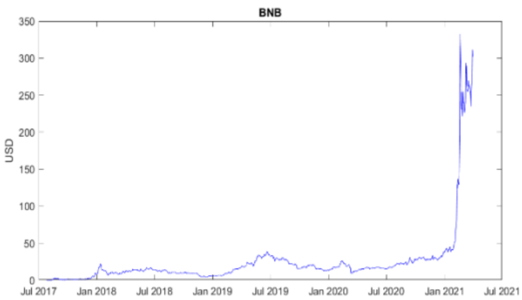

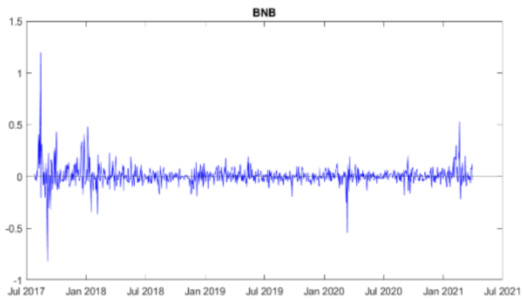

[29] The scale of the different plots differs, to display the data as detailed as possible.

Cardano:

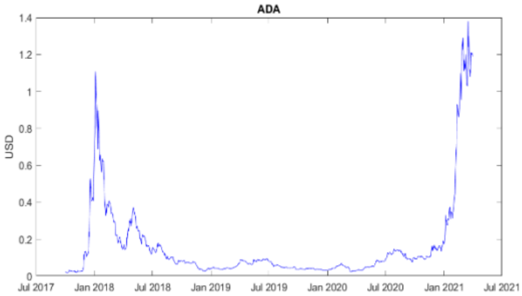

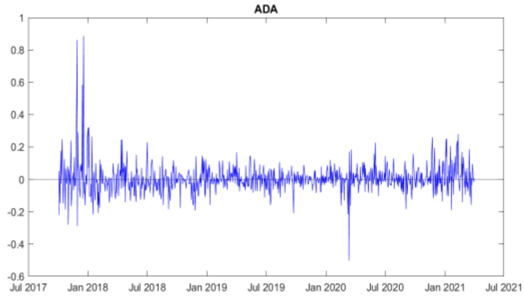

Euro:

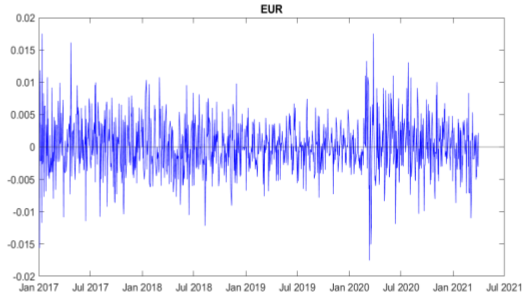

APPENDIX C: Histograms

Bitcoin:

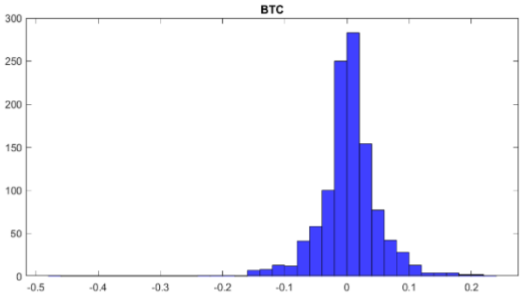

Ethereum:

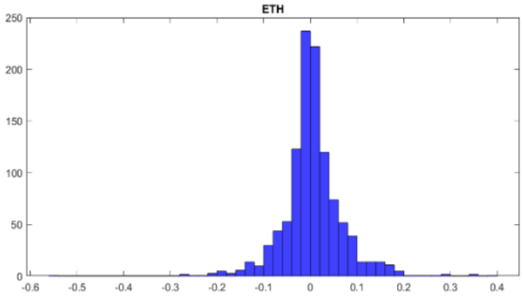

Ripple:

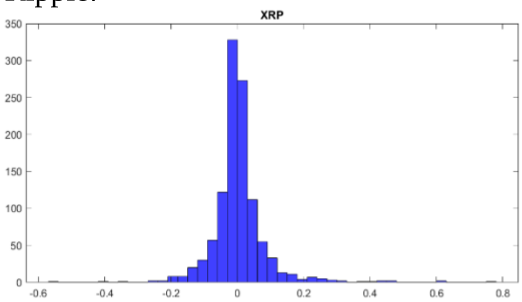

Litecoin:

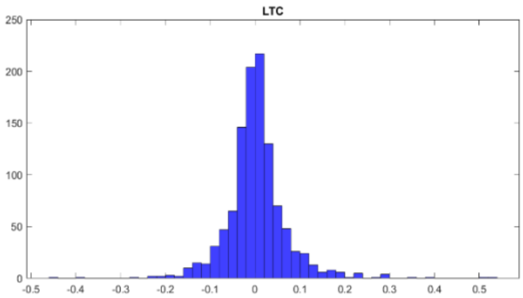

Binance Coin:

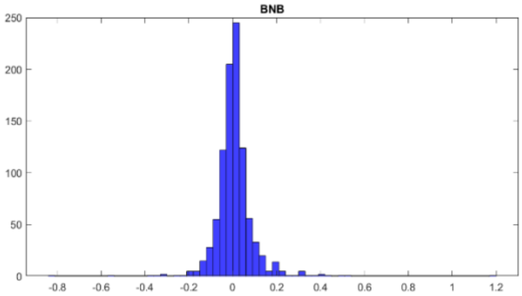

Cardano:

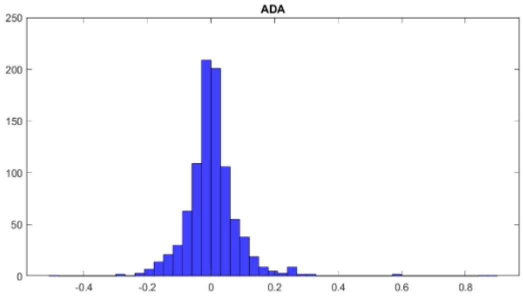

Euro:

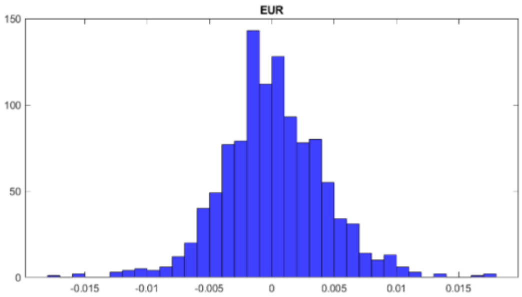

APPENDIX D: Bai-Perron Test

Table D.1 Bitcoin:

m	0	1	2
RSS	0.00268	0.002242	0.00192
BIC	-11160	-11340	-11500
Breakdates:			
m = 1	2018-03-20		
m = 2	2017-08-22	2018-04-11	

Table D.2 Ethereum:

m	0	1	2
RSS	0.007558	0.005131	0.0047
BIC	-10010	-10430	-10510
Breakdates:			
m = 1	2018-02-20		
m = 2	2018-02-23	2020-01-30	

Table D.3 Ripple:

m	0	1	2
RSS	0.06117	0.04317	0.03781
BIC	-7697	-8069	-8201
Breakdates:			
m = 1	2018-02-26		
m = 2	2018-02-27	2020-08-11	

Table D.4 Litecoin:

m	0	1	2
RSS	0.01307	0.006169	0.006
BIC	-9405	-10220	-10250
Breakdates:			
m = 1	2018-01-25		
m = 2	2018-01-29	2020-01-30	

Table D.5 Binance Coin:

m	0	1	2
RSS	0.2154	0.1043	0.1024
BIC	-5327.86	-6010.93	-6014.4
Breakdates:			
m = 1	2018-02-13		
m = 2	2018-02-13	2020-09-10	

Table D.6 Cardano:

m	0	1	2
RSS	0.1145	0.05153	0.0502
BIC	-5591	-6305	-6315
Breakdates:			
m = 1	2018-04-11		
m = 2	2018-04-11	2020-09-21	

Table D.7 Euro:

m	0	1	2
RSS	8.55E-08	7.436E-08	5E-08
BIC	-22620	-22760	-23180
Breakdates:			
m = 1	2020-01-30		
m = 2	2018-09-21	2020-01-29	

Note: Tables D.1 to D.7 report the results of the Bai-Perron tests for multiple structural breaks on the seven sample volatility series. We present the **(1) RSS** (residual sum of squares) as well as the **(2) BIC** (Bayesian information criteria) for each of the seven series. The choice of the number of breakpoints is made according to the lowest BIC value for each volatility series. The RSS displays the discrepancy between the sample / model and the actual data and could also be used for the selection of number of breakpoints (the lower the RSS the better the fit). Furthermore, the dates of the structural breaks are displayed for either one or two structural breaks within the data.

APPENDIX E: Structural Breaks - Sample Volatility and Return Series[30]

Bitcoin:

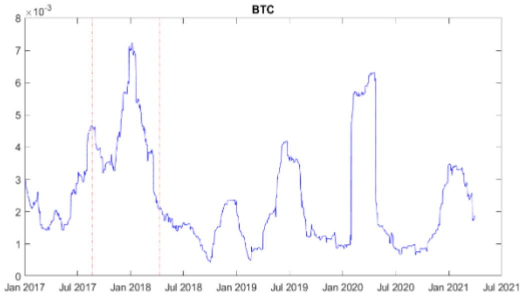

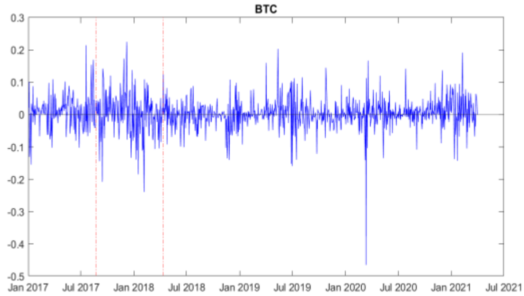

Ethereum:

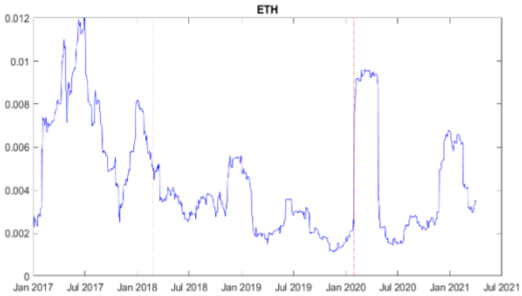

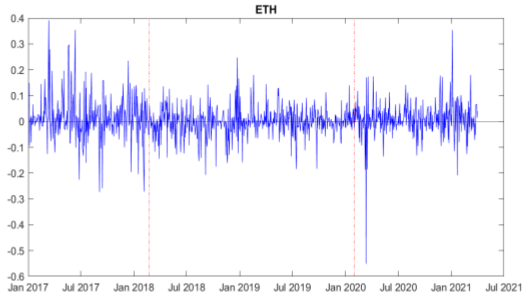

Ripple:

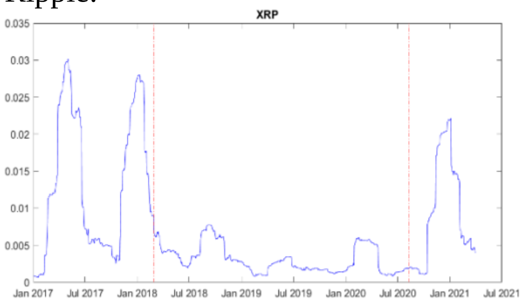

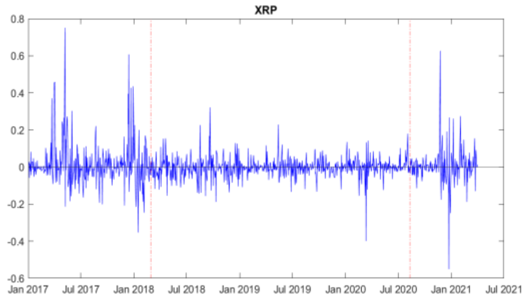

Litecoin:

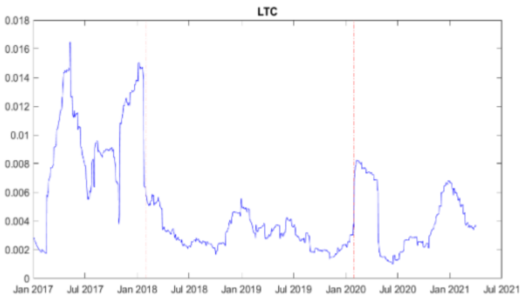

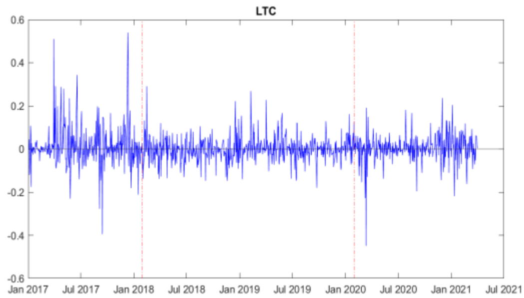

Binance Coin:

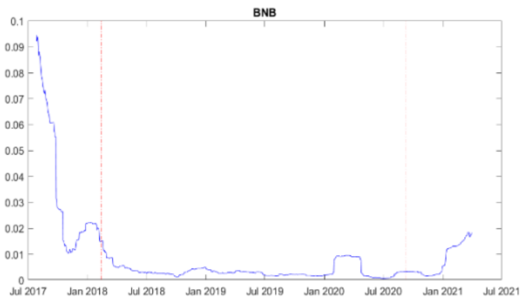

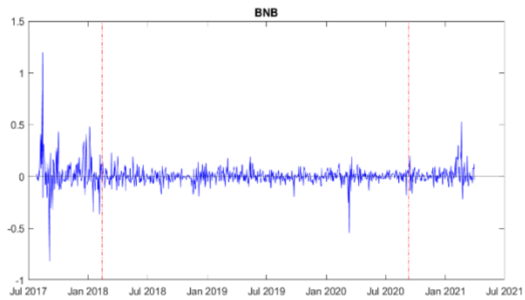

[30] The red dash-dotted lines display the structural breaks indicated by the Bai-Perron test.

Cardano:

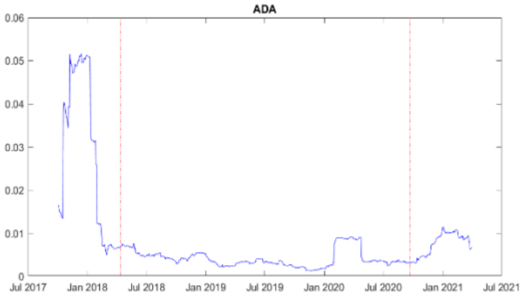

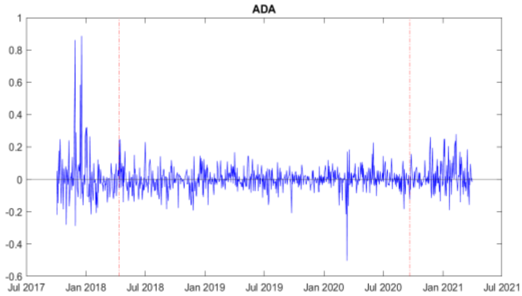

Euro:

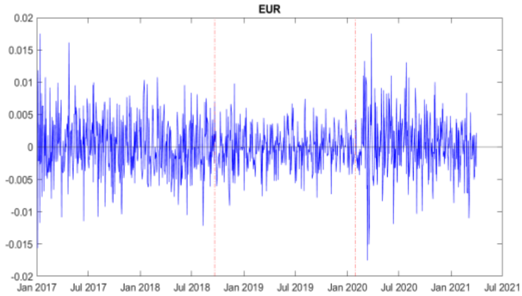

APPENDIX F: FIGARCH - Conditional Volatility and Return Series

Bitcoin:

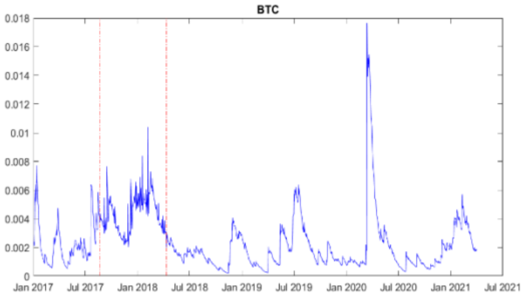

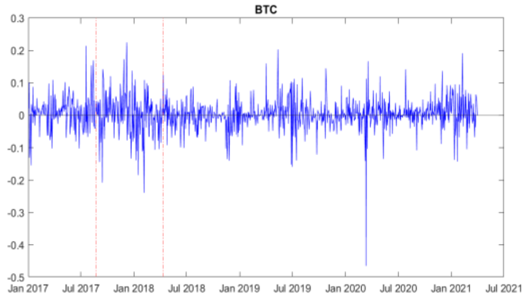

Ethereum:

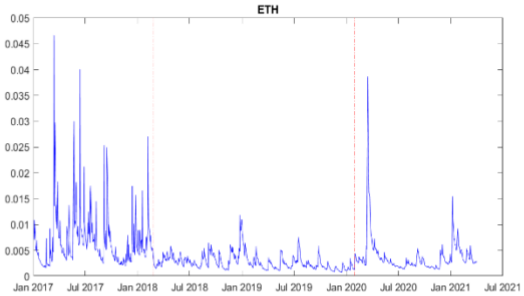

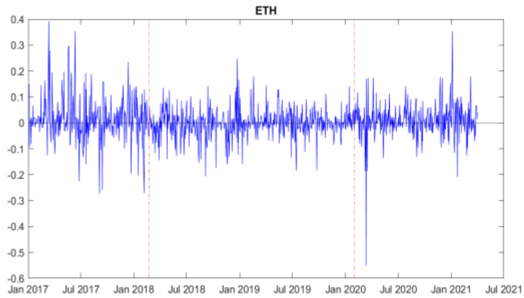

Ripple:

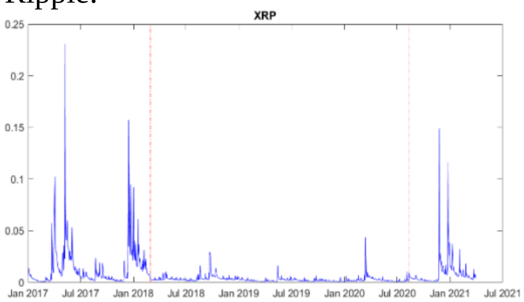

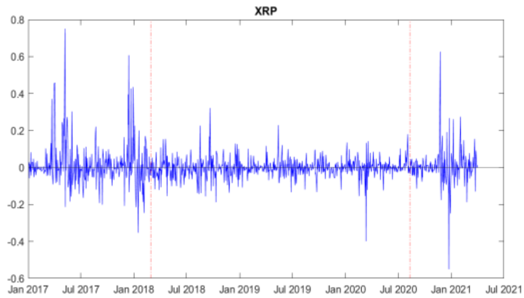

Litecoin:

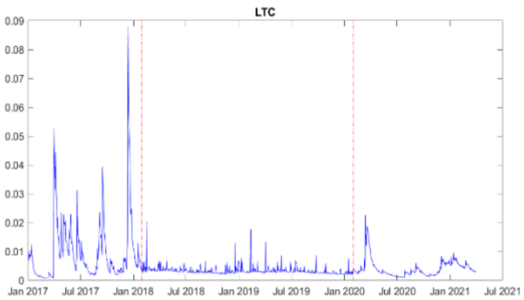

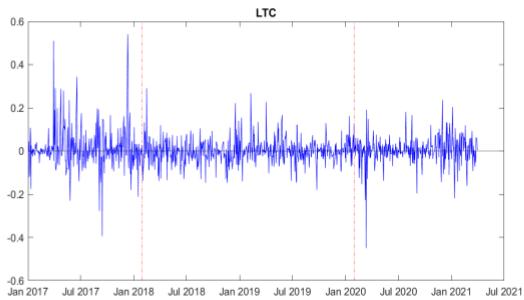

Binance Coin:

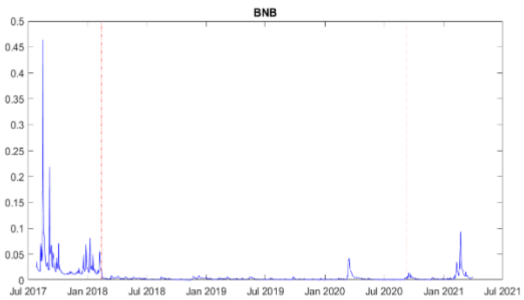

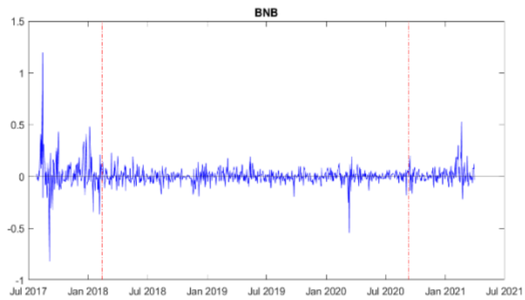

Cardano:

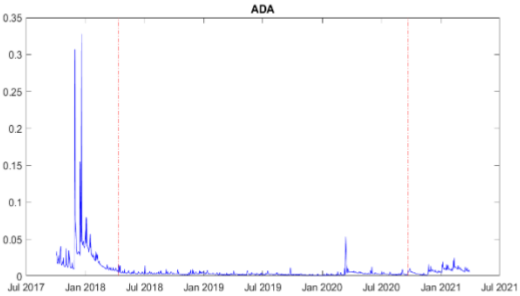

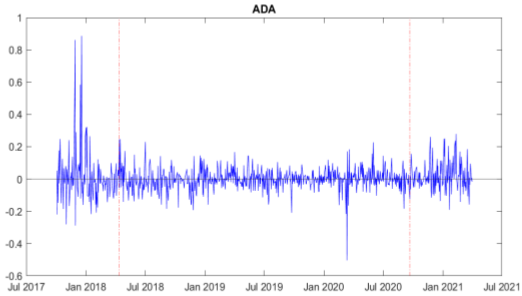

Euro:

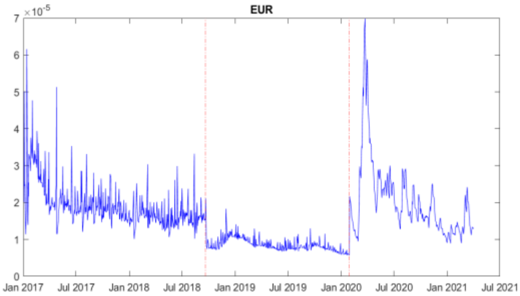

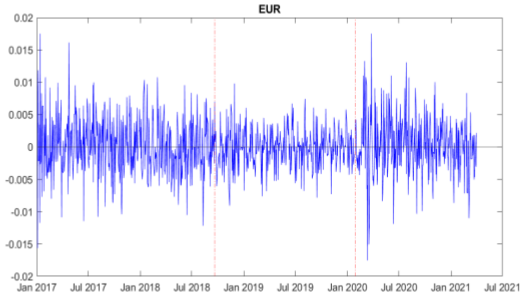

APPENDIX G: FIGARCH - Parameter Estimates

Table G.1 Bitcoin

	BTC	BTC_1	BTC_2	BTC_3
μ	0.0028***	0.0107***	0.0057	0.0014
	(0.0009)	(0.0027)	(0.0046)	(0.0009)
ϑ (AR1)	-0.8616***	0.4613	-0.4832*	-0.8454***
	(0.0952)	(1.2865)	(0.2898)	(0.1337)
ζ (MA1)	0.8393***	-0.4433	0.5624**	0.8046***
	(0.1020)	(1.2973)	(0.2685)	(0.1489)
ω	0.0000	0.0000	0.0000	0.0000
	NA	NA	NA	NA
ϕ (ARCH)	0.2193***	0.1412	0.3849**	0.0469
	(0.0699)	(0.1649)	(0.1954)	(0.0990)
β (GARCH)	0.7379***	0.6316***	0.8554***	0.9088***
	(0.1032)	(0.1450)	(0.1051)	(0.0348)
d-FIGARCH	0.6040***	0.6144***	0.5717**	0.9340***
	(0.1325)	(0.1897)	(0.2811)	(0.1133)

Table G.2 Ethereum

	ETH	ETH_1	ETH_2	ETH_3
μ	0.0017	0.0068**	-0.0021**	0.0068***
	(0.0013)	(0.0032)	(0.0010)	(0.0024)
ϑ (AR1)	-0.5520**	-0.5199*	0.9799***	-0.5224**
	(0.2696)	(0.2844)	(0.0080)	(0.2322)
ζ (MA1)	0.5050*	0.5891**	-0.9751***	0.4120*
	(0.2787)	(0.2606)	(0.0027)	(0.2441)
ω	0.0000	0.0000	0.0000	0.0000
	NA	NA	NA	NA
ϕ (ARCH)	0.2930*	0.5657***	0.2289	0.0000
	(0.1541)	(0.1715)	(0.1791)	(0.2861)
β (GARCH)	0.5108***	0.7672***	0.4985***	0.2677
	(0.1503)	(0.1205)	(0.1932)	(0.3243)
d-FIGARCH	0.3403***	0.4956***	0.3906***	0.2909***
	(0.0609)	(0.1752)	(0.0977)	(0.0932)

Table G.3 Ripple

	XRP	XRP_1	XRP_2	XRP_3
μ	-0.0021**	-0.0031	-0.0029**	0.0020
	(0.0010)	(0.0032)	(0.0012)	(0.0032)
ϑ (AR1)	0.0318	0.8008***	-0.4627	-0.9592***
	(0.2808)	(0.1797)	(0.3167)	(0.0193)
ζ (MA1)	-0.1127	-0.7710***	0.3429	1.0000***
	(0.2773)	(0.1932)	(0.3426)	(0.0082)
ω	0.0000	0.0000	0.0000	0.0000
	NA	NA	NA	NA
ϕ (ARCH)	0.7006***	0.4419**	0.6472***	0.6912***
	(0.0946)	(0.1839)	(0.1132)	(0.2102)
β (GARCH)	0.8243***	0.8038***	0.8480***	0.8283***
	(0.0588)	(0.0804)	(0.0590)	(0.1317)
d-FIGARCH	0.4425***	0.7648***	0.4606***	0.5191**
	(0.0882)	(0.1636)	(0.1282)	(0.2141)

Table G.4 Litecoin

	LTC	LTC_1	LTC_2	LTC_3
μ	-0.0005	0.0011	-0.0039***	0.0026
	(0.0012)	(0.0019)	(0.0008)	(0.0023)
ϑ (AR1)	-0.4028	0.5524	0.9678***	-0.2959
	(0.5405)	(0.3756)	(0.0066)	(0.2894)
ζ (MA1)	0.3634	-0.6274*	-0.9622***	0.1953
	(0.5512)	(0.3463)	(0.0039)	(0.2925)
ω	0.0000	0.0000	0.0000	0.0000
	NA	NA	NA	NA
ϕ (ARCH)	0.4724***	0.1469	0.1957***	0.2619
	(0.1062)	(0.1541)	(0.0755)	(0.2093)
β (GARCH)	0.7042***	0.6410***	0.9919***	0.7885***
	(0.0916)	(0.1512)	(0.0003)	(0.1998)
d-FIGARCH	0.4084***	0.6929***	1.0000***	0.6105*
	(0.0879)	(0.1381)	(0.0099)	(0.3613)

Table G.5 Binance Coin

	BNB	BNB_1	BNB_2	BNB_3
μ	0.0025*	0.0021	0.0015	0.0057
	(0.0014)	(0.0086)	(0.0016)	(0.0035)
ϑ (AR1)	0.6357**	-0.7568***	0.6676***	0.6202*
	(0.2721)	(0.2058)	(0.2450)	(0.3488)
ζ (MA1)	-0.6589**	0.8169***	-0.6952***	-0.6827**
	(0.2630)	(0.1670)	(0.2348)	(0.3162)
ω	0.0000	0.0000	0.0000	0.0000
	NA	NA	NA	NA
ϕ (ARCH)	0.0000	0.9654***	0.0259	0.8898***
	(0.6081)	(0.0292)	(0.3811)	(0.0961)
β (GARCH)	0.1056	0.9886***	0.2362	0.8421***
	(0.6059)	(0.0002)	(0.3990)	(0.1107)
d-FIGARCH	0.2910***	0.3226***	0.3174***	0.2232*
	(0.0633)	(0.0578)	(0.0668)	(0.1158)

Table G.6 Cardano

	ADA	ADA_1	ADA_2	ADA_3
μ	-0.0011	-0.0101	-0.0014	0.0033
	(0.0017)	(0.0071)	(0.0018)	(0.0057)
ϑ (AR1)	-0.5616***	0.7220**	-0.5615***	-0.6635***
	(0.1396)	(0.3191)	(0.1723)	(0.2113)
ζ (MA1)	0.4805***	-0.7451**	0.4811***	0.5397**
	(0.1486)	(0.2963)	(0.1833)	(0.2368)
ω	0.0000	0.0000	0.0000	0.0000
	NA	NA	NA	NA
ϕ (ARCH)	0.6779***	0.6455***	0.7137***	0.2072
	(0.0866)	(0.1543)	(0.0795)	(0.4571)
β (GARCH)	0.8260***	0.8929***	0.8676***	0.9130***
	(0.0567)	(0.0656)	(0.0382)	(0.1554)
d-FIGARCH	0.3451***	0.6054***	0.3525***	0.9093*
	(0.0763)	(0.2329)	(0.0676)	(0.5451)

Table G.7 Euro

	EUR	EUR_1	EUR_2	EUR_3
μ	0.0000	0.0002	-0.0002	0.0001
	(0.0001)	(0.0002)	(0.0002)	(0.0003)
ϑ (AR1)	-0.7388***	-0.6417***	0.3281	0.0212
	(0.1152)	(0.2407)	(0.2616)	(0.6070)
ζ (MA1)	0.8016***	0.7051***	-0.2062	0.0518
	(0.1011)	(0.2182)	(0.2686)	(0.6061)
ω	0.0000	0.0000	0.0000	0.0000
	NA	NA	NA	NA
ϕ (ARCH)	0.3829***	0.5191***	0.1297	0.1482
	(0.0715)	(0.0848)	(0.1289)	(0.1323)
β (GARCH)	0.7555***	0.8945***	0.9797***	0.5434***
	(0.0575)	(0.0247)	(0.0123)	(0.1479)
d-FIGARCH	0.4447***	0.5066***	0.9363***	0.3952***
	(0.0773)	(0.0898)	(0.0863)	(0.1143)

Note: Tables G.1 to G.7 report the parameter estimates with the ARMA(1,1)-FIGARCH(1,d,1) model on the seven cryptocurrencies and the Euro, respectively. The second column of each table reports the parameter estimates fitted on each of the asset for the entire sample period; the third, fourth and fifth columns report the estimates fitted on each of the sub-period as has been detected by the Bai-Perron test previously. The parameter estimates reported in the tables, from top to bottom, are for **(1) $\boldsymbol{\mu}$** (the constant term in the mean equation), **(2) $\boldsymbol{\vartheta}$** (the parameter for the autoregressive term in the mean equation), **(3) $\boldsymbol{\zeta}$** (the parameter for the moving average term in the mean equation), **(4) $\boldsymbol{\omega}$** (the constant term in the variance equation, set to zero in each model), **(5) $\boldsymbol{\phi}$** (the ARCH effect parameter in the variance equation), **(6) $\boldsymbol{\beta}$** (the GARCH effect parameter in the variance equation) and **(7) $\boldsymbol{d}$-FIGARCH** (the fractional differencing parameter of a FIGARCH model). The standard error for every corresponding parameter estimate is also reported in brackets below that very estimate. */**/*** after each estimate indicates statistical significance level at the 10/5/1 percent level.

APPENDIX H: 250 day - VaR and ES Series at a 97.5% Confidence Level

Bitcoin VaR:

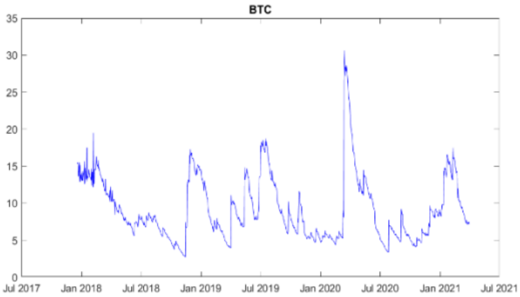

Bitcoin ES:

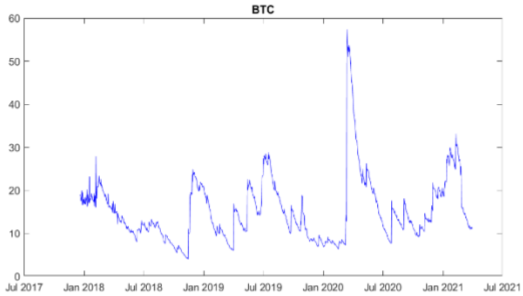

Ethereum VaR:

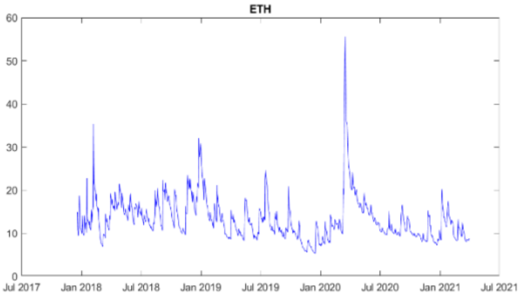

Ethereum ES:

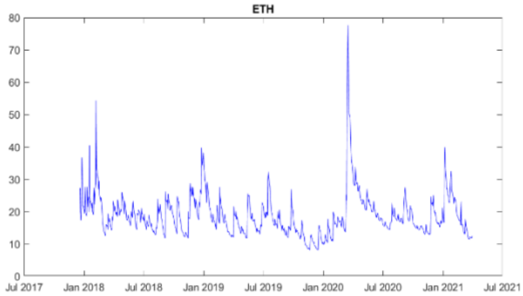

Ripple VaR:

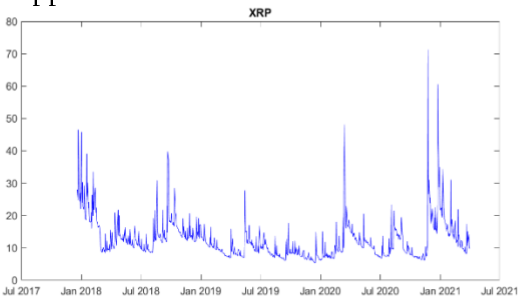

Ripple ES:

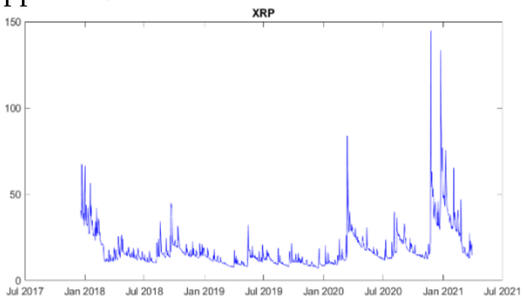

Litecoin VaR:

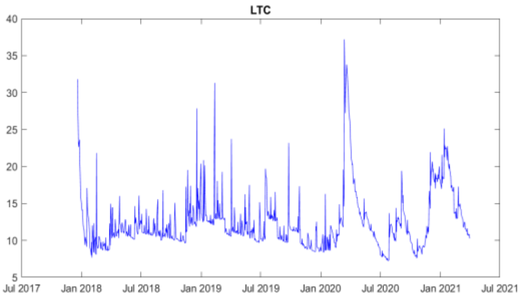

Litecoin ES:

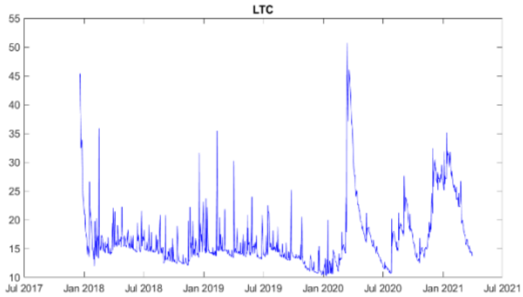

Binance Coin VaR:

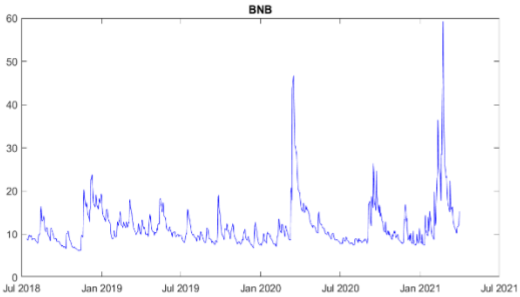

Binance Coin ES:

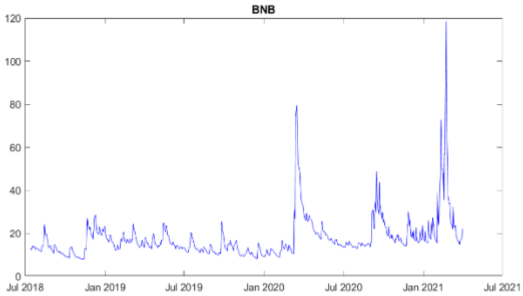

Cardano VaR:

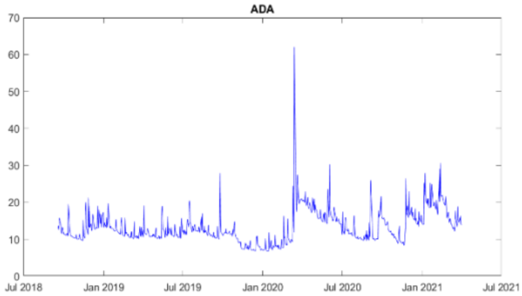

Cardano ES:

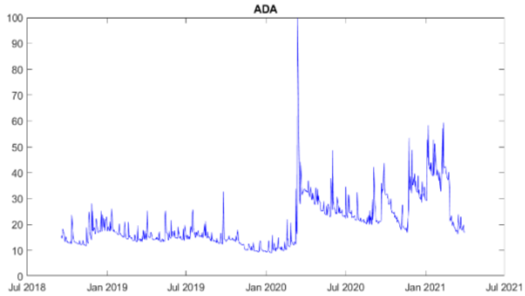

Euro VaR:

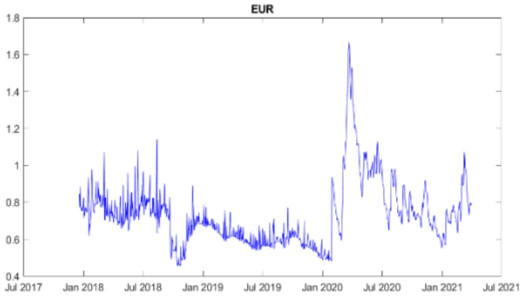

Euro ES:

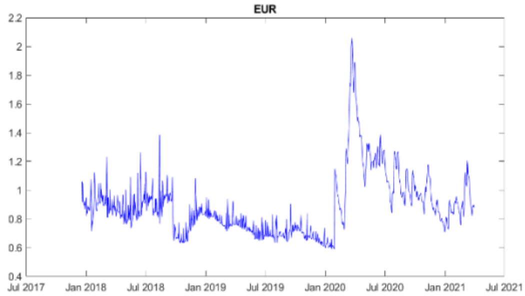

APPENDIX I: Supplementary Risk Measure Results

Table I.1 Risk Measures on the 2021-01-01 over a 250 day period

	BTC	ETH	XRP	LTC	BNB	ADA	EUR
VaR	8.49	9.06	21.62	17.75	8.44	14.12	0.58
ES	19.06	17.91	47.64	26.23	17.77	29.51	0.74

Table I.2 Risk Measures on the 2020-01-01 over a 250 day period

	BTC	ETH	XRP	LTC	BNB	ADA	EUR
VaR	4.92	7.36	6.21	8.75	7.68	7.09	0.49
ES	7.13	10.26	8.42	10.73	9.48	9.52	0.60

Table I.3 Risk Measures on the 2019-01-01 over a 250 day period

	BTC	ETH	XRP	LTC	BNB	ADA	EUR
VaR	14.30	27.98	14.47	14.60	18.80	13.58	0.67
ES	20.54	34.72	15.84	16.60	22.65	18.00	0.82

Table I.4 Risk Measures on the 2018-01-01 over a 250 day period

	BTC	ETH	XRP	LTC	BNB	ADA	EUR
VaR	13.12	10.33	45.80	14.36	NA	NA	0.82
ES	16.87	20.32	66.35	21.37	NA	NA	0.95

Note: Tables I.1 to I.4 report the reports the risk measures for all seven loss series on different dates throughout our sample. The displayed results are **(1) VaR** (the value at risk estimate for the date in the table heading) and **(2) ES** (the expected shortfall estimate for the date in the table heading)**.** For Table I.4 we do not obtain results for Binance Coin and Cardano since their respective start date is later than for the other cryptocurrencies.