Università
Ca'Foscari
Venezia

Master's Degree

In Economics and Finance

Final Thesis

Autoregressive Models for Cryptocurrencies Volatility Forecasting

Supervisor
Prof. Giovanni Angelini

Graduand
Tomaso Brunello
Matriculation number 859104

Academic Year
2019 / 2020

Index

Introduction

Given the increasing importance of cryptocurrencies for financial systems, the focus of this work is the analysis of their volatilities, which have proved to be large and difficult to predict, so far. A better understanding of risk related to cryptocurrencies' price movements would be useful for investment purposes.

As of 30 September 2014, Bitcoin has volatility seven times greater than gold, eight times greater than the S&P 500, and 18 times greater than the US dollar (Williams Mark T., World Bank Conference Washington DC, 11 November 2014), and this high variability persists over time.

Moreover, the large volatility is combined with clusters phenomena, which make forecasting attempts more difficult and investments in cryptocurrencies far riskier than investments in other financial assets.

Fluctuations of the value of these new cyber assets have been difficult to predict so far because they are not linked to any fundamental. In theory, cryptocurrencies' value is a representation of their usefulness as medium of exchange, which is still low, today. This leads to the hypothesis that their value is mainly influenced by sentiments of the market.

As showed in the next Chapters, the price of Bitcoins and many other cryptocurrencies has shown cyclicality patterns (also referred to as bubbles and busts) in recent years. (Moore, Heidi, "Confused about Bitcoin? It's 'the Harlem Shake of currency'", The Guardian, 1 March 2014).

Many studies tried to forecast cryptocurrencies volatility through different models and techniques. Catania et al. proposed a forecasting analysis of the four major cryptocurrencies (Bitcoin, Ethereum, Litecoin and Ripple) with univariate and multivariate models for volatility. (Catania, Grassi, Ravazzolo, "Forecasting Cryptocurrencies financial timeseries, Predicting the Volatility of Cryptocurrency Time–Series", 2018).

Their research was extended in subsequent Rossini and Bothe paper, in which the authors applied different multivariate Vector Autoregressive models to cryptocurrencies' returns and compared their forecasting powers. In their research, US stock market volatility as well as other measures of global economic activity were considered as possible exogenous drivers of cryptocurrencies' prices. (Bothe and Rossini, "Comparing the forecasting of cryptocurrencies by Bayesian time-varying volatility models", 2019).

Conrad et al. analysed cryptocurrencies' volatility through the lens of GARCH-MIDAS model to extract the long and short-term volatility components. (C. Conrad, A. Custovic, E. Ghysels, "Long- and Short-Term Cryptocurrency Volatility Components: A GARCH-MIDAS Analysis", J. Risk Financial Manag., 2018).

A similar research was conducted by Walther et al. who applied the GARCH-MIDAS framework to forecast the volatilities of five highly capitalized cryptocurrencies as well as the Cryptocurrency index CRIX, investigating the effect of Global Real Economic Activity as a major driver (T. Walther, T. Klein, E. Bouri "Exogenous drivers of Bitcoin and Cryptocurrency volatility – A mixed data sampling approach to forecasting", Elsevier B.V., 2019).

Kristjanpoller and Minutolo studied the behaviour of cryptocurrencies proposing a hybrid Artificial Neural Network-Generalized Autoregressive Conditional Heteroskedasticity (ANN-GARCH) model to forecast the price volatility of the most traded cryptocurrency: the Bitcoin (Werner Kristjanpoller, Marcel C. Minutolo, "A hybrid volatility forecasting framework integrating GARCH, artificial neural network, technical analysis and principal components analysis", Elsevier Ltd. 2018).

This thesis investigates the efficiency of both univariate models, in particular, different flavours of GARCH and Stochastic Volatility (SV) models and multivariate models, like VAR and VEC models combined with DCC GARCHs. The multivariate analysis, takes into account possible cointegrating relations among cryptocurrencies and includes the effect of exogenous variables on their returns.

The original Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model were created in 1986 by Tim Bollerslev, and today is one of the most used models applied in financial timeseries for volatility forecasting. In this thesis, Bollerslev's GARCH results are adopted as benchmark for evaluating richer GARCH family's models.

Several authors (see Sadorsky, 2005; Vo, 2009; Trolle and Schwartz, 2009; Larsson and Nossaman, 2011; Brooks and Prokopczuk, 2013) suggested that SV models are better in predicting the volatility of financial variables because they assume the variance of the stochastic process describing the volatility of prices to be itself randomly distributed, while in GARCH models the volatility follows an Autoregressive Moving Average process (ARMA), which is a deterministic function of model parameters and past data. This research applies Stochastic Volatility to univariate series of cryptocurrencies' log-returns, to evaluate its performances with respect to GARCH models.

For what concerns the multivariate analysis, a number of VARs and VECs models, including exogenous variables, are fitted to cryptocurrencies' log-returns. The predictors are added to the models to see whether they help understanding cryptocurrencies' volatility movements. The S&P 500, Nikkei 225 and Stoxx 600, are chosen as possible predictors since they are three of the largest stock indices in the world. Also, Gold and Silver are included in the models as they represent the most exchanged commodities. Finally, the US 3 months, US 10 years bonds interest rates and the VIX index are used as possible predictors.

Since their first introduction by C. Sims (Christopher Sims, "Macroeconomics and Reality", Econometrica, 1980), Vector autoregressions (VARs) have been widely adopted for forecasting. Their large success is due to their simplicity and efficiency, which make them a good benchmark in this thesis to compare performances of other multivariate models.

DCC GARCHs are finally fitted to VARs and VECs residuals to model cryptocurrencies' conditional volatilities and conditional correlations.

The best fitting models, selected with information criteria or Likelihood Ratio, are compared to other models in a pseudo-forecasting exercise, in order to determine whether they would have been good in predicting future cryptocurrencies' returns in the days subsequent to the "training" period of the models.

I Cryptocurrencies

Cryptocurrencies were born in the wake of the 2008 global financial crisis as a way for people to have a direct control over their savings, without having to rely on banks, governments or other financial intermediaries' management services.

They are digital assets designed to be internet-based mediums of exchange. The real innovation brought by cryptocurrencies is the use of cryptography to control the creation of new units of currency, to secure financial transactions and to prove the ownership of any amount of currency.

The resiliency of cryptography is guaranteed by a blockchain, which is a list of records called blocks, that cannot be modified retroactively once they are recorded. Blockchains are resistant to modifications of the data because they rely on a peer-to-peer network validating each new block. To put it simply, blockchain system could resemble an open ledger recording verifiable transactions between parties.

Blockchain method promises to make cryptocurrencies decentralised, unaffected by the influence of governments and central banks. While in centralized economies governments control the supply of currency by printing money, in cryptocurrencies' decentralized systems, companies or governments do not control it. In fact, most cryptocurrencies are designed to gradually decrease the production of units, placing a cap on the total amount of that currency in the market (Jon Matonis, "How Cryptocurrencies Could Upend Banks' Monetary Role", American Banker, 26 May 2013).

High transactions costs and high price volatility make cryptocurrencies not currently good for retail payments. Today, their main use is speculative, they represent opportunities for single investors or capital ventures, which hope to obtain a capital gain by holding directly units of currency or investing through specialized funds.

In 2019, cryptocurrencies' market capitalization was 237.1 billion U.S. dollars, a net increase from the 2018 value of 128.78 billion (M. Szmigiera, Cryptocurrency market capitalization 2013-2019, Statista, 20 January 2020). Although the market is composed of almost 3.000 different cryptocurrencies, the ten most capitalized cryptocurrencies make the 91% of the total. In this thesis only four cryptocurrencies are analysed: Bitcoin, Ethereum, Litecoin and Ripple, which together, represent approximately 76% of the market capitalization.

All operations involving cryptocurrencies are executed in cryptocurrency exchanges, privately held firms allowing customers to trade cryptocurrencies for conventional money or other digital currencies. A cryptocurrency exchange typically charges fees on transactions or earns the dealer's bid-ask spreads for its services.

Being private and not subject to same rules as other larger exchanges, cryptocurrency exchanges created concerns about security and trading fees regulation (M. Szmigiera, "Traders: biggest problems of cryptoexchanges 2018 | Statistic", Statista, 3 September 2018. "Expectations of traders from cryptocurrency exchanges 2018 | Statistic", Statista, 3 September 2018).

Despite these issues, these particular exchanges are the major sources of information on cryptocurrencies. Data used in this thesis are freely available on “cryptodatadownload.com”, a web service which gathers price time series from all major cryptocurrency exchanges. In particular, time series of daily prices were provided by Kraken and Bittrex, two of the largest crypto-exchanges in the US and in the world. This work analyses the value of cryptocurrencies in terms of US$, the most exchanged currency in the world.

1. Bitcoin

Created in 2008 by an unknown group of developers using the name Satoshi Nakamoto, Bitcoin is the older cryptocurrency in the market. Today it is still the most exchanged in the world, despite many others have been created in subsequent years.

On October 2019, Bitcoin represents 62.12% of total cryptocurrencies market with its $147.3bn capitalization (Rick Bagshaw, “Top 10 cryptocurrencies by market capitalisation”, Coin Rivet, 8 October 2019).

2. Ethereum

Often the name Ethereum is erroneously attributed to the currency, while it is instead the name of the platform generating cryptocurrency as a reward to mining nodes (by performing computations). The name of the currency generated is Ether and it is the only one accepted in the payment of transaction fees (Usman W. Chohan, “Cryptocurrencies: A Brief Thematic Review”, Social Science Research Network, 28 August 2017).

In 2013 Vitalik Buterin, a cryptocurrency programmer, proposed Ethereum for the first time as an alternative to the more famous Bitcoin. Ethereum platform development was financed through a crowdfunding campaign which allowed the system to go online on 30 July 2015. At that time, 72 million Ether were available in the market, ("Ether Supply Growth Chart", etherscan.io, 9 January 2020).

What makes Ethereum attractive with respect to the Bitcoin is the time required to make a transaction (to generate each block in the blockchain). While this operation takes about 14 to 15 seconds in the case of Ethereum, 10 minutes are required by Bitcoin. This makes Ether much more useful in transactions with merchants than other cryptocurrencies.

Ethereum (ETH) is the second cryptocurrency by market capitalization with $19.4bn circulating on October 2019 (R. Bagshaw, Coin Rivet, 8 October 2019).

3. Litecoin

On October 7, 2011 Charlie Lee, a Google employee, released Litecoin via an open-source client (Robert McMillan, "Ex-Googler Gives the World a Better Bitcoin", Wired).

Litecoin was created as an alternative to Bitcoin in the early days of cryptocurrencies’ markets. Despite the two cryptocurrencies being almost identical, the former system is way faster in validating blocks. Litecoin needs only 2.5 minutes per operation, rather than Bitcoin's 10 minutes. While Litecoin is faster than Bitcoin it is still much less efficient than Ethereum and Ripple in confirming transactions. (Ian Steadman, "Wary of Bitcoin? A guide to some other cryptocurrencies", Ars Technica, 11 May 2013).

Litecoin Market Capitalization on October 2019 was $3.6bn (Rick Bagshaw, "Top 10 cryptocurrencies by market capitalisation", Coin Rivet, 8 October 2019).

4. Ripple

Ripple financial system was proposed by Jed McCaleb and subsequently developed by Ryan Fugger, as a mean of payment for participants to the system. Mc Caleb, Fugger and other partners created also their cryptocurrency, called XRP, which works similarly to Bitcoin. They aimed at creating a currency allowing fast and cheap transactions between financial institutions.

Since 2013, Ripple system has been increasingly adopted by many financial institutions in order to provide an alternative to usual remittance methods (Matt Scully, "Alternative Money Mover Ripple Labs Enters U.S. Banking System", American Banker, 24 September 2014). By 2018, more than 100 banks started using Ripple's technology, but avoided using the XRP cryptocurrency, because of its price instability. Slowly, also the XRP were adopted by financial institutions and in late 2019, XRP was the third cryptocurrency in terms of market capitalization, as $11.7bn were circulating in the market.

Like Bitcoin and other cryptocurrencies, Ripple relies on a blockchain system connected through servers in which common shared database record all transactions. Although Ripple validation operations take only few seconds, much less than Bitcoin, (Tom Simonite, "Big-name investors back effort to build a better Bitcoin", MIT Technology Review, 11 April 2013) no possibilities of chargebacks make Ripple still not appropriate for purchases and sales.

II Bayesian Models Comparison for Univariate Models

In analysing univariate cryptocurrencies series, both GARCH family models and Stochastic Volatility models were considered in the attempt of capturing the behaviour of data volatilities. While in the GARCH class models, conditional volatility σ_t is known if the parameters in its equation are specified, in Stochastic Volatility models σ_t is itself stochastic, with an innovation term that is not known at time t.

Univariate models' estimates are obtained through Bayesian estimation techniques, and model comparison is carried out by evaluating Marginal Likelihoods of each model, obtained by fitting the models to historical data time series of each cryptocurrency.

Marginal Likelihood is a good measure of fit when comparing non-nested models like GARCHs and SVs, because it considers the number of parameters to be estimated, penalizing model complexity. (Koop, 2003, for a detailed discussion on Bayesian model comparison).

Bayesian model comparison is implemented evaluating Bayes factor, computed using importance sampling[1]. Let {M1, . . . ,MK} to represent the set of models to evaluate. Each model Mk has two

[1] Importance sampling is a Bayesian estimation technique which estimates a parameter by drawing from a specified importance function rather than a posterior distribution. Importance sampling is used in Bayesian

separate components: a Likelihood function $p(y|\theta_k, M_k)$ that depends on the model-specific parameter vector θ_k, and a prior density $p(\theta_k|M_k)$.

Bayes factor is defined as

$$BF_{ij} \ = \frac{p(y\,|Mi)}{p(y\,|Mj)} \tag{1}$$

Where

$$p(y|M_k) = \int p(y|\theta_k, M_k) p(\theta_k|M_k) d\theta_k \tag{2}$$

$$\text{With} \quad k = i, j$$

is the Marginal Likelihood function for both models.

Marginal Likelihood is a density forecast obtained observing actual data y under model M_k. The Marginal Likelihood of a model would be "large" if the observed data are likely under that model. Hence, the ratio defines whether model M_i or model M_j fits better the data.

When BF_{ij} > 1 observed data are more likely under model M_i, when BF_{ij} < 1, the opposite holds true and observed data are more likely under model M_j. The log Bayes factor is asymptotically equivalent to the Schwarz information criterion (SIC). The SIC for model M_k is defined as

$$SIC_k = log\, p\big(y\big|\hat{\theta}_k, M_k\big) - \frac{pk}{2}\, logT \tag{3}$$

where $\hat{\theta}_k$ is the maximum likelihood estimate and T is the sample size. Both the Bayes factor and SIC are consistent model selection criteria. Moreover, for a finite sample, the Bayes factor is related to the posterior odds ratio between the two models as follows

$$\frac{\mathrm{P}(M_i\,|\,\mathrm{y})}{\mathrm{P}(M_j\,|\,\mathrm{y})} = \frac{\mathrm{P}(M_i)}{\mathrm{P}(M_j)} \times BF_{ij} \tag{4}$$

where $\mathrm{P}(M_i)/\mathrm{P}(M_j)$ is the prior odds ratio[2]. If the models are equally probable a priori, the prior odds ratio is one and the posterior odds ratio between the two models is equal to the Bayes factor. In this work only the Marginal Likelihoods are reported, as the Bayes factor is just a ratio between two Marginal Likelihoods.

Methods to calculate the Marginal Likelihood are complex because the integral in (2) cannot be obtained analytically, as it is high-dimensional. One method used to address this, were developed

inference and rare event simulation in finance or insurance. (Aptech Systems Inc, "Bayesian Importance Sampling", Aptech.com)

[2] The prior probability distribution of a parameter express researcher's beliefs about that parameter before some evidence is obtained. It is in fact the unconditional probability assigned before any relevant evidence is considered.

by Chan and Eisenstat, who implemented an improved version of the cross-entropy method, (Chan and Eisenstat, “Marginal Likelihood Estimation with the Cross-Entropy Method”, September 2012), an adaptive importance sampling method to compute marginal likelihood developed by Rubinstein (1997), Rubinstein and Kroese (2004).

Advantages of this method are its easy implementation and the fact that it only requires the evaluation of the priors and the Likelihood.

All the GARCH and stochastic volatility models are estimated using Markov chain Monte Carlo (MCMC) methods. Specifically, the Metropolis-Hastings algorithms are used to sample from the posterior distributions of the models. The posterior draws obtained are used to compute the posterior means and the Marginal Likelihoods. Technical details of estimation technique are available in Chan and Grant paper “Modeling Energy Price Dynamics: GARCH versus Stochastic Volatility”, November 2015.

III Univariate Models

Firstly, seven common GARCH models are fitted to data: GARCH (1,1), GARCH (1,2), GARCH-j with jumps, GARCH-M or GARCH in mean, GARCH-MA with moving average (MA) innovations, GARCH-t with t distributed innovations, GARCH-GJR with leverage effects. In a second step the Stochastic Volatility models (SV), counterparts of GARCH considered before, are fitted to the same data: SV, SV (2), SV-j with jumps, SV-M or SV in mean, SV-MA, SV-t and SV-L with leverage effects.

The last section of this Chapter is dedicated to a discussion of residual analysis techniques adopted. They are important because they can be used to define whether the conditional volatility models are correctly specified and whether the selected models fit properly to data.

1. GARCH Models

In this section are illustrated the GARCH models applied to cryptocurrencies series. The first model considered is the GARCH (1,1), which can be represented as

$$y_t = \mu + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_t^2),$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \qquad (5)$$

where $\varepsilon_0 = 0$ and σ_t^2 are constant. The conditional volatility follows an AR (1) process and two constraints are set to ensure model stationarity and positiveness of σ_t^2:

$$\alpha_1 + \beta_1 < 1 \text{ for stationarity,}$$

$$\alpha_0 > 0,\ \alpha_1 > 0,\ \beta_1 > 0 \text{ for positive conditional variance.}$$

In the second model, σ_t^2 follows an AR (2) process. The GARCH (1,2) can be defined as

$$y_t = \mu + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_t^2),$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 + \beta_2 \sigma_{t-2}^2 \tag{6}$$

The constraints assumed for the GARCH (1,1) are also applied in this second model,

$$\alpha_1 + \beta_1 + \beta_2 < 1 \quad \text{to ensure stationarity,}$$

$$\alpha_0 > 0, \ \alpha_1 > 0, \ \beta_1 > 0, \ \beta_2 > 0 \quad \text{for positive conditional variance.}$$

The third model is a GARCH (1,1) allowing for large movements (or jumps) in the series of data.

$$y_t = \mu + k_t q_t + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_t^2),$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 (y_{t-1} - \mu)^2 + \beta_1 \sigma_{t-1}^2 \tag{7}$$

q_t is a dichotomous variable taking the value 1 with probability $\mathrm{P}(q_t = 1) = \kappa$, when a jump occurs, and 0 when it does not. k_t defines the size of the jump and it is distributed as $k_t \sim \mathrm{N}(\mu_k, \sigma_k^2)$.

The fourth model considered is the famous GARCH in mean (GARCH-M) introduced by Engle et al. in 1987. This model investigates whether the mean of the returns depends on some function of the conditional variance. In this specific case the model considered is

$$y_t = \mu + \lambda \sigma_t^2 + \varepsilon_t, \qquad \varepsilon_t \sim \mathrm{N}(0, \sigma_t^2)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 (y_{t-1} - \mu - \lambda \sigma_{t-1}^2)^2 + \beta_1 \sigma_{t-1}^2 \tag{8}$$

Where λ defines the impact of the conditional variance on the return at time t.

The fifth model considered is a GARCH-MA in which the innovation is supposed to follow a moving average (MA) process which links it to the conditional variance in time t and in time t-1. This model allows the data series to be correlated over time.

$$y_t = \mu + \varepsilon_t,$$

$$\varepsilon_t = u_t + \psi u_{t-1}, \tag{9}$$

$$u_t \sim \mathrm{N}(0, \sigma_t^2)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2$$

Invertibility of the MA process is ensured by imposing $|\psi|<1$, and the same conditions of stability and positiveness described before, hold for the GARCH (1,1) components of the model.

GARCH-t model considered next is simply a GARCH (1,1) in which the innovation is a Student-t distributed random variable. Its conditional variance is still defined by a GARCH (1,1) model. Student-t distribution has heavier tails than the normal distribution and makes more likely to observe extreme values of the innovation. The usual restrictions are imposed to the GARCH parts of the model.

$$y_t = \mu + \varepsilon_t,$$

$$\varepsilon_t \sim \mathrm{tv}(0, \sigma_t^2) \tag{10}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2$$

The last GARCH family model considered is a GARCH-GJR also known as threshold-GARCH. First introduced by Glosten, Jagannathan, and Runkle in 1993, the model allows large impact of negative returns on the conditional variance. It is defined as

$$y_t = \mu + \varepsilon_t,$$

$$\varepsilon_t \sim \mathrm{N}(0, \sigma_t^2)$$

$$\sigma_t^2 = \alpha_0 + (\alpha_1 + \delta_1 \mathbb{1}(\varepsilon_{t-1} < 0))\varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \qquad (11)$$

where $\mathbb{1}(\cdot)$ is the indicator function, when the condition between parentheses holds true, the asymmetric leverage effect controlled by the parameter δ_1 is considered by the model. If $\delta_1 = 0$, this variant becomes the standard GARCH.

2. Stochastic Volatility Models

Stochastic Volatilities models considered are close counterparts of GARCH models in the previous section. While in GARCH models the conditional variance is defined by a deterministic function, under stochastic volatility models, it is a random variable itself. The first model considered is the Stochastic Volatility (SV) model, which is defined as

$$y_t = \mu + \varepsilon_t^y, \qquad \varepsilon_t^y \sim N(0, e^{h_t}),$$

$$h_t = \mu_h + \varphi_h(h_{t-1} - \mu_h) + \varepsilon_t^h, \qquad \varepsilon_t^h \sim N(0, \omega_h^2) \qquad (12)$$

h_t defines the log-volatility of y_t and follows a stationary AR (1) process with $|\varphi_h| < 1$ and mean μ_h. The process starts with $h_1 \sim \mathrm{N}(\mu_h, \omega_h^2/(1 - \varphi_h^2))$. In $(\omega_h^2)/(1 - \varphi_h^2)$ can be recognized the formulation of the unconditional variance of an AR (1) process. In the second model considered, y_t is defined like in the previous case but the log volatility follows a stationary AR (2) process. The model can be written as

$$y_t = \mu + \varepsilon_t^y, \qquad \varepsilon_t^y \sim N(0, e^{h_t}),$$

$$h_t = \mu_h + \varphi_h(h_{t-1} - \mu_h) + \rho_h(h_{t-2} - \mu_h) + \varepsilon_t^h, \qquad \varepsilon_t^h \sim N(0, \omega_h^2) \qquad (13)$$

SV (2) model starts with h_1 and h_2 which are assumed to follow

$$h_1,\ h_2 \sim \mathrm{N}\left(\mu_h, \frac{(1-\rho\mathrm{h})\omega_h^2}{(1+\rho_h)((1-\rho_h)^2 - \varphi_h^2)}\right),$$

where $\frac{(1-\rho\mathrm{h})\omega_h^2}{(1+\rho_h)((1-\rho_h)^2 - \varphi_h^2)}$ is the unconditional variance of the AR (2) process.

The third Stochastic Volatility model is the SV-j model, counterpart of the GARCH-j model considered before. It is defined as

$$y_t = \mu + k_t q_t + \varepsilon_t^y, \qquad \varepsilon_t^y \sim N(0, e^{h_t}), \qquad (14)$$

$$h_t = \mu_h + \varphi_h(h_{t-1} - \mu_h) + \varepsilon_t^h \qquad \varepsilon_t^h \sim N(0, \omega_h^2).$$

q_t is defined similarly to the GARCH-j case, it takes the value 1 with probability ($\mathrm{P}(q_t = 1) = \kappa$, when a jump occurs and 0 when it does not. Exactly as before, k_t defines the size of the jump and it is normally distributed as $k_t \sim \mathrm{N}(\mu_k, \sigma_k^2)$.

Stochastic Volatility in mean (SV-M) model, developed by Koopman and Hol Uspensky in 2002, is considered next. Like in its GARCH counterpart, the observation at time t is affected by stochastic volatility at time t as follows

$$y_t = \mu + \lambda e^{h_t} + \varepsilon_t^y \qquad \varepsilon_t^y \sim N(0, e^{h_t}), \tag{15}$$

Where the parameter λ defines how much the volatility affects the log-return at time t. h_t is defined as in previous models as

$$h_t = \mu_h + \varphi_h(h_{t-1} - \mu_h) + \varepsilon_t^h \qquad \varepsilon_t^h \sim N(0, \omega_h^2).$$

The fifth model is SV-MA in which the returns' innovation follows a first order Moving Average process,

$$y_t = \mu + \varepsilon_t^y$$
$$\varepsilon_t^y = u_t + \psi u_{t-1} \tag{16}$$
$$u_t \sim N(0, e^{h_t}),$$
$$h_t = \mu_h + \varphi_h(h_{t-1} - \mu_h) + \varepsilon_t^h \qquad \varepsilon_t^h \sim N(0, \omega_h^2).$$

Where u_0 = 0. Moreover, in order to ensure the invertibility of the MA process | ψ | < 1 is assumed. The usual stationary AR (1) process is assumed for the log volatility, as before.

SV-t model is considered next, as Stochastic Volatility counterpart of the GARCH-t model. As in the GARCH case, the innovation follows a Student-t distribution, which increases the probability of large values of y_t.

$$y_t = \mu + \varepsilon_t^y$$
$$\varepsilon_t^y \sim tv(0, e^{h_t}) \tag{17}$$
$$h_t = \mu_h + \varphi_h(h_{t-1} - \mu_h) + \varepsilon_t^h, \qquad \varepsilon_t^h \sim N(0, \omega_h^2)$$

In the SV-t model, the log volatility follows again the stationary AR (1) process defined in (12). Finally, a proper counterpart of GARCH-GJR model, is the SV-L model, that allows for leverage effect. As in the GARCH case, negative returns produce larger values of the conditional volatility

$$y_t = \mu + \varepsilon_t^y$$
$$h_{t+1} = \mu_h + \varphi_h(h_t - \mu_h) + \varepsilon_t^h$$

In this case the innovations ε_t^y and ε_t^h follow the bivariate normal distribution

$$\begin{pmatrix} \varepsilon_t^y \\ \varepsilon_t^h \end{pmatrix} \sim \text{ MVN } (0, \begin{pmatrix} e^{ht} & \rho e^{\frac{1}{2}ht}\omega_h \\ \rho e^{\frac{1}{2}ht}\omega_h & \omega_h^2 \end{pmatrix}) \tag{18}$$

If $y_t < 0$ and $\rho < 0$, then the volatility at time t+1 tends to be larger than that at time t. Models adopted to carry out the analysis can be summarized as follow

GARCH model	**Model description**
GARCH (1,1)	The conditional volatility follows a stationary AR (1) process
GARCH (1,2)	The conditional volatility follows a stationary AR (2) process
GARCH-j	GARCH (1,1) allowing for large movements (or jumps) in the returns
GARCH-m	GARCH (1,1) in which returns are directly affected by the volatility at time t
GARCH-MA	GARCH (1,1) in which the innovation follows an invertible MA (1) process
GARCH-t	GARCH (1,1) in which the innovation is distributed as a Student-t
GARCH-GJR	GARCH (1,1) allowing for leverage effects
Stochastic Volatility model	**Model description**
SV	The log-volatility follows a stationary AR (1) process
SV (2)	The log-volatility follows a stationary AR (2) process
SV-j	SV allowing for large movements (or jumps) in the returns
SV-m	SV with returns directly affected by the volatility observed at time t
SV-MA	SV in which the innovation follows an invertible MA (1) process
SV-t	SV in which the innovation is distributed as a Student-t
SV-l	SV allowing for leverage effects

Table 1: Univariate models summary

3. Residuals Analysis

The presence of autocorrelation in the standardized residuals of a model, biases the estimators and makes them less efficient, suggesting that some information have not been accounted for in the fitted model (Eliezer Bose, Marilyn Hravnak and Susan M. Sereika, “Vector Autoregressive (VAR)

Models and Granger Causality in Time Series Analysis in Nursing Research: Dynamic Changes Among Vital Signs Prior to Cardiorespiratory Instability Events as an Example", 2017 Wolters Kluwer Health, Inc.).

Thus, it is important to determine the presence of autocorrelation in the standardized residuals of any fitted univariate and multivariate model. Several tests exist in order to assess whether residuals are autocorrelated, the most famous are the Ljung-Box and the Breusch-Godfrey tests. Today, a discussion is still open among several authors in the field of Econometrics to determine which of the two tests is more appropriate in the analysis of residuals from autoregressive models. While many are against the validity of the Ljung-Box Q-statistic in an autoregressive model residuals analysis (Maddala, "Introduction to Econometrics (3d edition), ch 6.7, and 13. 5 p 528, 2001 and Hayashi, ch. 2.10 "Testing For serial correlation", 2000), it is still the most used in these cases and it is also the one used in this thesis. The test is also useful when applied to standardized squared residuals. In this case it suggests whether conditional heteroskedasticity is left in the standardized residual series.

Ljung–Box test is a portmanteau test and it is run to check whether any of a group of autocorrelations of a time series are different from zero. Under the null hypothesis, the data are independently distributed (there is no autocorrelations in the series). The alternative hypothesis is that data are not independently distributed since they exhibit serial correlation.

The test statistic is: $\mathrm{Q} = \mathrm{n}(\mathrm{n}+2) \sum_{k=1}^{h} \frac{\widehat{\rho_k^2}}{n-k}$, where

- n = sample size;
- $\widehat{\rho_k}$ = sample autocorrelation at lag k;
- h = number of lags being tested

(G. Ljung; G. E. Box, "On a Measure of a Lack of Fit in Time Series Models", Biometrika, 1978).

Under the null hypothesis, Q asymptotically follows a $\chi^2_{(h)}$. The hypothesis (H_0) is rejected at significance level α when: Q > $\chi^2_{1-\alpha,h}$,where $\chi^2_{1-\alpha,h}$ is the 1-α quantile of a chi-square distribution with *h* degrees of freedom. (Wikipedia.org, "Augmented Dickey–Fuller test", 27 February 2020). However, when the portmanteau test is applied to estimated residuals, asymptotic distribution of the test statistics is unknown. In such a case, portmanteau tests should be interpreted with care even if simulation results reported by Tse and Tsui (1999) suggest that they provide a useful diagnostic in many situations. (S. Laurent, L. Bauwens, J. Rombouts, "Multivariate GARCH models: a survey", Journal of Applied Econometrics, 2006).

Assessing the absence of autocorrelation in the residuals estimated from an autoregressive model, though, would not be a complete analysis. One should investigate whether the assumptions made for the innovations in the models are correct or not. There exist many tests to check the distribution of residuals. The most famous are the Jarque-Bera, or the Shapiro-Wilk test, which are used to determine if a model's residuals are normally distributed. Though, they are not used in this thesis, because they do not support other distributions, which are often preferred in financial series analysis.

A more complete test is the Kolmogorov-Smirnov. It is a non-parametric test which compares the shape of two distributions. It may be used to compare both a sample distribution with a theoretical one, and two different sample distributions.

In this thesis analysis, the test statistic is computed as the distance between the empirical Cumulative Distribution Function (CDF) and the theoretical one. The empirical CDF F_n for n independent and identically distributed (i.i.d.) ordered observations X_i is defined as

$$F_n(x) = \frac{1}{n}\sum_{i=1}^{n} I_{[-\infty,x]}(X_i) \quad (19)$$

where $I_{[-\infty,x]}$ is the indicator function, equal to 1 if $X_i \leq x$ and equal to 0 otherwise. The null hypothesis for this test is

$$H_0: F(x) = F_0(x), \forall x \quad (20)$$

So, under the null hypothesis the empirical distribution is equal to the theoretical one. The Kolmogorov–Smirnov statistic for a given cumulative distribution function $F(x)$ is

$$D_n = sup_x|F_n(x) - F(x)| \quad (21)$$

The idea behind the Kolmogorov-Smirnov test is quite intuitive because the test statistic simply represent the distance between the empirical CDF and the theoretical CDF. When D_n is large the null hypothesis that the two distributions are the same is rejected. It is not rejected instead, for small values of D_n (Wikipedia.org, “Test di Kolmogorov-Smirnov”, 24 February 2018).

4. Forecasting

The GARCHs and SVs models main use in finance is forecasting. Once the best models are selected through the marginal Likelihood results, it is interesting to check their actual forecasting performances against the performances of the discarded models.

It would be easy to forecast conditional volatilities and returns at time t+1 with GARCHs and SV models, since it is enough to build the models with data available at time t. In order to estimate data at time t+2, the same models are built on previously estimated data for time t+1. In the case of a simple GARCH (1,1) the model can be rewritten as

$$\hat{y}_{t+1} = \hat{\mu} + \hat{\varepsilon}_{t+1}, \quad \hat{\varepsilon}_{t+1} \sim N(0, \hat{\sigma}^2_{t+1})$$

$$\hat{\sigma}^2_{t+1} = \hat{\alpha}_0 + \hat{\alpha}_1\hat{\varepsilon}^2_{t-1} + \hat{\beta}_1\hat{\sigma}^2_t$$

Where the condition $\hat{\alpha}_1 + \hat{\beta}_1 < 1$ must hold to ensure stationarity, and $\hat{\alpha}_0 > 0$, $\hat{\alpha}_1 > 0$, $\hat{\beta}_1 > 0$ for granting positive conditional variance. The forecast estimations of Stochastic Volatility models could be obtained in a similar way. For example, when the aim is forecasting one period ahead, the SV model can be rewritten as

$$\hat{y}_{t+1} = \hat{\mu} + \hat{\varepsilon}^y_{t+1}, \quad \hat{\varepsilon}^y_{t+1} \sim N(0, e^{\hat{h}_{t+1}}),$$

$$\hat{h}_{t+1} = \hat{\mu}_h + \hat{\varphi}_h(\hat{h}_t - \hat{\mu}_h) + \hat{\varepsilon}^h_{t+1}, \quad \hat{\varepsilon}^h_{t+1} \sim N(0, \hat{\omega}^2_h)$$

Where $\hat{h}_t$ defines the log-volatility of $\hat{y}_t$ and follows a stationary AR (1) process with $|\hat{\varphi}_h| < 1$ and mean $\hat{\mu}_h$.

The procedure of comparing estimated forecasts with actual observations in the market, which were not used to train a model, is called out of sample forecasting. In other words, in out of sample forecasting, observations are not a subsample of historical data used for estimations. The forecasted sample is usually generated through a rolling forecast procedure. In practice, the estimated model is used to generate a forecast of the next daily log-return. The model is then re-estimated on the entire available sample (historical data plus the first observation out of the sample). A forecast for the second day is generated and the model estimated again. The algorithm proceeds in this way until a forecasted period of the desired length is produced.

While this is often the approach used to evaluate forecasting performances, its results are not reported here. The measures used to evaluate models in out of sample forecasting are usually based on the Mean Squared Error, which, is the mean of the squared differences of forecasted observations and actual observations. The MSE computed using the approach described above were not reliable because of the scarce number of observations available in the markets at the time of the analysis (Roughly 100 actual daily prices were available to make the comparison).

The approach chosen in order to determine whether the models are good for forecasting, is comparing the predictive marginal Likelihood for each model. This method is completely different from the MSE measures of performance, as it does not need that forecasts are produced. In fact, the predictive marginal Likelihood can be obtained exactly like the marginal Likelihood computed on historical data. The difference is that it is the marginal Likelihood obtained filtering the newly available data with a model computed on the full sample available (historical data + out of sample data). The result obtained can be interpreted as the likelihood of observing the actual log-returns under the last estimated model. The larger the marginal Likelihood obtained, the better the model in explaining the returns observed in the market.

Finally, the forecasting performances obtained with predictive marginal Likelihood can be compared to fitting performances obtained on historical data. In this way, one can observe whether the models selected to fit the data would have been also the best for forecasting.

IV Univariate Models: Implementation and results

The first step in this analysis is the gathering of the daily historical exchange rate of Bitcoin, Ethereum, Litecoin and Ripple. Data used in this dissertation are freely available on cryptodatadownload.com, which gathers historical data from many cryptocurrencies' exchanges across the world. In particular, the data used in this work are provided by Kraken and Bittrex exchanges.

The period considered starts on 9th August 2015 and terminates on 7th February 2020 for a total of 1632 daily observations. The value of cryptocurrencies is compared against the $US which is the more stable currency in the World. Moreover, more data are available for this exchange rate compared to others. Data transformation is necessary to obtain stationary time series and meaningful GARCH and SV models estimations results.

In fact, before starting the analysis with conditional volatility models, it is important to investigate the stationarity of the series. This can be done by implementing the augmented Dickey-Fuller test[3] on each cryptocurrency's exchange rate. All the three versions of the test are run and the P-values obtained for cryptocurrencies' series are much larger than the 0.05 critical value, except for Ripple, in which the test p-value resulted slightly above the critical value in the version of the test with no drift and no trend.

Bitcoin	p-value	Test result
AR	0.5429	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift	0.5716	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift and trend	0.4567	p-value > 0.05 critical value, do not reject the null hypothesis
Ethereum		
AR	0.2332	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift	0.4138	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift and trend	0.7379	p-value > 0.05 critical value, do not reject the null hypothesis
Litecoin		
AR	0.1776	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift	0.2792	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift and trend	0.4756	p-value > 0.05 critical value, do not reject the null hypothesis
Ripple		
AR	0.0634	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift	0.1202	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift and trend	0.2749	p-value > 0.05 critical value, do not reject the null hypothesis

Table 2: Augmented Dickey-Fuller test results on exchange rate series.

These results suggest not rejecting the null hypothesis that a unit root is detected in each of the time series, meaning that all the series are not stationary. One could expect such result observing the exchange rate time series plot obtained with Matlab software,

---

[3] Given the AR(1) process $y_t = \rho y_{t-1} + u_t$. A unit root is present if ρ = 1, and in this case the model would be non-stationary. The model can be rewritten as $\Delta y_t = \alpha + \beta_t + \Upsilon y_{t-1} + \delta \Delta y_{t-1} + \cdots + \delta_{p-1} \Delta y_{t-p+1} + \varepsilon_t$, where Δ is the first difference operator, β is the coefficient of the time trend and δ = (ρ-1).

Augmented Dickey Fuller tests the null hypothesis that a unit root is present by checking whether δ = 0. Since the test is done over the residuals rather than raw data, the statistic t has a specific distribution expressed by the Dickey–Fuller table.

There are three main versions of the test:

1. Test for a unit root: when α = 0 and β = 0,
2. Test for a unit root with drift: when β = 0,
3. Test for a unit root with drift and deterministic time trend: with no restrictions imposed.

(Wikipedia.org, "Augmented Dickey–Fuller test", 27 February 2020)

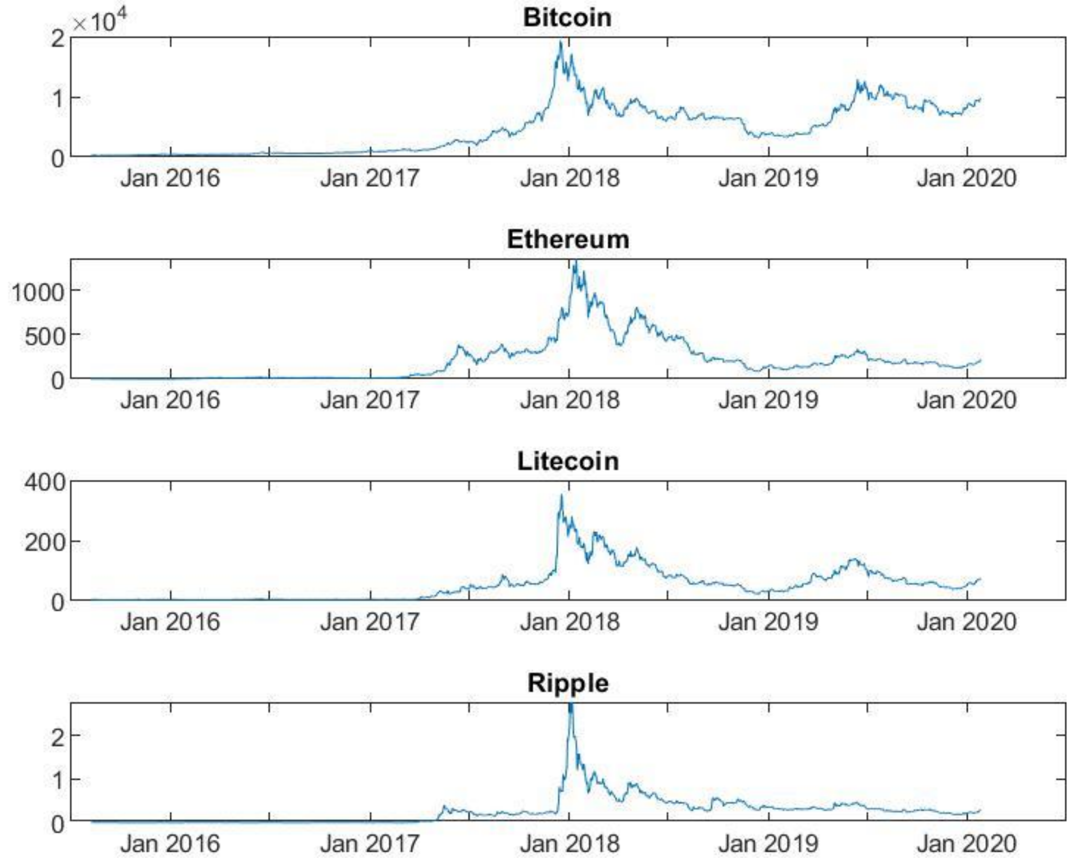

Figure 1: Cryptocurrencies daily prices.

Instead of analysing the raw series one could work on log returns defined as $\mathrm{r}_t = \log\left(\frac{\mathrm{price}_t}{\mathrm{price}_{t-1}}\right)$.

By manipulating the data in this way, all the series appear to be stationary. Formal results about stationarity can be obtained again by assessing the stationarity of the series with Augmented Dickey Fuller test. This time, the three versions of the test produced very small P-values for all series. The null hypothesis is rejected at 95% confidence level, meaning that the transformed series are stationary.

Bitcoin	p-value	Test result
AR	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis
AR with drift	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis
AR with drift and trend	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis
Ethereum		
AR	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis
AR with drift	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis
AR with drift and trend	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis
Litecoin		
AR	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis
AR with drift	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis
AR with drift and trend	1.00e^{-03}	p-value < 0.05 critical value, reject the null hypothesis

Ripple		
AR	1.00e-03	p-value < 0.05 critical value, reject the null hypothesis
AR with drift	1.00e-03	p-value < 0.05 critical value, reject the null hypothesis
AR with drift and trend	1.00e-03	p-value < 0.05 critical value, reject the null hypothesis

Table 3: Augmented Dickey-Fuller test results on log-returns series.

All the series are integrated order 1 because they become stationary after they are differenced one time.

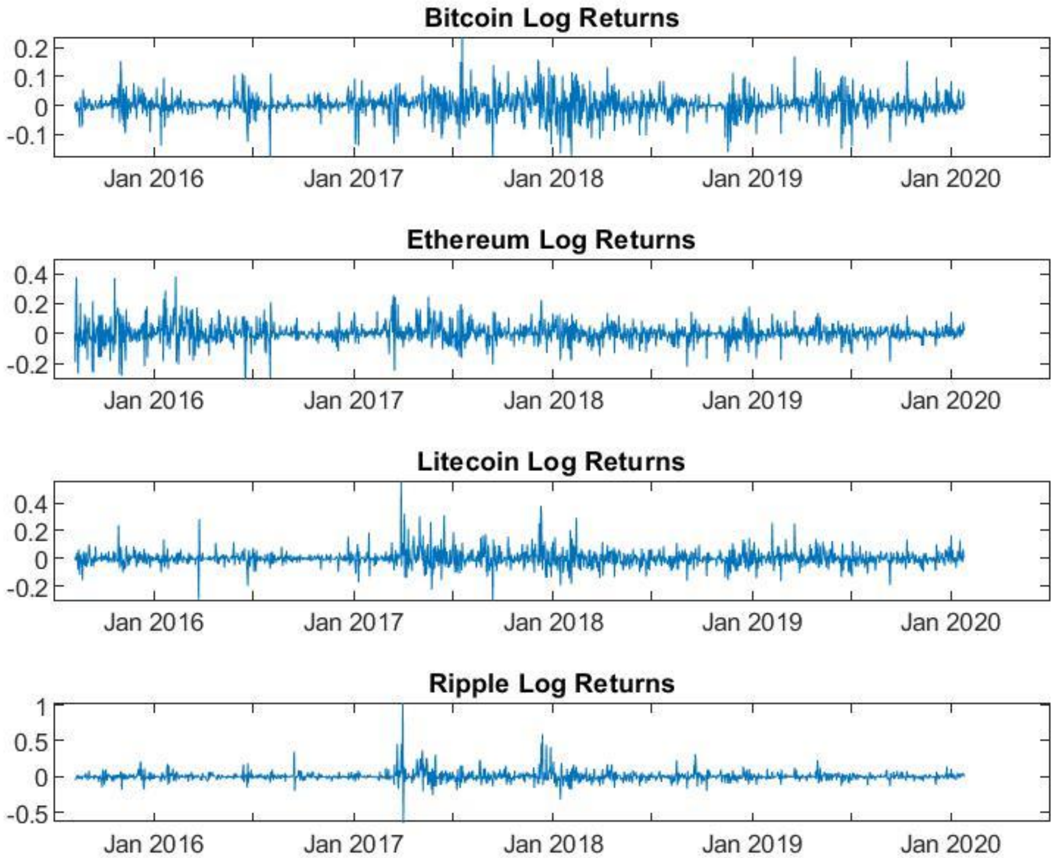

Figure 2: Cryptocurrencies daily log-returns.

Also, the timeseries plots suggest that data appear to be stationary. Plots can provide good information about variables distribution, too. As one can observe, all the series exhibit volatility clusters. Ethereum shows high volatility at the beginning of the series while all the series shows a high volatility period in the cryptocurrencies' market during the period 2017-first quarter of 2018, in concomitance with one major bubble burst in the cryptocurrencies market. Ripple seems to exhibit the highest volatility and Bitcoin the lowest, being the more stable.

In order to have more exhaustive information about their main features, it is convenient to comment also descriptive statistics of the timeseries.

	Bitcoin	Ethereum	Litecoin	Ripple
Max	0,235879838	0,382992252	0,559615788	1,032208312
Min	-0,178691789	-0,307175867	-0,307704504	-0,651701781
Mean	0,002200428	0,003187278	0,001806045	0,002105126
Median	0,00182127	0	-0,00055213	-0,002898394
Standard Dev.	0,039281269	0,066749208	0,05807876	0,071800619
Skewness	-0,093466355	0,218301413	1,314025591	2,824486052
Kurtosis	6,816250439	7,399530464	13,95258852	41,57607971

Table 4: log-returns cryptocurrencies series descriptive statistics.

Ripple seems to be the less stable among the cryptocurrencies for the period observed. It shows the highest maximum value corresponding to a daily 103,22% return and the lowest minimum of -65,17% daily log-return.

These results are justified by the fact that Ripple has also the highest standard deviation and kurtosis. While all the series observed have a kurtosis far higher than 3 (kurtosis of a standard normal distribution), Ripple exhibit a kurtosis of 41,57, almost three times that of Litecoin and six times that of Ethereum.

It is important to look also at distributions' skewness. Bitcoin exhibits negative skewness, meaning that values lower than the mean of the distribution are more likely to be observed than in the normal case (normal distribution has 0 skewness). Ethereum, Litecoin and Ripple, all exhibit positive skewness, with Ripple showing the highest value of 2.82. The probability of observing values higher than the mean, is the highest in the Ripple series.

1. Model Selection

The code used to carry out the research is based on the work of Joshua Chan and Angelia L. Grant, who studied the volatility of several commodities by applying the same univariate GARCH and SV models described in this thesis. (Joshua Chan and Angelia L. Grant, "Modeling Energy Price Dynamics: GARCH versus Stochastic Volatility", Elsevier B.V. All, November 2015).

Although the original code was efficiently written and worked perfectly with original paper's data, a revision was necessary to adapt it to cryptocurrencies' series.

The log marginal Likelihood of each model is computed using the improved cross-entropy method that Chan and Eisenstat discuss in their paper. All the results obtained are gathered and collected in Table 5, where the best models, are evidenced in grey.

	Bitcoin	Ethereum	Litecoin	Ripple
GARCH (1,1)	3,087.9 (0.02)	2,245.7 (0.01)	2,279.8 (0.03)	2,416.1 (0.03)
GARCH (2,1)	3,089.3 (0.05)	2,245.7 (0.03)	2,280.3 (0.05)	2,422.2 (0.04)
GARCH-j	3,226.2 (0.13)	2,195.8 (0.08)	2,380.7 (0.16)	2,591.3 (0.12)

GARCH-m	3,085.2 (0.04)	2,244.5 (0.01)	2,164.4 (0.20)	2,412.2 (0.03)
GARCH-MA	3,086.8 (0.02)	2,242.6 (0.02)	2,287.0 (0.06)	2,413.4 (0.12)
GARCH-t	3,280.8 (0.01)	2,354.4 (0.02)	2,526.9 (0.02)	2,655.9 (0.03)
GARCH GJR	3,086.0 (0.02)	2,242.9 (0.02)	2,273.7 (0.3)	2,417.1 (0.03)
SV	3,268.3 (0.27)	2,356.4 (0.23)	2,512.9 (0.22)	2,644.4 (0.47)
SV (2)	3,275.6 (0.38)	2,353.6 (0.28)	2,516.1 (0.24)	2,651.8 (0.42)
SV-j	3,108.3 (2.68)	2,345.0 (0.54)	2,411.6 (1.52)	2,630.9 (0.35)
SV-m	3,266.0 (0.21)	2,355.7 (0.16)	2,512.0 (0.20)	2,646.5 (0.26)
SV-MA	3,267.1 (0.29)	2,359.1 (0.19)	2,521.6 (0.25)	2,660.8 (0.19)
SV-t	3,301.2 (0.15)	2,365.9 (0.02)	2,536.0 (0.16)	2,649.2 (0.08)
SV-l	3,266.0 (0.34)	2,354.2 (0.19)	2,511.3 (0.23)	2,642.1 (0.24)

Table 5: estimated log marginal Likelihoods for all models considered. Standard deviation is reported between parentheses.

The SV-t model showed better performances than other models for Bitcoin, Ethereum and Litecoin, while the SV-MA performed best for Ripple series. The second models in terms of performances are GARCH-t for Bitcoin, Litecoin and Ripple, and SV-MA for Ethereum. It is important to notice that SV family models performed better than their GARCH counterparts in every case, with the exceptions of GARCH-j for the Bitcoin series, which outperformed the SV-j and GARCH-t which performed better than SV-t for Ripple.

This can be explained by the fact that under stochastic volatility models the log-volatility is a random variable. This make SV models more robust to changes in the time series with respect to GARCH family models. It is possible however to make GARCH models more flexible to misspecifications and changes in the series by assuming that the innovation is distributed following a Student-t or another distribution with heavy tails, or by including jump components. In this way the advantages of Stochastic volatility models become less evident.

The series considered show several changes in the volatility over time. This is why both the GARCH-J and GARCH-t perform way better than the GARCH (1,1) model. The results in the table show clearly that the same holds true for stochastic volatility models. The estimated log marginal likelihoods indicate that SV-j and SV-t performed much better than the SV model.

It is interesting now to determine which features in the model proposed are relevant to describe cryptocurrencies' series volatility. By comparing the results of GARCH (2,1) with GARCH (1,1) and SV

(2) with SV, one can determine whether the AR (2) process defines the volatility behaviour better than the AR (1) considered in the first model of each family.

Both in GARCH and SV models the advantage of assuming an AR (2) process for the conditional volatility is negligible and, in the case of Ethereum, the SV (2) performance is even worse than that of the SV. These results suggest that one could avoid complicating the models and simply maintain an AR (1) assumption for the conditional volatility.

By comparing the results of GARCH-m with GARCH (1,1) and SV-m with SV it is possible to determine whether including the volatility in the returns process provides significant improvements. Results show that this kind of models never provide better results for cryptocurrencies series. The SV-m applied to the Ripple series represents the only exception because the marginal Likelihood increases slightly in this case.

Comparison of GARCH-MA with GARCH and the SV-MA with SV allows to determine whether assuming a moving average process for the innovation of the returns is relevant in the analysis of cryptocurrencies. While in GARCH models the addition of MA component always decreases the marginal likelihood, in SV models it improved the model performances for Ethereum, Litecoin and Ripple series. This indicates that, at least for SV models, the returns errors exhibit serial correlation.

It is finally considered the effect of leverage by comparing the performances of GARCH-GJR with GARCH and the SV-L with SV. While the presence of leverage effect in volatility models provides better results when applied to stocks series, this is not true for cryptocurrencies. The addition of the leverage component produced worse performances in all cases except when GARCH-GJR is applied to Ripple timeseries, where however, the improvement is negligible.

2. Estimation Results

Both the GARCH and SV family models' estimation results are reported for all the cryptocurrencies series. The first four tables contain estimations for the GARCH models. Each table is dedicated to a cryptocurrency timeseries and contains the parameters estimates produced by each model.

	GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
μ	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
α_0	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
α_1	0.21 (0.02)	0.25 (0.03)	0.09 (0.01)	0.22 (0.03)	0.22 (0.02)	0.07 (0.01)	0.2 (0.03)
β_1	0.75 (0.02)	0.54 (0.05)	0.86 (0.01)	0.75 (0.02)	0.74 (0.02)	0.84 (0.02)	0.73 (0.03)
β_2		0.18 (0.04)					
k			0.10 (0)				
μ_k			0 (0)				
σ_k^2			0.01 (0)				
λ				0.72			

				(0.88)			
ψ					0.06 (0.03)		
v						2.60 (0.20)	
δ_1							0.06 (0.04)

Table 6: GARCH models estimation results for Bitcoin timeseries.

	GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
μ	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
α_0	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
α_1	0.20 (0.02)	0.21 (0.03)	0.17 (0.02)	0.19 (0.02)	0.2 (0.02)	0.11 (0.02)	0.2 (0.03)
β_1	0.73 (0.03)	0.62 (0.05)	0.72 (0.02)	0.73 (0.03)	0.73 (0.03)	0.73 (0.04)	0.73 (0.03)
β_2		0.09 (0.03)					
k			0.09 (0.01)				
μ_k			-0.02 (0)				
σ_k^2			0.02 (0)				
λ				1.30 (0.55)			
ψ					0.01 (0.03)		
v						3.21 (0.29)	
δ_1							0 (0.03)

Table 7: GARCH models estimation results for Ethereum timeseries.

	GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
μ	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
α_0	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
α_1	0.06 (0.01)	0.07 (0.01)	0.05 (0)	0.07 (0)	0.06 (0.01)	0.05 (0.01)	0.07 (0)
β_1	0.89 (0.01)	0.57 (0.04)	0.92 (0)	0.85 (0)	0.89 (0.01)	0.86 (0.04)	0.88 (0.01)
β_2		0.3 (0.03)					
k			0.08 (0.01)				
μ_k			0.04				

			(0.01)				
σ_k^2			0.02 (0.01)				
λ				1.52 (0)			
ψ					-0.05 (0.03)		
v						2.96 (0.24)	
δ_1							-0.02 (0)

Table 8: GARCH models estimation results for Litecoin timeseries.

	GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
μ	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
α_0	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
α_1	0.3 (0.03)	0.38 (0.04)	0.18 (0.01)	0.27 (0.03)	0.28 (0.03)	0.14 (0.02)	0.3 (0.03)
β_1	0.67 (0.03)	0.32 (0.05)	0.69 (0.02)	0.68 (0.03)	0.68 (0.03)	0.59 (0.07)	0.68 (0.03)
β_2		0.27 (0.04)					
k			0.08 (0.01)				
μ_k			0.04 (0.01)				
σ_k^2			0.02 (0)				
λ				0.05 (0.32)			
ψ					0.02 (0.03)		
v						2.56 (0.20)	
δ_1							-0.10 (0.03)

Table 9: GARCH models estimation results for Ripple timeseries.

The mean parameter for the cryptocurrencies' return μ, is estimated to be 0 in each series by each model. This result is expected, since the mean of the series were very close to zero as reported in descriptive statistics in Table 4.

The constant parameter defining the conditional volatility in the GARCH expression, α_0 is estimated to be always 0 across all series and models. The parameter weighting the returns' innovation at time t-1, α_1, is positive and varies across the series. Its highest value is estimated for the Ripple series in which results varies across models between 0.14 and 0.38. The lower values instead are estimated for the Litecoin series in which results varies between 0.05 and 0.07.

The β_1 posterior mean estimates range between 0.59 and 0.92, meaning that conditional volatility values are strongly affected by their own previous values. Among all the cryptocurrencies, the lowest values are estimated for the GARCH (2,1) model, in which the sum of β_1 and β_2 is estimated to be always lower than the estimated β_1 in other models. Moreover, β_2 is not that large, suggesting that an AR(2) process does not provide any improvement with respect to the more parsimonious AR(1). This conclusion is supported by the results obtained for the marginal Likelihood.

Looking at GARCH-j models results, it is interesting to notice that the average jump size μ_k is estimated to be negative for Ethereum at -0.02, while it is positive for Litecoin and Ripple and 0 for Bitcoin. The GARCH-j is the worst performing model in the case of Ethereum. The jump probabilities are estimated to be between 0.08 and 0.10, meaning that for the cryptocurrencies considered, about 29-37 jumps per year are observed (8%-10% of 365 trading days in a year, as cryptocurrencies markets are open every day), depending on the series considered.

The posterior estimates of λ, ψ and $\delta 1$ should reflect the marginal likelihood results of GARCH-m, GARCH-MA, GARCH-GJR respectively. This in general holds. The 95% credible interval for most of the estimates of these parameters contains 0, meaning that the parameters may be not relevant, as in that case the models would reduce to a GARCH (1,1). GARCH-m is the most evident exception, because in both the Ethereum and the Litecoin cases, estimated λ are quite large and cannot be 0 with 95% probability. Despite this, the marginal likelihood ranking does not indicate any improvement in the fitting to data with respect to the GARCH (1,1) model. GARCH-MA and GARCH-GJR produce some statistically significant parameters at 0.05 level, but since their estimates are close to 0, they might be not relevant when a lower significance level is considered. The marginal Likelihood ranking does not indicate better performances also in these cases.

The most interesting results are those obtained for Student-t degrees of freedom estimates. The v posterior mean estimates are between 2.56 and 3.21, indicating that the tails of the t distribution followed by the returns' innovations are very heavy, meaning that outliers in cryptocurrencies series are frequent. In fact, this is observed in log-returns timeseries plot, in Figure 2. For this reason, the marginal Likelihood indicates the GARCH-t model as the best in the GARCH family.

Like for GARCH models, the results for SV models are collected in four tables containing parameters posterior estimates.

	SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-l
μ	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
μ_h	-7.22 (0.16)	-7.23 (0.21)	-7.42 (0.32)	-7.21 (0.17)	-7.20 (0.17)	-7.46 (0.55)	-7.21 (0.17)
ϕ_h	0.91 (0.01)	0.94 (0.05)	0.97 (0.01)	0.91 (0.02)	0.92 (0.01)	0.98 (0.01)	0.92 (0.01)
ω_h^2	0.32 (0.03)	0.34 (0.04)	0.10 (0.02)	0.32 (0.07)	0.27 (0.03)	0.07 (0.03)	0.28 (0.01)
ρ_h		-0.04 (0.06)					
k			0.08 (0.01)				
μ_k			0 (0.01)				
σ_k^2			0.01				

			(0)				
λ				0.51 (0.70)			
ψ					-0.05 (0.02)		
v						4.89 (2.91)	
ρ							0.01 (0.06)

Table 10: Stochastic Volatility models estimation results for Bitcoin timeseries.

	SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-l
μ	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
μ_h	-6.04 (0.14)	-5.78 (0.69)	-6.06 (0.15)	-6.04 (0.14)	-6.07 (0.14)	-6.23 (0.23)	-6.01 (0.14)
ϕ_h	0.91 (0.02)	0.79 (0.09)	0.91 (0.02)	0.91 (0.02)	0.90 (0.02)	0.96 (0.01)	0.91 (0.02)
ω_h^2	0.23 (0.05)	0.23 (0.05)	0.22 (0.04)	0.23 (0.04)	0.27 (0.07)	0.10 (0.03)	0.23 (0.04)
ρ_h		0.16 (0.08)					
k			0 (0.01)				
μ_k			0.03 (0)				
σ_k^2			12.52 (7.55)				
λ				1.08 (0.48)			
ψ					-0.09 (0.02)		
v						6.23 (1.67)	
ρ							-0.03 (0.06)

Table 11: Stochastic Volatility models estimation results for Ethereum timeseries.

	SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-l
μ	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
μ_h	-6.41 (0.12)	-6.42 (0.10)	-6.48 (0.21)	-6.44 (0.11)	-6.41 (0.12)	-6.30 (0.37)	-6.42 (0.12)
ϕ_h	0.87 (0.02)	1.00 (0.11)	0.96 (0.01)	0.85 (0.03)	0.87 (0.02)	0.98 (0.01)	0.85 (0.02)
ω_h^2	0.36 (0.02)	0.37 (0.06)	0.08 (0.02)	0.39 (0.08)	0.37 (0.02)	0.04 (0.01)	0.43 (0.03)
ρ_h		-0.18 (0.15)					
k			0.05				

			(0.01)				
μ_k			-0.01 (0)				
σ_k^2			0.02 (0)				
λ				1.21 (0.56)			
ψ					-0.13 (0.03)		
v						9.62 (4.12)	
ρ							-0.10 (0.04)

Table 12: Stochastic Volatility models estimation results for Litecoin timeseries.

	SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-l
μ	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
μ_h	-6.5 (0.15)	-6.50 (0.12)	-6.51 (0.15)	-6.51 (0.13)	-6.49 (0.15)	-6.69 (0.21)	-6.51 (0.14)
ϕ_h	0.89 (0.01)	1.11 (0.06)	0.90 (0.02)	0.86 (0.01)	0.89 (0.01)	0.93 (0.02)	0.88 (0.02)
ω_h^2	0.40 (0.03)	0.34 (0.05)	0.26 (0.08)	0.50 (0.03)	0.37 (0.04)	0.24 (0.07)	0.46 (0.05)
ρ_h		-0.24 (0.08)					
k			0 (0)				
μ_k			0.04 (0.01)				
σ_k^2			3.89 (3.95)				
λ				1.46 (0.43)			
ψ					-0.17 (0.02)		
v						6.24 (2.13)	
ρ							-0.04 (0.04)

Table 13: Stochastic Volatility models estimation results for Ripple timeseries.

Exactly like for GARCH results, Stochastic Volatility models estimate μ to be 0 in each series. The constant parameter within the log conditional volatility expression, μ_h is estimated to be about -6.50, which means that its effect on returns conditional volatility is almost irrelevant across all series and models ($e^{-6.50} \cong 0$).

ϕ_h instead, is estimated to be significant across the various models and series. It ranges from 0.79 to 1.11, indicating that the stochastic volatility process is highly persistent.

The posterior mean of the variances of stochastic volatility processes do not provide much information about the true variances, because estimates change considerably across models and series. In fact, ω_h^2 is estimated to be between 0.07 - 0.34 for the Bitcoin series, 0.10 - 0.27 for the Ethereum, 0.04 - 0.43 for the Litecoin and 0.24 - 0.50 for the Ripple.

The SV (2) model marginal Likelihoods suggest that the model performs slightly better than the simple SV, except for the Ethereum series. The estimates observed for the former model indicate that the log conditional volatility depends negatively on its own values at time t-2. It must be considered that, apart for Ripple timeseries, ρ_h estimates have a large standard deviation, so it is difficult to understand their true impact.

In the SV-j model, the posterior mean estimate of probability of jumps is 0 both in the case of Ripple and Ethereum, meaning that the inclusion of the jump parameter is not relevant for the two series. For Bitcoin the estimated value of k ranges between 0.07 and 0.09 indicating that the jumps are present in the data, but their average jump size μk is just between -0.01 and 0.01 and σ_k^2 estimate is 0.01. These results indicate that the model allows for very small jumps in the data, meaning that it is not providing better results than the simple SV. The same conclusions can be drawn for the Litecoin timeseries, while the probability is relevant and significant, the average value of jumps and their variance are estimated to be very small. These results are supported by the marginal likelihood ranks for the SV-j model in Table 5.

The SV-m parameter λ is estimated to be positive and significant for all timeseries except for the Bitcoin in which its 95% credible interval is between -0.19 and 1.21. This would suggest that volatility feedback is relevant for modelling cryptocurrencies but, like for GARCH, the marginal Likelihood ranking does not suggest any significant improvement in performances with respect to the benchmark model, the SV model, in this case.

ψ represents the estimate of the moving average parameter in SV-MA models, which is negative and significant for all timeseries. The most evident result is observed for Ripple series, in which the posterior mean of ψ takes the smaller value of -0.17 (with standard deviation 0.02). The second lowest estimate is observed for the Litecoin series in which a value of -0.13 (standard deviation of 0.03) is estimated. These results suggest that the innovations in the SV equation depends negatively on their own past values. For Ripple and Litecoin the results are also supported by marginal likelihood ranking, the SV-MA model is preferred to SV and, for Ripple it is the best performing model.

Next, 95% credible interval of the coefficient ρ excludes 0 only for the Litecoin series, indicating that allowing for jumps in the data does not produce better results in SV models. In fact, the marginal Likelihood estimates for the SV-j are slightly worse than for the SV model.

As for GARCH models, the best results according to marginal Likelihood are obtained by assuming t-distributed innovations for the returns (the only exception being the SV-MA model for the Ripple series). In the SV-t, ν posterior mean estimates are larger than those in the GARCH-t models, but they still indicate fat tails.

3. Univariate GARCH Models Residual Analysis

In this section are reported the Q statistics from Ljung-Box tests and Q^2 statistics from McLeod-Li tests, run respectively on 20 lags of autocorrelation of models standardized residuals and squared standardized residuals. For both tests, the 5% and 1% critical values are 31.41 for the former and 37.57 for the latter. These values can be gathered by the table of a Student-t distribution with 20 degrees of freedom. As one could observe by looking at results in the next tables, they vary across the different models and series considered.

	GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
Q	32.2858 (0.6979)	32.5077 (0.6115)	39.6781 (0.6045)	33.0562 (1.2003)	28.4461 (1.4689)	38.8544 (1.5018)	32.4558 (0.7256)
Q^2	19.5012 (0.8275)	20.0042 (0.7651)	18.4740 (0.3002)	20.9697 (0.8633)	22.0689 (0.8974)	20.4649 (0.8547)	22.6820 (0.9244)

Table 14: Ljung-Box Q statistics and McLeod-Li Q^2 statistics estimation results for Bitcoin timeseries.

	GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
Q	34.8392 (0.7616)	35.0072 (0.8198)	31.9889 (0.7362)	32.4519 (0.8375)	35.4426 (2.0383)	32.2896 (0.8825)	34.8474 (0.8060)
Q^2	9.1690 (0.9611)	10.2392 (0.8516)	16.1436 (0.7658)	9.3887 (0.8451)	10.5956 (0.7031)	13.3114 (2.1044)	9.5278 (0.6978)

Table 15: Ljung-Box Q statistics and McLeod-Li Q^2 statistics estimation results for Ethereum timeseries.

	GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
Q	24.1774 (0.4892)	24.2902 (0.3836)	19.5682 (0.2109)	23.3046 (0.0605)	23.3228 (1.5321)	19.3438 (0.6448)	24.5783 (0.4205)
Q^2	5.9412 (0.3937)	4.9469 (0.3015)	5.7361 (0.1798)	5.0054 (0.1530)	5.6751 (0.4454)	6.0510 (0.2471)	6.1729 (0.3886)

Table 16: Ljung-Box Q statistics and McLeod-Li Q^2 statistics estimation results for Litecoin timeseries.

	GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
Q	17.3874 (0.5694)	16.3361 (0.7603)	13.8867 (0.3399)	17.4572 (1.7071)	18.4758 (1.4762)	15.5212 (1.2340)	18.0971 (0.6822)
Q^2	8.4802 (1.1014)	12.6454 (1.3101)	5.9016 (0.1570)	9.2438 (0.8133)	8.3146 (0.7481)	6.7883 (1.7078)	7.3464 (0.8529)

Table 17: Ljung-Box Q statistics and McLeod-Li Q^2 statistics estimation results for Ripple timeseries.

For the Bitcoin series, the null hypothesis of no correlation among the residuals is rejected at 0.05 and 0.01 significance level for GARCH-j, and GARCH-t, it is not rejected at all levels for the GARCH-MA while it is not rejected only at 1% significance level for other models. In the Ethereum case, the null hypothesis must be rejected for all models at 5% significance level, but it is not rejected at lower levels. Q-statistics in both Litecoin and Ripple case suggest to not reject the null hypothesis of no serial correlations among the residuals at all significance level.

These results support the ranking of the marginal likelihood, if autocorrelations among the residuals were detected, the most appropriate model would have been the GARCH-MA model. In this case there are not clear autocorrelation evidences and the marginal Likelihood indicates that GARCH-t model is the best performing model in the GARCH family. The only exception is in the Bitcoin analysis, where the marginal Likelihood favours the GARCH-t over the GARCH-MA but the presence of autocorrelation among residuals of GARCH-t would suggest that GARCH-MA is the best fitting model.

Next, the estimated Q^2 statistics strongly indicate to not reject the null hypothesis of no autocorrelation in the squared residuals for all cryptocurrencies at all significance level. The absence of autocorrelation in the squared residuals can be interpreted as adequacy of fitted models in describing the volatility process.

Finally, for the best GARCH model selected by the marginal Likelihood criterium, the Kolmogorov-Smirnov test is run on the residuals to verify whether the assumed distributions for the innovations in the models are correct. The GARCH-t model is the best fitting model for all cryptocurrencies so, the estimated residuals from the models are compared to Student t distributions. The degree of freedom assumed for the theoretical distributions are those estimated in the GARCHs models (parameter v) The following table reports the results for the Kolmogorov-Smirnov tests.

Cryptocurrency	Model	Distribution	p-value	Test result
Bitcoin	GARCH-t	Student-t Df = 2.60 (0.20)	0.4588	p-value > 0.05 critical value, do not reject the null hypothesis: residuals follow a Student t distribution
Ethereum	GARCH-t	Student-t Df = 3.21 (0.29)	0.03569	p-value < 0.05 critical value, reject the null hypothesis: residuals do not follow a Student t distribution
Litecoin	GARCH-t	Student-t Df = 2.96 (0.24)	0.2644	p-value > 0.05 critical value, do not reject the null hypothesis: residuals follow a Student t distribution
Ripple	GARCH-t	Student-t Df = 2.56 (0.20)	0.09099	p-value > 0.05 critical value, do not reject the null hypothesis: residuals follow a Student t distribution

Table 18: Kolmogorov-Smirnov test on GARCHs residuals.

The test suggests that assumptions made for the innovations distribution in the models are correct, except for the GARCH-t fitted to Ethereum timeseries. When looking at these results though, one should consider that they are affected by the standard deviations of the degrees of freedom estimates. For example, in the Litecoin case, it is different to assume that the degree of freedom of the theoretical distribution are 2.96, 2.72 (2.96 – 0.24) or 3.2 (2.96 + 0.24).

To investigate further the results obtained for the Litecoin GARCH-t, it is convenient to observe the Q-Q plot of these residuals, which graphically compare them with their theoretical distribution.

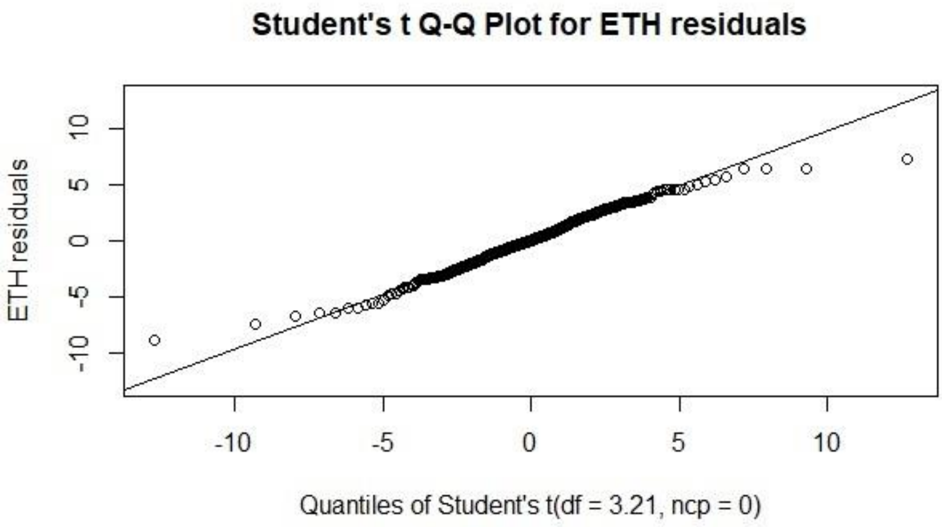

Figure 3: Ethereum residuals from GARCH-t model, compared to their theoretical distribution.

The plot indicates that the Student t distribution has fatter tails than the standardized residuals distribution. This means that the Student t assumption overestimate the probability of observing extreme values in the residuals' series.

4. Univariate Stochastic Volatility Models Residual Analysis

Q statistics from Ljung-Box tests and Q^2 statistics from McLeod-Li tests are reported also for the SV family of models. The 5% and 1% critical values are respectively 31.41 and 37.57 as before. The standard deviations of the estimated Q-statistics and Q^2-statistics are quite large for SV models, and interpretations of results are not clear cut.

	SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-l
Q	24.8546 (2.5590)	25.4544 (3.7033)	23.2393 (5.8970)	23.6703 (3.9804)	27.0618 (4.4565)	26.8618 (0.2829)	24.6850 (2.9886)
Q^2	38.7468 (8.7988)	25.0911 (8.9953)	25.2512 (6.5026)	23.0317 (3.7751)	26.4592 (7.7758)	17.8909 (16.0824)	27.3793 (5.9818)

Table 19: Ljung-Box Q statistics and McLeod-Li Q^2 statistics estimation results for Bitcoin timeseries.

	SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-l
Q	31.1057 (2.3880)	31.3811 (3.1420)	31.8397 (3.7699)	30.3274 (3.2486)	31.0076 (4.9093)	30.6452 (3.8673)	31.4408 (3.8324)
Q^2	31.6783 (4.0682)	14.3015 (7.4103)	22.3617 (5.6444)	15.6958 (3.6713)	20.1308 (1.1163)	18.8441 (8.2430)	20.5291 (5.5939)

Table 20: Ljung-Box Q statistics and McLeod-Li Q^2 statistics estimation results for Ethereum timeseries.

	SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-l
Q	27.5394 (3.6370)	29.1090 (3.7842)	31.5680 (4.8832)	29.1151 (2.8438)	20.1886 (3.5728)	23.5707 (3.4888)	28.1480 (3.1011)
Q^2	34.6285 (6.5406)	20.5710 (4.5459)	15.2383 (6.1840)	21.6038 (3.2909)	12.3443 (5.2621)	49.1955 (54.3918)	23.0824 (4.7394)

Table 21: Ljung-Box Q statistics and McLeod-Li Q^2 statistics estimation results for Litecoin timeseries.

	SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-l
Q	37.9999 (4.9305)	35.4964 (3.8899)	35.1578 (6.5922)	37.6165 (6.7897)	20.0068 (2.9938)	28.8156 (4.0648)	33.1626 (4.8709)
Q^2	26.4222 (7.9600)	32.4781 (5.9926)	23.5247 (4.7598)	45.9198 (0.6014)	18.3643 (4.1382)	52.8331 (15.7990)	20.0238 (9.0316)

Table 22: Ljung-Box Q statistics and McLeod-Li Q^2 statistics estimation results for Ripple timeseries.

For Bitcoin series, the null hypothesis of no serial correlation among the standardized residuals is not rejected at 0.05 and 0.01 significance level for all the models.

In the Ethereum case, the null hypothesis must be rejected at 5% significance levels for the SV-j and SV-l, but it cannot be rejected at 1% level for all models.

Next, the Litecoin Q-statistics indicate H_0 of no autocorrelation should not be rejected at all significance levels for all models. SV-j residuals is the only exception as for this series the null hypothesis must be rejected at 5% level.

For the Ripple series, the null hypothesis is not rejected at all significance level for the SV-MA and SV-t, it is rejected at 5% but it is not rejected at 1% significance level for the SV (2), SV-j and SV-l. Finally, H_0 is rejected at all levels for SV and SV-M.

The results, again, support the ranking of the marginal Likelihood. More clues about possible correlations among the residuals is detected only in the case of the Ripple series, while it is less evident for other cryptocurrencies. The marginal Likelihood, in fact, favour the GARCH-MA model for the Ripple and the GARCH-t for other series.

The estimated Q^2 statistics strongly suggest to not reject the null hypothesis of no autocorrelation in the squared residuals for Bitcoin, except for the test run on SV standardized residuals, for which the null hypothesis must be rejected at both 5% and 1% level.

In the case of Ethereum the null hypothesis cannot be rejected for all models except the SV where the McLeod-Li statistic is slightly above the 5% critical value and hence, the null must be rejected.

The null for SV and SV-t should be rejected in the Litecoin squared residuals analysis and finally, in Ripple squared residuals it is detected autocorrelation in residuals of the SV (2), SV-M and SV-t. In interpreting these results, one should be very careful because the standard deviations of the McLeod-Li are very large sometimes, making the result not statistically significant. Overall, the squared residuals autocorrelation results support the marginal Likelihood ranking but, since many statistics are not reliable it is not possible a correct interpretation of the tests for all models.

As already done for the GARCH models, the Kolmogorov-Smirnov test is run on the residuals of the best fitting SV models, to verify whether the assumed distributions for the errors are correct. In the case of Bitcoin, Ethereum and Litecoin the best model is the SV-t, while it is the SV-MA for the Ripple series. The SV-t models residuals are compared to the estimated Student t distributions and the SV-MA residuals are compared to a standard Normal distribution. The degree of freedom assumed for the theoretical distributions are again, those estimated in the models (parameter v). The following table reports the results for the Kolmogorov-Smirnov test.

Cryptocurrency	Model	Distribution	p-value	Test result
Bitcoin	SV-t	Student-t Df = 4.89 (2.91)	0.07377	p-value > 0.05 critical value, do not reject the null hypothesis: residuals follow a Student t distribution
Ethereum	SV-t	Student-t Df = 6.23 (1.67)	0.3725	p-value > 0.05 critical value, do not reject the null hypothesis: residuals follow a Student t distribution
Litecoin	SV-t	Student-t Df = 9.62 (4.12)	0.00129	p-value < 0.05 critical value, reject the null hypothesis: residuals do not follow a Student t distribution
Ripple	SV-MA	Standard Normal	0.3970	p-value > 0.05 critical value, do not reject the null hypothesis: residuals follow a Student t distribution

Table 23: Kolmogorov-Smirnov test on GARCHs residuals.

For the SV models, the standard deviations of the estimated degree of freedom are even larger than those estimated in the GARCH case. The results of the test would change dramatically if they are run considering the maximum or minimum value of the degree of freedom indicated by their standard deviation. With such large standard deviations of the estimated degree of freedom, the test results are not much indicative. The residuals Q-Q plots are more helpful in this case to analyse their sample distributions.

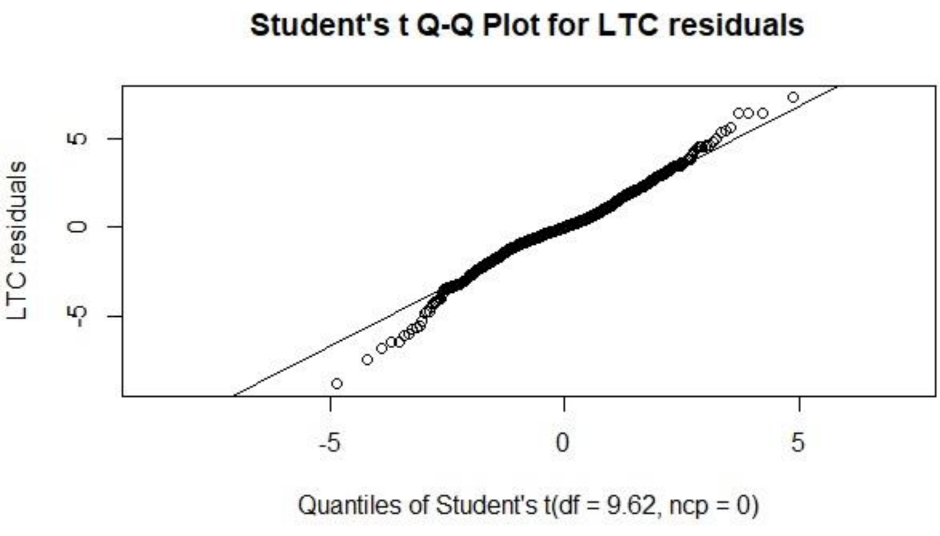

Figure 4: Litecoin residuals from SV-t model, compared to their theoretical distribution.

The quantile-quantile plot shows clearly that the theoretical distribution assumed for innovations, is overestimating the probability of extreme events. Residuals are fat tailed in comparison with the theoretical distribution and the left tail is larger than the right one, meaning that maybe, a skewed distribution for the innovations would provide better results. One should take these results with a grain of salt, as this is the plot obtained considering a theoretical distribution with 9.62 degrees of freedom. By considering 5.5 degrees of freedom (9.62 – 4.12), the minimum value in the credible interval of *v* for the Litecoin residuals distribution, the result seems to improve, even though there is still evidence that a skewed distribution for the innovations would be more appropriate to fit the data. In other words, the problem might be the unprecise degrees of freedom estimate.

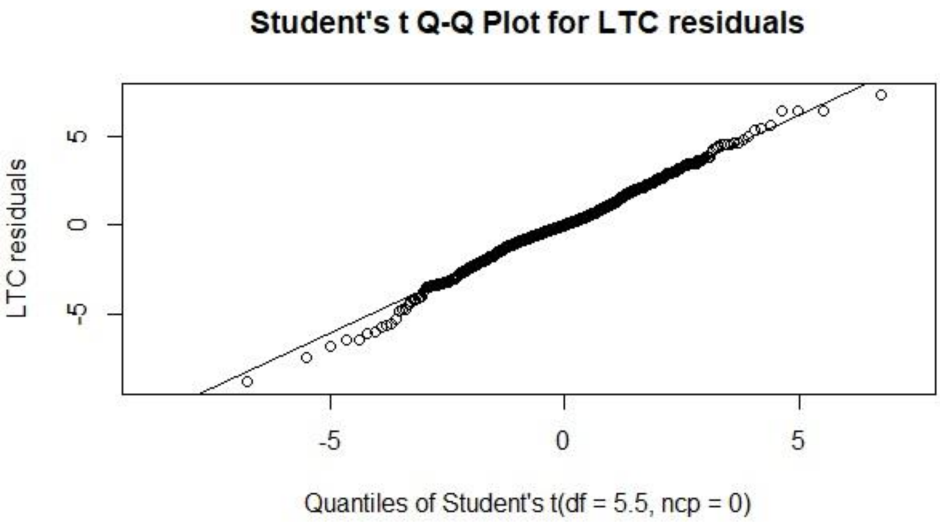

Figure 5: Litecoin residuals from SV-t model, compared to their theoretical distribution.

5. Univariate GARCH Models Forecasting

The following table gathers the predictive marginal Likelihood (m-LK). They are obtained by filtering the out of sample data, the cryptocurrencies' log-returns for the period 08/02/2020 – 17/05/2020 with the fully specified models estimated on all data available (the historical data + the out of sample data).

		GARCH (1,1)	GARCH (2,1)	GARCH-J	GARCH-M	GARCH-MA	GARCH-t	GARCH-GJR
BTC	m-LK	76.6 (0.30)	75.6 (0.27)	73.2 (0.43)	79.0 (0.46)	75.0 (0.41)	139.8 (0.19)	98.4 (0.33)
ETH	m-LK	86.0 (0.22)	83.6 (0.27)	82.9 (0.30)	83.9 (0.31)	83.9 (0.33)	116.5 (0.38)	88.2 (0.27)
LTC	m-LK	103.7 (0.29)	103.1 (0.32)	101.8 (0.18)	103.3 (0.17)	103.8 (0.18)	127.3 (0.19)	103.7 (0.20)
XRP	m-LK	116.0 (0.13)	114.2 (0.17)	109.0 (0.20)	113.0 (0.17)	114.1 (0.11)	137.2 (0.21)	112.6 (0.27)

Table 24: Predictive marginal Likelihoods of univariate GARCH models.

The best model according to predictive marginal Likelihood is the GARCH-t for all cryptocurrencies timeseries. This is the same result obtained with marginal Likelihood computed fitting historical data. This means that by using the models chosen by the marginal Likelihood in the step of fitting the historical data, one could have obtained the best predictions for cryptocurrencies daily log-returns.

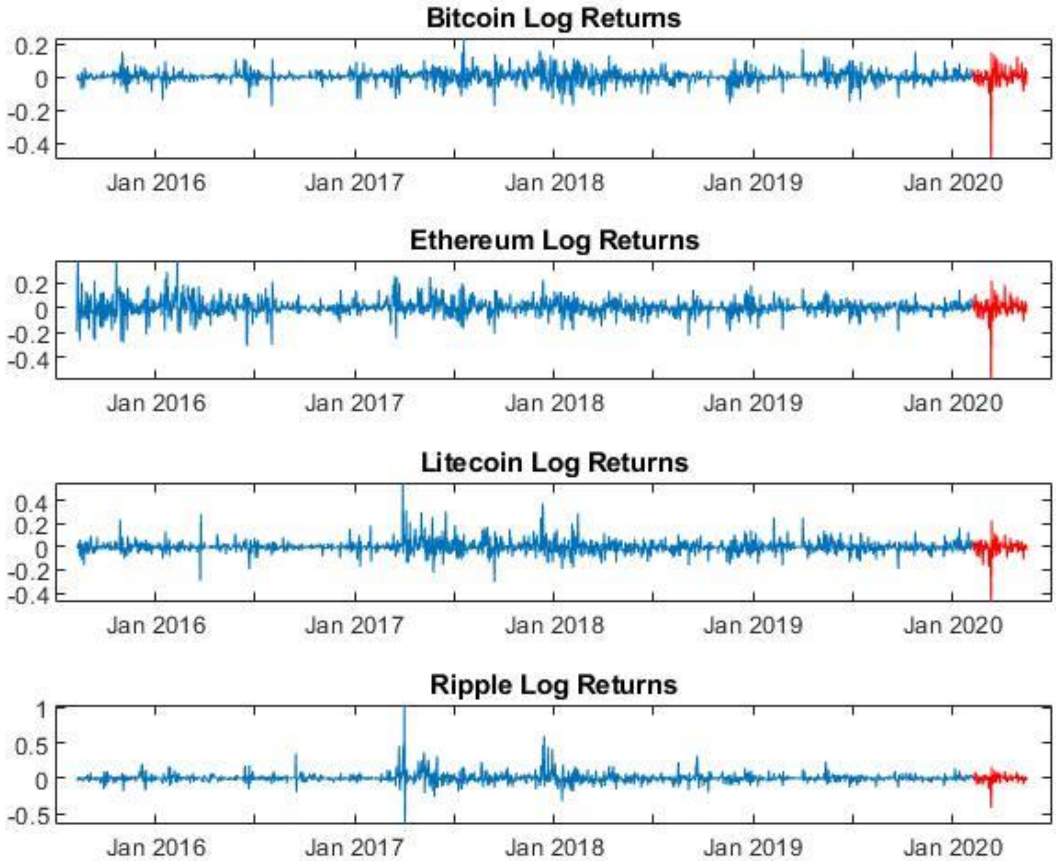

Figure 6: Cryptocurrencies daily log-returns.

This result suggests that cryptocurrencies series continued to exhibit large volatility with extreme values observations.

In fact, it can be observed from the plots that volatility increased during the out of sample period. This is likely due to the impact that recent global coronavirus emergency had on markets. Though, cryptocurrencies' market has always been more volatile than other financial markets and even larger fluctuations were observed in the training period.

6. Univariate Stochastic Volatility Models Forecasting

Since volatility in markets increased substantially since the end of the training sample, one could expect that Stochastic Volatility models perform better than their GARCH counterparts. Stochastic Volatility models results of predictive marginal Likelihood (m-LK) are reported in the next table.

		SV	SV (2)	SV-J	SV-M	SV-MA	SV-t	SV-L
BTC	m-LK	138.7 (0.13)	145.1 (0.18)	105.3 (0.26)	136.4 (0.18)	136.2 (0.20)	147.8 (0.22)	138.4 (0.23)
ETH	m-LK	115.2 (0.24)	115.6 (0.14)	98.2 (0.25)	113.06 (0.16)	112.3 (0.15)	120.6 (0.29)	112.2 (0.17)
LTC	m-LK	127.0 (0.14)	136.0 (0.22)	113.4 (0.32)	125.58 (0.10)	123.5 (0.18)	131.7 (0.14)	126.0 (0.09)
XRP	m-LK	134.6	143.4	123.5	133.23	132.1	141.9	131.8

		(0.19)	(0.35)	(0.19)	(0.13)	(0.15)	(0.06)	(0.23)

Table 25: Predictive Marginal Likelihoods of univariate Stochastic Volatility models.

The best SV models according to predictive marginal Likelihood are the SV-t for Bitcoin and Ethereum series and the SV (2) for Litecoin and Ripple series. For Bitcoin and Ethereum series, this is the same result obtained with marginal Likelihood obtained fitting historical data, meaning that SV-t would have produced the best daily forecasts. For Litecoin and Ripple series, SV (2) models perform better in the out of sample data than the SV-t and SV-MA, which were the best performers on historical data. The models selected fitting past observations would have not produced the best forecasts for those two series.

More interestingly, every SV-model perform better than its GARCH counterpart when both are fitted to the out of sample data. This founding is in line with results obtained on historical data. Stochastic Volatility processes are more appropriate to model the extremely large volatility in cryptocurrencies' markets.

V Multivariate Models

In this thesis two famous models are used in the attempt of modelling the conditional mean of cryptocurrencies log-returns: Vector Autoregressive models (VARs) and the Vector Error Correction models (VECs). While VAR models are applied to cryptocurrencies in order to capture linear interdependencies between a set of time series, VECs models can be applied only in case cointegrating relations are detected through apposite tests. Though, if such relation is detected they can provide much more information about the behaviour of the timeseries with respect to VARs.

Next, the DCC GARCH, a famous multivariate GARCH model, is fitted to the residuals of VAR and VEC models to describe the dynamic dependencies of the conditional volatility matrix. All these models and cointegration phenomenon are introduced more in detail in the following sections.

1. Vector Autoregression Models

The Vector Autoregressive model or VAR (n,P) model is basically a multivariate Autoregression (AR) which allows for more than one response variable. It is a stationary multivariate stochastic process model composed of a system of *n* equations describing the linear interdependencies between *n* variables, *n*P* lagged responses and other terms.

The most common expression for a VAR (n,P) model is the reduced form (MathWorks Inc, "varm", 1994-2020) :

$$Y_t = c + \Phi_1 Y_{t-1} + \Phi_2 Y_{t-2} + \ldots + \Phi_p Y_{t-p} + \beta x_t + \delta_t + \varepsilon_t \qquad (22)$$

Where t = 1, ... ,T.

- Y_t is a n x 1 vector containing the values of the n response variables at time t;
- c is a n x 1 vector of constants;

- Φ_j is a n x n matrix of autoregressive coefficients, where j = 1,…,p;
- β is a n x k matrix of regression coefficients;
- x_t is a k x 1 vector of exogenous variable values at time t;
- δ_t is a n x 1 vector of linear time-trend values;
- ε_t is a n x 1 vector of random Gaussian innovations, each with 0 mean and a n x n covariance matrix Σ. ε_t and ε_s are independent for t ≠ s;

A VAR model needs to be stationary to properly examine the statistical significance of estimated coefficients. The conditions to be satisfied to grant that the model is stationary are:

- The expected value of Yt must be equal to its unconditional mean μ at any time t;
- The covariance matrix of the model must not depend on time.

A VAR model could be also represented in matrix form. While this representation is easier to understand it is not appropriate to express rich models because the VAR equation becomes very large as parameters are added to the model. Here it is shown the example of a model with just two response variables:

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} + \begin{bmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{bmatrix} \begin{bmatrix} y_{1,t-1} \\ y_{2,t-1} \end{bmatrix} + \cdots + \begin{bmatrix} a_{1,1}^p & a_{1,2}^p \\ a_{2,1}^p & a_{2,2}^p \end{bmatrix} \begin{bmatrix} y_{1,t-p} \\ y_{2,t-p} \end{bmatrix} + \begin{bmatrix} b_{1,1} & \cdots & b_{1,k} \\ b_{2,1} & \cdots & b_{2,k} \end{bmatrix} \begin{bmatrix} x_{1,t} \\ \vdots \\ x_{k,t} \end{bmatrix} + \begin{bmatrix} \delta_1 \\ \delta_2 \end{bmatrix} t + \begin{bmatrix} e_{1,t} \\ e_{2,t} \end{bmatrix}$$

Each of the n variables at time t is described by one equation, which express it as a function of its own lagged values and the lagged values of each other variables in the model.

VAR models are quite simple to understand and have several advantages over other multivariate timeseries models because their only required input is a set of variables which can be hypothesized to affect each other over time.

2. Cointegration Analysis

Before introducing VEC models it is essential a review of cointegration phenomenon. If detected by appropriate statistical tests, the presence of cointegration among two or more variables, suggests that a multivariate VEC model would be more convenient than a VAR model to fit the data, because the former provides useful information about both short run and long run relationships among variables.

A set of time series which are integrated of order d are said to be cointegrated if a linear combination of those variables is integrated of order less than d. The simplest case is when individual series are I (1) (integrated of order 1) but a linear combination of them is stationary (integrated of order 0).

Cointegration has acquired increasing importance in time series analysis because, if detected among two or more non-stationary variables, it makes the usual hypotheses testing procedure based on OLS regressions on differenced data to be biased. Cointegration is useful to detect spurious correlation in the case of integrated series, (this is the case whenever two series show a significant correlation, but they are not causally related).

Two very famous tests for detecting cointegrating relations among timeseries variables are used in this thesis: the Engle-Granger test and the Johansen test.

Engle-Granger Cointegration Test

The Engle-Granger test takes its name by two statisticians who first proposed it, Robert Engle and Clive Granger. The test is implemented through two steps.

From the definition of cointegration, if two timeseries ($y_{1,t}$ and $y_{2,t}$) are non-stationary and integrated of order 1, and a linear combination of them like $y_{1,t}$ = $\beta y_{2,t}$ + u_t is found to be stationary for some value of β and u_t, then they are cointegrated. In case of cointegration, u_t, the residuals of the regression, are stationary.

The first step of the test is thus to estimate β through OLS. Next the stationarity of estimated residuals is tested with a unit root test such as the Dickey Fuller (Wikipedia, the free encyclopedia, "Cointegration", 17 December 2019).

While the Engle Granger is easy to understand and straight forward to implement it has several weaknesses. Firstly, it is a test for cointegration that allows for only one cointegrating relationship to be tested at time. Secondly, it has the problem of requiring two steps: any error which may occur in the first step will be carried into the second step.

Johansen Cointegration Test

The Johansen test avoids using two step estimators, since it relies on maximum Likelihood estimation technique. Moreover, it can test for the presence of multiple cointegrating vectors at once. There are several empirical evidences suggesting that Johansen test performs significantly better than the simpler Engle-Granger test (Munich Personal RePEc Archive, "Stationarity and cointegration tests: Comparison of Engle - Granger and Johansen methodologies", 1998).

The Johansen test takes its name from its creator Søren Johansen. Johansen test can be carried out in two ways, with eigenvalues or with trace. The estimates for the parameters and the inference result may vary slightly, using a test or the other. (Hänninen, R. (2012). "The Law of One Price in United Kingdom Soft Sawnwood Imports – A Cointegration Approach". Modern Time Series Analysis in Forest Products Markets. Springer. p. 66.).

In both the trace test and the eigenvalues test, the null hypothesis is that the number of cointegration vectors is r = r* < n, where r* = 1,2,…n-1 and n is the number of variables to be tested. The test is run for each different rank r* and the estimate of r produced by the test is taken from the first non-rejection of the null hypothesis. The alternative hypothesis changes in the two cases. In the trace test the alternative hypothesis is r = n, while in the eigenvalues test is r = r* + 1 (Wikipedia, the free encyclopedia, "Johansen test", 15 April 2019).

The tests run with Matlab software also produce maximum Likelihood estimates for the parameters in a Vector Error-Correction model (VEC), which will be reviewed in the next section (1994-2020 The MathWorks, Inc, "jcitest documentation", 2011).

3. Vector Error Correction Models

The Vector Error Correction Model (VEC) can be written in a similar way as a VAR model but it includes a number of cointegrating relations existing among the n response timeseries variables.

These relations are included in the models in the form of an error-correction term, which are linear functions of the responses in levels. The number of cointegrating relations in the model is expressed by the rank r that can be estimated through the Johansen test for cointegration.

The components of VEC model can be: an autoregressive polynomial for each differenced response variable, in which each variable is function of its own differenced past values and those of other response series (this is the short run component of VEC), a constant, a time trend, exogenous predictor variables, and a constant and time trend in the error-correction term. A generalized VEC (n,r,q) model can be written in component reduced form as

$$\Delta y_t = A(B' y_{t-1} + c_0 + d_0 t) + c_1 + d_1 t + \Phi_1 \Delta y_{t-1} + \cdots + \Phi_{p-1} \Delta y_{t-(p-1)} + \beta x_t + \varepsilon_t \quad (23)$$

Where $t = 1,...,T$, n is the number of differenced response timeseries, r is the cointegrating rank and q is the lag for the autoregressive system of polynomial of the differenced response variables (q = p-1, where p is the lag of a corresponding VAR model).

In the equation:

- y_t is an n x 1 vector of values corresponding to n response variables at time t;
- $\Delta y_t = y_t - y_{t-1}$;
- r is the number of cointegrating relations, 0 < r < n;
- A is an n x r matrix of adjustment speeds coefficients;
- B is an n x r cointegration matrix;
- c_0 is an r x 1 vector of constants in the cointegrating relations;
- d_0 is an r x 1 vector of linear time trends in the cointegrating relations;
- c_1 is an n x 1 vector of constants, representing linear trends in the level of data;
- d_1 is an n x 1 vector of linear time-trend values, representing deterministic quadratic trends in the level of data;
- Φ_j is an n x n matrix of short-run coefficients, for j = 1,...,$p - 1$;
- β is an n x k matrix of regression coefficients;
- x_t is a k x 1 vector of exogenous variable values at time t;
- ε_t is a n x 1 vector of random Gaussian innovations, each with 0 mean and a n x n covariance matrix Σ. ε_t and ε_s are independent for t ≠ s.

It is important to notice that if n = r, then the response variables in levels are stationary, meaning that the VEC (n,r,P-1) model can be rewritten as a stable VAR (n,P) model. If r = 0 instead, the VEC (n,r,P-1) model is a stable VAR (n,P–1) model. Both these two cases are not of interest if one is aiming at implementing a cointegration analysis. (1994-2020 The MathWorks, Inc, “vecm documentation”, 2017).

Johansen Form

The VEC model can be rewritten also as

$$\Delta y_t = C y_{t-1} + \Phi_1 \Delta y_{t-1} + \cdots + \Phi_q \Delta y_{t-q} + DX + \varepsilon_t \quad (24)$$

In the case 0 < r < n there is at least one cointegrating relation among the variables and C is an n x n impact matrix obtained from the multiplication of adjustment matrix by cointegration matrix. It is defined as C = AB'. DX is a general expression for any deterministic trend in the data that is included in the model, while the other elements are defined as before. Johansen Forms are different expressions of Cy_{t-1} + DX that it is possible to define while running the Johansen test for cointegration or in a VEC model. Possible values of Johansen Form that can be estimated in Johansen test are:

H2 $\quad \mathrm{AB}'y_{t-1}$ There are no intercepts or trends in the cointegrated series and there are no deterministic trends in the levels of the data.

H1* $\quad \mathrm{A}(\mathrm{B}'y_{t-1} + c_0)$ There are intercepts in the cointegrated series and there are no deterministic trends in the levels of the data.

H1 $\quad \mathrm{A}(\mathrm{B}'y_{t-1} + c_0) + c_1$ There are intercepts in the cointegrated series and there are deterministic linear trends in the levels of the data. This is the default value.

H* $\quad \mathrm{A}(\mathrm{B}'y_{t-1} + c_0 + d_0 t) + c_1$ There are intercepts and linear trends in the cointegrated series and there are deterministic linear trends in the levels of the data.

H $\quad \mathrm{A}(\mathrm{B}'y_{t-1} + c_0 + d_0 t) + c_1 + d_1 t$ There are intercepts and linear trends in the cointegrated series and there are deterministic linear and quadratic trends in the levels of the data.

(1994-2020 The MathWorks, Inc, "jcitest documentation", 2011).

For example, a simple VEC (2,1,1) with two timeseries variables, cointegration rank 1, 1 lag for the autoregressive system of polynomial, and with the single cointegration equation expressed in the Johansen 'H1' form, can be written in matrix form as

$$\begin{bmatrix} \Delta y_{1,t} \\ \Delta y_{2,t} \end{bmatrix} = \begin{bmatrix} c_1^0 \\ c_2^0 \end{bmatrix} + \begin{bmatrix} \alpha_{1,1} \\ \alpha_{2,1} \end{bmatrix} ([\beta_{1,1} \quad \beta_{1,2}]\, y_{t-1} + c^1) + \begin{bmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{bmatrix} \begin{bmatrix} \Delta y_{1,t-1} \\ \Delta y_{2,t-1} \end{bmatrix} + \begin{bmatrix} e_{1,t} \\ e_{2,t} \end{bmatrix} \quad (25)$$

4. Exogenous Variables in VAR and VEC Models

In the attempt of improving both the fit of models to data and their predictive power, different crypto predictors are included in multivariate models. The theoretical justification for the use of these exogenous variables is that correlations between cryptocurrencies and these assets might exist, since today, cryptocurrencies are mainly considered as investment assets. The predictors included in the autoregressive models are those suggested by Catania et al. (2019) as best indicators of market sentiments: international stock index prices, commodity prices, US Government Bonds interest rates, and the market volatility index (VIX).

While several studies investigated the relations between cryptocurrencies and other financial assets, currently there are no empirical evidences that there exist good predictors for cryptocurrencies (Muglia, Santabarbara and Grassi, "Is Bitcoin a Relevant Predictor of Standard & Poor's 500?" Journal of Risk and Financial Management.) (Stavroyiannis and Babalos, "Dynamic Properties of the Bitcoin and the US Market", SSRN Electronic Journal).

Similarly to the work done by Bothe and Rossini (Rick Bothe and Luca Rossini, “Comparing the forecasting of cryptocurrencies by Bayesian time-varying volatility models”, Vrije Universiteit Amsterdam, The Netherlands, 2019), this thesis investigates whether the main cryptocurrencies in the market are linked together and if they are correlated to other important macroeconomic variables. The predictors included in the multivariate models considered in this piece of work are:

- S&P 500 daily log-returns: the most famous and one of the major US stock indices;
- Nikkei 225 daily log-returns: the Japanese stock index, since the Japanese exchange is the second largest in the world;
- Stoxx 600 daily log-returns: the largest index representing European stock market;
- Gold and Silver daily log-returns: as they are the main investment commodities.
- US 3 months Treasury Bond yield;
- US 10 years Treasury Bond yield;
- VIX daily log-returns: the index measuring the volatility in the US stock market.

Cryptocurrency markets are open every day, so there are no breaks in the data corresponding to weekends or holidays. On the other hand, the markets where the predictors assets are exchanged close on weekends. For this reason, the timeseries considered in the analysis have not the same length and must be adjusted to implement the multivariate models.

A simple correction is made, the missing values in the crypto predictor series are substituted with the closing values in the last day of trading. In this way, a return of zero is generated in the missing days. This is the best way to make the models work without affecting the estimations too much. After daily prices are gathered for the period analysed, it is necessary to transform the prices into log-returns in order to add them to autoregressive models.

Both in a VAR (n,P) model and in a VEC (n,r,q) model, the predictors are represented by the exogenous variable x_t. VAR and VEC models including exogenous predictors are called respectively VARX and VECX from now on.

5. Residuals Analysis

The easiest check consists in plotting the correlograms and Quantile-Quantile plots of residuals, to assess the goodness of fit to data and to determine whether the assumptions on their distributions are correct. While this is a widely used approach in both univariate and multivariate framework, it is not completely appropriate when working with multivariate models. The reason they are used is that there exist only few tests to analyse multivariate models.

One could also rely on univariate tests applied to univariate models’ residuals. Although these tests can be applied to marginal residuals of multivariate models, this approach does not consider contemporaneous correlation of disturbances. Statistics from individual equations are not independent, thus even if the marginal residuals analysis might provide good results, the multivariate residuals analysis might not. (Massimo Guidolin, “Modelling and Forecasting Conditional Covariances: DCC and Multivariate GARCH”, Dept. of Finance, Bocconi University).

A proper multivariate residuals analysis can be implemented running the multivariate version of the Ljung-Box test on the multivariate standardized residuals, and the McLeod-Li test on the multivariate squared standardized residuals of VARs, VEC, VARX and VECX models. The multivariate

Portmanteau statistic (Ljung-Box Q statistic) tests the presence of autocorrelation in residuals up to lag h in a stable VAR (P) model. It is defined as:

$$Q_h = T\sum_{j=1}^{h} tr(\hat{C}'_j \hat{C}_0^{-1}\, \hat{C}_j \hat{C}_0^{-1}), \qquad (26)$$

Where $\hat{C}_i = \frac{1}{T}\sum_{t=i+1}^{T} \hat{\boldsymbol{u}}_t\, \hat{\boldsymbol{u}}'_t$. The test statistic is approximately distributed as $X^2_{(K2(h-p))}$. For what concerns the analysis of multivariate standardized residuals distributions, the test run to check the correctness of the distribution assumption of the errors is the Baringhaus-Franz multivariate test. It is an in-sample test which basically checks whether two datasets, A and B, are identically distributed or not. The test null hypothesis (H_0) is that sample A is distributed as sample B. The test statistic is:

$$T = \frac{mn}{m+n}\left\{\frac{1}{mn}\sum_{i=1}^{m}\sum_{j=1}^{n}\left\|X_i - Y_j\right\| - \frac{1}{2m^2}\sum_{i=1}^{m}\sum_{j=1}^{m}\left\|X_i - X_j\right\| - \frac{1}{2n^2}\sum_{i=1}^{n}\sum_{j=1}^{n}\left\|Y_i - Y_j\right\|\right\} \qquad (27)$$

where || · || is the Euclidian distance. The test is applied to residuals of multivariate GARCH models. $X_1, ..., X_m$ represent vectors of estimated residuals, while $Y_1, ..., Y_n$ are vectors of samples from a given theoretical distribution, it could be a multivariate Gaussian or a multivariate Student's t. The critical point is obtained by bootstrapping this statistic with a 95% confidence level. If the estimated statistic falls inside the confidence interval, the null hypothesis that the DCC-GARCH model explains the data well is not rejected (Elisabeth Orskaug, "Multivariate DCC-GARCH Model -With Various Error Distributions", NTNU Norwegian University of Science and Technology, June 2009).

6. Multivariate GARCH Models

While the VAR and VEC are multivariate models which aim at modelling linear interdependences among a set of response variables, multivariate GARCHs aim at modelling their conditional volatilities and correlations.

Multivariate GARCH models were introduced in the 1980s and are an extension of univariate GARCHs. They have become more and more important since then, because they are often used to model financial data. The most famous multivariate GARCH model, the one used in this thesis, is the DCC GARCH model (Dynamic Conditional Correlation GARCH).

Multivariate GARCH implementation have several major problems. Firstly, the so called "curse of dimensionality". A multivariate GARCH applied to *n* timeseries imply the estimation of $n(n+1)/2$ conditional variances and covariances, which change with passing of time. Secondly, all these variances and covariances are not observed in a direct way. Moreover, multivariate variance covariance matrix must always be positive definite, exactly like in the case of univariate models.

A financial timeseries can be decomposed as

$$y_t = \mu_t + a_t \qquad (28)$$

Where $\mu_t = E(z_t|F_{t-1})$ is the conditional expected value of y_t given F_{t-1}, the information available at time t-1. a_t represents instead the innovation, the part which is not autocorrelated. It can be written as

$$a_t = \Sigma_t \varepsilon_t \qquad (29)$$

Where ε_t is a sequence of casual iid vectors, such that $E(\varepsilon_t) = 0$, and Σ_t is the conditional covariance matrix of returns. Commonly used distribution for ε_t are the multivariate Student T or the multivariate normal.

DCC GARCH models

DCC GARCH models are used to describe the dynamic dependencies of volatility $\Sigma_t = [\sigma_{ij,t}]$. It is the volatility matrix of a_t given F_{t-1}. The conditional correlation matrix is then $\rho_t = D_t^{-1}\Sigma_t D_t^{-1}$ where D_t is the diagonal matrix containing n volatilities at time t modelled by univariate GARCH processes. Elements contained in D_t can be estimated by different GARCH (P,Q) processes, with different number of lags. It is sufficient that the GARCHs satisfy the conditions of stationarity and positiveness.

The idea behind DCC models is to separate the process in two different steps. In the first step the volatilities $\{\sigma_{ii,t}\}$ are computed, for i = 1, ..., n. Next, the dynamic dependencies among correlation matrices ρ_t are modelled. If $\eta_t = (\eta_{1t}, \dots, \eta_{nt})$ is the vector of standardized marginal innovations, where $\eta_{it} = \hat{a}_{it}/\sqrt{\sigma_{ii,t}}$, ρ_t is then the volatility matrix of η_t. The DCC model is defined as

$$\widehat{Q}_t = (1 - a - b)\overline{Q} + a\widehat{Q}_{t-1} + b\eta_{t-1}\eta'_{t-1} \qquad (30)$$

$$\rho_t = J_t Q_t J_t \qquad (31)$$

Where $\widehat{Q}_t$ is the estimated conditional covariance matrix, $\overline{Q}$ is the unconditional covariance matrix of η_t, and J_t is a normalization matrix, a and b are non-negative real numbers which satisfy the condition $0 < a + b < 1$. a e b in the first equation control the dynamic dependencies of correlations, independently on the number of timeseries n in the model. This makes the DCC GARCH very parsimonious, allowing the parameters estimation to be carried out easily. On the other hand, the small number of parameters in the model represents a weakness, because this implies that all the correlations should develop in the same way, which is difficult to justify from a theoretical point of view.

Since the estimation of parameters is divided in two steps, the log-Likelihood equation of the DCC-GARCH is separated into two components, a volatility component which is used in the first estimation step and a correlation component which produces estimates in the second step. The log-Likelihood equation can be expressed as:

$$LL = \tfrac{1}{2}\textstyle\sum_{i=1}^{T}(Nlog(2\pi) + 2log|D_t| + log|\rho_t| + \eta_t'\rho_t^{-1}\eta_t') \qquad (32)$$

$$= \frac{1}{2}\sum_{i=1}^{T}(Nlog(2\pi) + 2log|D_t| + \varepsilon_t' D_t^{-1} D_t^{-1}\varepsilon_t) - \frac{1}{2}\sum_{i=1}^{T}(\eta_t'\eta_t + log|\rho_t| + \eta_t'\rho_t^{-1}\eta_t')$$

$$= LL_V(\theta_1) + LL_R(\theta_1, \theta_2) \qquad (33)$$

where $LL_V(\theta_1)$ is the volatility component with parameters θ_1, and $LL_R(\theta_1, \theta_2)$ the correlation component with parameters θ_1 and θ_2 (Alexios Ghalanos, "The rmgarch models: Background and properties", Version 1.3-0, September 12, 2019).

This is important when the log-Likelihood is used for model selection purposes. The Likelihood Ratio is not reliable in this case because of the two steps estimation. Other selection criteria must be preferred. In particular, one could rely on Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC) as they are appropriate comparison methods for non-nested models. In brief, the steps required to build up a DCC GARCH model are:

- Estimating the conditional mean μ_t of timeseries through apposite models, such as a VAR (n,P) or a VEC (n,r,q);
- Estimating the residuals from the previous model. They are defined as $\hat{a}_t = z_t - \hat{\mu}_t$.
- Fitting univariate GARCH models to $\hat{a}_{it}$, obtaining the estimates of conditional volatilities $\hat{\sigma}_{ii,t}$.
- Standardizing the innovations as $\hat{\eta}_{it} = \hat{a}_{it}/\sqrt{\hat{\sigma}_{ii,t}}$ to fit a DCC model to $\hat{\eta}_t$.

The conditional distribution of η_t can be either standard multivariate normal or standard multivariate Student t with v degree of freedom. (Dedalo Invest Finance, "Modelli GARCH univariati e multivariati").

In this work, several DCC GARCHs are used to model the residuals of VARs and VEC models. Different flavours of univariate GARCHs are evaluated, vanilla GARCH (1,1), GARCH (1,2) with two lags of conditional variance, the GARCH-GJR already described in the univariate analysis at previous Chapters, GARCH-t with the innovations following a Student-t distribution.

The DCC equations considered are specified with 1 lag of symmetric innovation and 1 lag of correlation and, both a multivariate Normal distribution and a multivariate Student-t distribution are assumed for the standardized innovations in the DCC equation. The features of univariate GARCHs and DCC just described are combined in many DCC GARCH models, in order to find the best combination to fit the VARs and VEC residuals.

7. Forecasting

One could easily obtain daily forecasts of log-returns from multivariate models in an analogous way as for univariate models. The example of forecasting with a VAR (2,1) and a DCC (1,1) – N, GARCH (1,1)(1,1) – N is illustrated. The VAR model estimating the log-returns observation at time t+1 can be written as

$$\begin{bmatrix} \hat{y}_{1,t+1} \\ \hat{y}_{2,t+1} \end{bmatrix} = \begin{bmatrix} \hat{c}_1 \\ \hat{c}_2 \end{bmatrix} + \begin{bmatrix} \hat{a}_{1,1} & \hat{a}_{1,2} \\ \hat{a}_{2,1} & \hat{a}_{2,2} \end{bmatrix} \begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix}$$

While for the DCC GARCH model, the following equations estimate the future conditional volatilities

$$\hat{a}_{t+1} = \hat{\Sigma}_{t+1}\hat{\varepsilon}_{t+1} \qquad \hat{\varepsilon}_{t+1} \sim N(0, \hat{\Sigma}_{t+1}),$$

where $\hat{\Sigma}_t = [\hat{\sigma}_{ij,t}]$ is the estimated volatility matrix of $\hat{a}_{t+1}$ given F_t, the set of information available at time t.

Since the conditional correlation matrix is defined as $\rho_t = D_t^{-1}\Sigma_t D_t^{-1}$, $\hat{\Sigma}_t$ is obtained easily as $\hat{\rho}_t$ and elements contained in $\hat{D}_t$ are estimated. In the first step the volatilities $\{\hat{\sigma}_{ii,t}\}$ are estimated, for i = 1, ..., n. Next, the dynamic dependencies among correlation matrices are modelled as

$$\hat{Q}_{t+1} = (1 - \hat{a} - \hat{b})\bar{Q} + \hat{a}\hat{Q}_t + \hat{b}\hat{\eta}_t\hat{\eta}'_t$$

$$\hat{\rho}_{t+1} = \hat{J}_{t+1}\hat{Q}_{t+1}\hat{J}_{t+1}$$

where $\hat{\eta}_t = (\hat{\eta}_{1t}, \dots, \hat{\eta}_{nt})$, is the vector of standardized marginal innovations and each element is $\hat{\eta}_{it} = \sqrt{\hat{\sigma}_{ii,t}}$. Once both the VAR or VEC models and the DCC GARCH are estimated, one can get the estimated daily log returns as

$$\hat{y}_{t+1} = \hat{\mu}_{t+1} + \hat{a}_{t+1}$$

Further explanations on DCC GARCH models can be found in the previous section. The best way to obtain forecasts through the procedure just described would be implementing a rolling forecast procedure. Though, the same problem already experienced for univariate models arises in the multivariate framework. The out of sample timeseries are too short to compute meaningful performances indicators.

To determine whether the models are good for forecasting, the predictive log-Likelihoods are computed. In practice, they are obtained filtering the out of sample data with the models estimated on full data (historical + out of sample data). These are different from the predictive marginal Likelihoods. The predictive log-Likelihood cannot be interpreted directly, because it does not weight the number of parameters in the models like the marginal Likelihood does. Moreover, the log-Likelihoods obtained are not strictly increasing with richer and richer specifications. This is because the predictive log-Likelihood is not the maximum log-Likelihood given a set of data, it represents the Likelihood of observing the out of sample data with a model whose specification is obtained through estimation on a different sample. Though, the Likelihood Ratio comparison method used for model selection step cannot be used and information criteria must be preferred.

VI VAR and VEC Models Implementation and Results

1. Cointegration Analysis

In order to determine which multivariate framework is more convenient to model cryptocurrencies timeseries it is necessary to assess whether there is cointegration among them.

The first step in the multivariate analysis is therefore to run cointegration tests. One step carried out while fitting the univariate models, was to assess that all cryptocurrencies series are integrated of order 1. Therefore, it is possible to run both the Engle-Granger test and the Johansen test to search for cointegration relations.

Cointegration tests are run on log-prices. The advantage of working with logarithmic transformation of prices is that this allows to reason in terms of elasticities. In a multivariate environment, for a percentage variation in a cryptocurrency it becomes possible to estimate the percentage change of another. Log transformed series are still I(1), like raw data. This is confirmed by results of Augmented Dickey Fuller tests collected in Table 26.

Bitcoin log-prices	p-value	Test result
AR	0.9899	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift	0.5928	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift and trend	0.9109	p-value > 0.05 critical value, do not reject the null hypothesis
Ethereum log-prices		
AR	0.8973	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift	0.3370	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift and trend	0.9523	p-value > 0.05 critical value, do not reject the null hypothesis
Litecoin log-prices		
AR	0.8638	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift	0.7227	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift and trend	0.9160	p-value > 0.05 critical value, do not reject the null hypothesis
Ripple log-prices		
AR	0.1151	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift	0.7405	p-value > 0.05 critical value, do not reject the null hypothesis
AR with drift and trend	0.9426	p-value > 0.05 critical value, do not reject the null hypothesis

Table 26: Augmented Dickey Fuller test results on cryptocurrencies log-prices series.

Results of Engle-Granger Cointegration tests are collected in table 27, while results of Johansen tests are reported in table 28. The Johansen test is run for different estimated VEC models with lag from 1 to 4. For each model the test is run to determine whether there are 0 or 1 cointegrating relations (VEC model with rank 0 or rank 1). In order to conclude that a cointegration relation is detected it is first necessary to reject the hypothesis that the cointegration rank is 0. If the hypothesis cannot be rejected, then evaluating the results for rank 1 is not meaningful.

The unconditional Likelihood resulting from the tests is reported because it is the basis for the computation of the Likelihood ratio[4], used to determine the number of lags to consider in a VEC model, in case cointegration is detected.

[4] When a model nests inside another one, Likelihood ratio test can be used to assess if the parameters are statistically significantly different from 0. When a model is nested into another, it is strictly a subset of the other. Consider M_U and M_R, respectively the unrestricted and the restricted models, where it is assumed U is the number of parameters in the unrestricted model and R is the number of parameters in the restricted model. The restricted log-Likelihood minus the unrestricted Likelihood doubled, follows a chi square distribution with U-R degree of freedom (the number of restrictions). The Likelihood ratio can be written as:

$$\mathrm{LR} = 2(LU - LR) \sim X^2_{\text{(number of restrictions)}}$$

If LR statistic > $X^2_{\text{(number of restrictions)}}$ the unrestricted model fits better the data at a pre-determined significance level.

The test can be used alternatively to information criteria such as the Akaike Information Criteria or the Bayesian Information Criteria in model comparison (Jon Danielsson, “Financial Risk Forecasting”, John Wiley and Sons Ltd, 2011 edition).

Cryptocurrencies couples	p-value	Test result
Bitcoin-Ethereum	0.7127	p-value > 0.05 critical value, do not reject the null hypothesis of no cointegration
Bitcoin-Litecoin	0.2207	p-value > 0.05 critical value, do not reject the null hypothesis of no cointegration
Bitcoin-Ripple	0.3081	p-value > 0.05 critical value, do not reject the null hypothesis of no cointegration
Ethereum-Litecoin	0.2066	p-value > 0.05 critical value, do not reject the null hypothesis of no cointegration
Ethereum-Ripple	0.1698	p-value > 0.05 critical value, do not reject the null hypothesis of no cointegration
Litecoin-Ripple	0.0125	p-value < 0.05 critical value, reject the null hypothesis of no cointegration

Table 27: Engle-Granger test results on sets of cryptocurrencies pairs.

Cryptocurrencies couples	Model lags	Rank 0 p-value	Rank 1 p-value	Unrestricted Likelihood	Test result
Bitcoin-Ethereum	1 lag	0.7137	0.0965	5,323.7	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	2 lags	0.7179	0.0939	5,321.5	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	3 lags	0.6855	0.1006	5,330.5	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	4 lags	0.7224	0.0996	5,332.6	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
Bitcoin-Litecoin	1 lag	0.4540	0.3124	5,715.8	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	2 lags	0.4513	0.3137	5,720.4	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	3 lags	0.4362	0.3065	5,721.0	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	4 lags	0.4355	0.2655	5,723.8	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
Bitcoin-Ripple	1 lag	0.4733	0.1830	5,093.9	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	2 lags	0.4207	0.1737	5,104.5	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration

	3 lags	0.4211	0.1704	5,111.0	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	4 lags	0.4225	0.1659	5,111.8	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
Ethereum-Litecoin	1 lag	0.2067	0.0645	4,652.1	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	2 lags	0.2200	0.0718	4,658.8	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	3 lags	0.2078	0.0865	4,664.2	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	4 lags	0.1990	0.0865	4,665.6	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
Ethereum-Ripple	1 lag	0.1690	0.0693	4,227.1	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	2 lags	0.1732	0.0738	4,242.5	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	3 lags	0.1848	0.0863	4,250.8	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	4 lags	0.2104	0.0841	4,251.9	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
Litecoin-Ripple	1 lag	0.0477	0.4545	4,498.4	Rank 0 p-value < 0.05 critical value, reject the null hypothesis of 0 cointegrating relations, there is evidence of rank 1 cointegration
	2 lags	0.0394	0.4496	4,509.0	Rank 0 p-value < 0.05 critical value, reject the null hypothesis of 0 cointegrating relations, there is evidence of rank 1 cointegration
	3 lags	0.1937	0.4153	4,552.1	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration
	4 lags	0.1565	0.4112	4,554.3	Rank 0 p-value > 0.05 critical value, there are no evidences of cointegration

Table 28: Johansen test results on sets of cryptocurrencies pairs.

The results from Engle-Granger test and Johansen test lead to the same conclusion, there is some evidence of a cointegrating relation between Litecoin and Ripple. The Johansen test, though, is more complete, because it also indicates that a VEC model with 2 lags would be the best fitting model for the multivariate log-prices series.

The procedure used to select the best models with Likelihood Ratio is as follow. Starting from the more parsimonious model, the VEC (2,1,1), it is tested whether the addition of one more lag produces significant improvements in the log-Likelihood. If the VEC (2,1,1) model is favoured by the Likelihood Ratio test the procedure stops and it is the best model to fit the data. If instead, the VEC (2,1,2) is favoured, the procedure must continue because one needs to assess whether the VEC (2,1,3) is better (in this specific case this would make no sense, since no cointegrations are detected by a VEC(2,1,3) on Litecoin-Ripple series). To sum up, in general each VEC (n,r,q) or VAR (n,P) model is tested only against the VEC (n,r,q-1) or VAR (n,P-1), the model with one less lag. In this case, the VEC (2,1,2) is tested against the restricted VEC (2,1,1) model. The LR test statistic is computed as

$$\mathrm{LR} = 2(4{,}509.0 - 4{,}498.4) \sim X^2{}_{(4)}$$

$$\mathrm{LR} = 21.2 \sim X^2{}_{(4)}$$

Since the statistic is larger than the critical value (21.2 > 13.277), the test suggests that the unrestricted model fits better the data and by adding one lag the fit improves significantly.

For all other couples of timeseries, there is no evidence of cointegration. While a VEC model cannot be estimated in these cases, it is still possible to run VAR models on the differenced series.

2. VAR Models Selection

In order to obtain the best fitting to the data, several VAR models are fitted to the couples of time series by using Matlab software. While AIC, BIC and log Likelihood are reported in the tables, the latter is used to compute the Likelihood Ratio, which is used to select the best model for each pair.

Computations are made exactly like in the VEC model selection example and they are replicated for all multivariate log-returns series. The 1% quantile of the Chi Squared distribution used as critical value in the in the Likelihood Ratio test is always 13.277, taken from a distribution with 4 degrees of freedom, equal to the number of restrictions in the test. The best result is highlighted in the table for each comparison criterium.

Bitcoin-Ethereum	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-10,619.4	-10,587	5,315.71	6
VAR (2,2)	-10,612.8	-10,558.8	5,316.41	10
VAR (2,3)	-10,622.3	-10,546.7	5,325.15	14
VAR (2,4)	-10,617.7	-10,520.5	5,326.86	18
VAR (2,5)	-10,615.2	-10,496.4	5,329.58	22
VAR (2,6)	-10,610.9	-10,470.6	5,331.47	26
VAR (2,7)	-10,612.8	-10,450.8	5,336.39	30
VAR (2,8)	-10,606.0	-10,422.4	5,336.98	34

U-R	LR Critical value (α = 1%)
4	13.277

Table 29: Bitcoin-Ethereum VAR models selection criteria.

Bitcoin-Litecoin	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-11,390.6	-11,358.2	5,701.28	6
VAR (2,2)	-11,383.1	-11,329.1	5,701.56	10
VAR (2,3)	-11,376.1	-11,300.5	5,702.04	14
VAR (2,4)	-11,373.7	-11,276.5	5,704.86	18
VAR (2,5)	-11,379.5	-11,260.8	5,711.76	22
VAR (2,6)	-11,386.6	-11,246.3	5,719.32	26
VAR (2,7)	-11,381.5	-11,219.5	5,720.74	30
VAR (2,8)	-11,376.0	-11,192.5	5,722.01	34

Table 30: Bitcoin-Litecoin VAR models selection criteria.

Bitcoin-Ripple	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-10136.6	-10104.2	5074.29	6
VAR (2,2)	-10149.1	-10095.1	5084.57	10
VAR (2,3)	-10154.1	-10078.5	5091.04	14
VAR (2,4)	-10147.5	-10050.4	5091.77	18
VAR (2,5)	-10148.1	-10029.4	5096.06	22
VAR (2,6)	-10144.9	-10004.6	5098.46	26
VAR (2,7)	-10137.7	-9975.71	5098.84	30
VAR (2,8)	-10140.1	-9956.56	5104.06	34

Table 31: Bitcoin-Ripple VAR models selection criteria.

Ethereum-Litecoin	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-9283.86	-9251.47	4647.93	6
VAR (2,2)	-9277.97	-9223.98	4648.98	10
VAR (2,3)	-9285.56	-9209.97	4656.78	14
VAR (2,4)	-9280.58	-9183.41	4658.29	18
VAR (2,5)	-9278.01	-9159.23	4661.00	22
VAR (2,6)	-9283.1	-9142.73	4667.55	26
VAR (2,7)	-9286.55	-9124.59	4673.27	30
VAR (2,8)	-9280.46	-9096.9	4674.23	34

Table 32: Ethereum-Litecoin VAR models selection criteria.

Ethereum-Ripple	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-8433.99	-8401.6	4223.00	6
VAR (2,2)	-8444.98	-8390.99	4232.49	10
VAR (2,3)	-8457.99	-8382.41	4243.00	14
VAR (2,4)	-8452.06	-8354.88	4244.03	18
VAR (2,5)	-8449.39	-8330.62	4246.70	22

VAR (2,6)	-8455.96	-8315.59	4253.98	26
VAR (2,7)	-8452.78	-8290.82	4256.39	30
VAR (2,8)	-8457.34	-8273.78	4262.67	34

Table 33: Ethereum-Ripple VAR models selection criteria.

Different criteria might provide different results. One could notice how the BIC criteria penalizes model complexity with respect to AIC. For this analysis the Likelihood Ratio results are considered to select the best VAR models, as it is the most precise selection criterion for nested models.

Comparison criteria do not target minimizing the amount of autocorrelation in residuals of models fitted to data, so the problem of possible serial correlation in the residuals of VAR and VEC models remains not addressed. It is evaluated in a dedicated section.

3. VAR Models Estimations

In the next tables are gathered the Matlab software estimations for the VAR models parameters.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0023394	0.00097007	2.4115	0.015886
Constant (2)	0.0032949	0.0016109	2.0453	0.040822
AR{1} (1,1)	0.011522	0.027774	0.41483	0.67827
AR{1} (2,1)	0.017125	0.046123	0.37128	0.71043
AR{1} (1,2)	-0.037511	0.016718	-2.2437	0.024849
AR{1} (2,2)	-0.073112	0.027763	-2.6335	0.0084521

Table 34: Bitcoin-Ethereum VAR (2,1) estimation results.

It is important to recall that VAR models are simply multivariate form of Ordinary Least Square estimations. Thus, the coefficients must be interpreted *ceteris paribus*. Since the VAR models are run on log-returns, the estimated coefficients can be interpreted as elasticities. In this case there is only one AR matrix because a VAR with only one lag was fitted to the data.

The table above gathers the estimates for the Bitcoin-Ethereum pair VAR (2,1). The AR {1} matrix elements represents the short run relations that the response variables have with their own lagged values and other variables lagged values. The coefficients in the position (1,1) of the Auto Regressive matrix, for example can be interpreted as the percentage change in Bitcoin returns given 1% change in Bitcoin lagged returns. The element in position (1,2) instead, belongs to the same equation but it indicates that 1% change in lagged value of Ethereum returns, on average, everything else held constant, leads to a -0.037511% change in Bitcoin returns. Coefficients stored in AR {1} (2,1) and AR {1} (2,2) weights lagged Bitcoin returns and lagged Ethereum returns in the second equation of the system. In other words, they define their impact on Ethereum returns at time t.

The two constants in the system are statistically significant at 0.05 significance level but they do not have an immediate interpretation. As the VAR is built with log returns response variables, one could interpret the constants by taking their exponential, but this would not be correct, since the average of the log returns is not the logarithm of their average.

More interestingly, the coefficients in AR {1} (1,2) and in AR {1} (2,2) are statistically significant. The former can be interpreted as written few lines above, while the latter meaning is that on average, 1% change in lagged value of Ethereum returns, lead to a -0.073112% change on its own returns at time t.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0022821	0.00097115	2.3499	0.018778
Constant (2)	0.0017917	0.001434	1.2495	0.21148
AR{1} (1,1)	-0.0036997	0.031372	-0.11793	0.90612
AR{1} (2,1)	0.042889	0.046323	0.92586	0.35452
AR{1} (1,2)	-0.01466	0.021238	-0.69029	0.49001
AR{1} (2,2)	-0.058101	0.031359	-1.8528	0.063918

Table 35: Bitcoin-Litecoin VAR (2,1) estimation results.

The VAR model estimated for Bitcoin and Litecoin returns timeseries does not provide good estimations. Only the constant in the Bitcoin equation is significant, indicating that there are no clues that the two cryptocurrencies returns are affected by their own past values or by the other currency past returns.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0022246	0.00096992	2.2936	0.021812
Constant (2)	0.0022462	0.0017644	1.2731	0.20299
AR{1} (1,1)	-0.0049457	0.026326	-0.18787	0.85098
AR{1} (2,1)	-0.046561	0.04789	-0.97225	0.33093
AR{1} (1,2)	-0.014984	0.014437	-1.0379	0.29931
AR{1} (2,2)	-0.056296	0.026263	-2.1436	0.032067
AR{2} (1,1)	-0.011454	0.026332	-0.43496	0.66359
AR{2} (2,1)	-0.030656	0.047902	-0.63998	0.52218
AR{2} (1,2)	0.042393	0.014428	2.9382	0.0033016

AR{2} (2,2)	0.10894	0.026247	4.1507	3.314e-05

Table 36: Bitcoin-Ripple VAR (2,2) estimation results.

By fitting a VAR (2,2) to Bitcoin and Ripple returns timeseries, more significant parameters were estimated. In this case the results suggest that the lagged returns of Ripple at time t-1 and t-2 have a significant influence on its own returns at time t. 1% increase in Ripple returns at time t-1 leads, on average, to -0.056296% variation in its returns at time t, while 1% increase in Ripple returns at time t-2 leads to 0.10894% increase in returns at time t.

Also the parameter AR{2} (1,2) is strongly significant, meaning that 1% increase in Ripple returns at time t-1 leads on average, ceteris paribus, to 0.042393% increase in Bitcoin returns.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0032885	0.0016086	2.0444	0.040919
Constant (2)	0.0019241	0.0014331	1.3426	0.1794
AR{1} (1,1)	-0.087025	0.02721	-3.1983	0.0013824
AR{1} (2,1)	-0.030168	0.024241	-1.2445	0.2133
AR{1} (1,2)	0.049573	0.030602	1.6199	0.36914
AR{1} (2,2)	-0.02588	0.027263	-0.94928	0.34248

Table 37: Ethereum-Litecoin VAR (2,1) estimation results.

Next, the estimated parameters in the VAR (2,1) for the Ethereum-Litecoin pair, suggest that Ethereum returns at time t are significantly influenced by their own values at time t-1, since the p-value is lower than 0.05. An increase of 1% in Ethereum returns at time t-1 is associated with -0.087025% in Ethereum returns at time t.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.002985	0.001608	1.8563	0.063405
Constant (2)	0.0018124	0.0017593	1.0302	0.30292
AR{1} (1,1)	-0.072248	0.025926	-2.7867	0.0053251
AR{1} (2,1)	0.032203	0.028366	1.1353	0.25626
AR{1} (1,2)	0.010863	0.023699	0.45837	0.64668
AR{1} (2,2)	-0.082373	0.025929	-3.1768	0.0014888
AR{2} (1,1)	0.029551	0.025999	1.1366	0.25569

AR{2} (2,1)	0.0066792	0.028445	0.23481	0.81436
AR{2} (1,2)	0.0012981	0.02364	0.054911	0.95621
AR{2} (2,2)	0.10587	0.025865	4.0934	4.2518e-05
AR{3} (1,1)	0.082353	0.025787	3.1936	0.0014049
AR{3} (2,1)	0.015179	0.028213	0.53801	0.59057
AR{3} (1,2)	-0.0029669	0.023674	-0.12532	0.90027
AR{3} (2,2)	0.076192	0.025902	2.9416	0.0032654

Table 38: Ethereum-Ripple VAR (2,3) estimation results.

The VAR suggested by the Likelihood Ratio to fit the pair Ethereum-Ripple is the VAR (2,3). In this case there are many significant parameters explaining the two cryptocurrencies returns at time t. There are evidences that Ethereum returns at time t are significantly influenced by their own values at time t-1 and time t-3. Ripple returns at time t, instead seems to be affected by their own past values at time t-1, t-2 and t-3.

1% variation in Ethereum returns at time t-1 leads to an average change in the Ethereum returns at time t of -0.072248%. For 1% variation in Ethereum returns at time t-3, the Ethereum returns at time t change on average of 0.082353%. 1% increase in Ripple returns at time t-1 leads to an average change in its own returns at time t of -0.082373%, the same variation in Ripple returns at time t-2 leads to an average change in the Ripple returns at time t of 0.10587% and 1% increase in Ripple returns at time t-3 leads to an average change in its own returns at time t of 0.076192%.

It is important to notice that in each model estimated, the returns of a cryptocurrency at time t is never affected by the returns of the other currency in the system of equations.

4. VEC Models Estimations

Both the Engle-Granger test and the Johansen test indicates the presence of a cointegrating relation between Litecoin and Ripple log-prices. A VEC model must be preferred over a VAR in this case, because it considers the long-run relations detected by the tests.

Next table gathers the Matlab software outputs with the estimation results for the VEC model. The model is specified according to lags and ranks suggested by the Johansen test in the first section of this Chapter.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.017185	0.022559	0.76177	0.4462
Constant (2)	-0.086911	0.027596	-3.1494	0.001636

Adjustment (1,1)	0.00097629	0.0014314	0.68205	0.49521
Adjustment (2,1)	-0.0056553	0.001751	-3.2297	0.0012394
Impact (1,1)	-0.0030355	0.0044505	-0.68205	0.49521
Impact (2,1)	0.017583	0.0054443	3.2297	0.0012394
Impact (1,2)	0.0023098	0.0033866	0.68205	0.49521
Impact (2,2)	-0.01338	0.0041429	-3.2297	0.0012394
ShortRun{1}(1,1)	-0.04781	0.027198	-1.7579	0.07877
ShortRun{1}(2,1)	0.059409	0.033271	1.7856	0.074164
ShortRun{1}(1,2)	0.019037	0.021954	0.8671	0.38589
ShortRun{1}(2,2)	-0.082626	0.026857	-3.0765	0.0020943
ShortRun{2}(1,1)	0.0056867	0.027167	0.20932	0.8342
ShortRun{2}(2,1)	-0.01077	0.033234	-0.32408	0.74588
ShortRun{2}(1,2)	-0.0040846	0.021919	-0.18635	0.85217
ShortRun{2}(2,2)	0.10519	0.026814	3.9232	8.739e-05

Cointegration Matrix (Estimated):

-3.1092
2.3659

Cointegration Constant:

15.4328

Table 39: Litecoin-Ripple VEC (2,1,2) estimation results.

The estimation output of the VEC model is richer than those of VARs, since it provides estimations for both parameters involved in modelling the long run and the short run relations among response variables. The estimated coefficients contained in the short run matrices can be interpreted like VAR parameters, since they represent the short run relation between the response variables. The cointegration constant and the impact matrix are, instead, the parameters in the long run equation, where the Impact matrix (C) is just the multiplication of Adjustment matrix (A) by Cointegration matrix (B) transposed (C = AB').

The estimated parameters must be interpreted ceteris paribus, like in the VAR models. In the short run equation, there are two significant estimates at 0.05 level. The model suggests that for 1% increase in Ripple returns at time t-1, its own returns at time t change on average of -0.082626%

and 1% increase in Ripple returns at time t-2, is associated on average with 0.10519% variation in Ripple returns at time t.

The other significant parameter is the adjustment coefficient of Ripple returns. It is negative and significant meaning that the process converges towards long run equilibrium. In particular, the model predicts that about 0.005% of departure of Ripple price from long run equilibrium is corrected each day. It is possible to conclude nothing in the case of Litecoin long-run equilibrium, as the adjustment coefficient is not significant.

5. Exogenous Variables

Like in the case of VAR and VEC models, different VARXs and VECXs are fitted to each pair of cryptocurrencies by using Matlab software. Again, the model selection criterium used in this work is the Likelihood Ratio. Despite this, also the AIC and BIC criteria are reported in the following table to make comparisons. The best models must be identified again because including the predictors in the models increases the informative set on which they are estimated, and this might change the dynamic specification of the model. In other words, if previously a VAR (2,1) resulted the best model for two cryptocurrencies returns series, this does not mean that the VARX (2,1) would be the best, once predictors are included.

The predictors daily prices series are transformed into log-returns series, in order to ensure stationarity, the same transformation applied to crypto-currencies series used in the VAR models. For interest rates predictors, the simple returns are considered, because sometimes they take negative values and computing their logarithm is impossible. The interpretation of their coefficient changes with respect to that of other exogenous variables in the model. They must be interpreted as semi-elasticities in the case they are statistically significant.

Computations to determine the LR are made like in the VARs and VECs case. The critical value used is 13.277 like before, since the number of restrictions between VARX (n,P) and VARX (n,P-1) are always 4. The best results are highlighted in the table for each comparison criterium.

Bitcoin-Ethereum	AIC	BIC	Log Likelihood	Number of parameters
VARX (2,1)	-10604.6	-10485.8	5324.29	22
VARX (2,2)	-10598.1	-10457.8	5325.07	26
VARX (2,3)	-10607.6	-10445.6	5333.8	30
VARX (2,4)	-10613.3	-10451.3	5336.63	34
VARX (2,5)	-10600.6	-10395.4	5338.29	38
VARX (2,6)	-10596.6	-10369.8	5340.28	42
VARX (2,7)	-10599.1	-10350.7	5345.54	46
VARX (2,8)	-10592.1	-10322.2	5346.06	50

U-R	LR Critical value ($\alpha = 1\%$)
4	13.277

Table 40: Bitcoin-Ethereum VARX models selection criteria.

Bitcoin-Litecoin	AIC	BIC	Log Likelihood	Number of parameters
VARX (2,1)	-11380.3	-11261.6	5712.16	22

VARX (2,2)	-11373.2	-11232.8	5712.59	26
VARX (2,3)	-11365.9	-11204	5712.96	30
VARX (2,4)	-11363.7	-11180.1	5715.83	34
VARX (2,5)	-11368	-11162.9	5722.01	38
VARX (2,6)	-11375.4	-11148.6	5729.7	42
VARX (2,7)	-11370.2	-11121.8	5731.09	46
VARX (2,8)	-11364.8	-11094.9	5732.42	50

Table 41: Bitcoin-Litecoin VARX models selection criteria.

Bitcoin-Ripple	AIC	BIC	Log Likelihood	Number of parameters
VARX (2,1)	-10124.5	-10005.7	5084.25	22
VARX (2,2)	-10137.2	-9996.81	5094.59	26
VARX (2,3)	-10142.4	-9980.42	5101.19	30
VARX (2,4)	-10135.8	-9952.26	5101.91	34
VARX (2,5)	-10136.7	-9931.57	5106.36	38
VARX (2,6)	-10133.3	-9906.58	5108.66	42
VARX (2,7)	-10126	-9877.68	5109.01	46
VARX (2,8)	-10127.9	-9857.94	5113.94	50

Table 42: Bitcoin-Ripple VARX models selection criteria.

Ethereum-Litecoin	AIC	BIC	Log Likelihood	Number of parameters
VARX (2,1)	-9272.3	-9153.52	4658.15	22
VARX (2,2)	-9266.09	-9125.72	4659.05	26
VARX (2,3)	-9272.97	-9111.01	4666.49	30
VARX (2,4)	-9267.98	-9084.42	4667.99	34
VARX (2,5)	-9265.18	-9060.03	4670.59	38
VARX (2,6)	-9270.47	-9043.72	4677.24	42
VARX (2,7)	-9273.66	-9025.32	4682.83	46
VARX (2,8)	-9268.1	-8998.16	4684.05	50

Table 43: Ethereum-Litecoin VARX models selection criteria.

Ethereum-Ripple	AIC	BIC	Log Likelihood	Number of parameters
VARX (2,1)	-8419.91	-8301.14	4231.96	22
VARX (2,2)	-8431.13	-8290.76	4241.57	26
VARX (2,3)	-8443.49	-8281.53	4251.75	30
VARX (2,4)	-8437.85	-8254.3	4252.93	34
VARX (2,5)	-8435.1	-8229.94	4255.55	38
VARX (2,6)	-8442.07	-8215.32	4263.04	42
VARX (2,7)	-8439.29	-8190.95	4265.65	46
VARX (2,8)	-8443.36	-8173.42	4271.68	50

Table 44: Ethereum-Ripple VARX models selection criteria.

As before, the selection criteria suggest often different results but, in this step, those provided by the Likelihood Ratio are preferred. One could observe that including the exogenous predictors to VAR models, did not change the results of the Likelihood Ratio and thus, the number of lags included in VARX are the same of those in the models without exogenous variables.

In order to obtain a comparison between autoregressive models, with and without crypto predictors, it is necessary to analyse residuals of the fitted models. This step is run in the dedicated section.

Matlab software outputs provide the estimation results for the previously selected VARX models. Estimates are reported in the following tables.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	-0.038381	0.070481	-0.54456	0.58605
Constant (2)	0.012699	0.11715	0.1084	0.91368
AR{1} (1,1)	0.014239	0.027739	0.51334	0.60772
AR{1} (2,1)	0.017116	0.046106	0.37122	0.71048
AR{1} (1,2)	-0.039723	0.016719	-2.3758	0.017509
AR{1} (2,2)	-0.076264	0.02779	-2.7443	0.0060639
Beta (1,1)	-0.15556	0.24624	-0.63173	0.52756
Beta (2,1)	-0.25332	0.40929	-0.61892	0.53597
Beta (1,2)	-0.23225	0.10126	-2.2937	0.021807
Beta (2,2)	-0.17306	0.1683	-1.0282	0.30384
Beta (1,3)	0.10498	0.16647	0.63063	0.52828
Beta (2,3)	0.14588	0.2767	0.52723	0.59804
Beta (1,4)	0.27589	0.27477	1.0041	0.31534
Beta (2,4)	0.73238	0.45671	1.6036	0.1088
Beta (1,5)	0.019644	0.14873	0.13208	0.89492
Beta (2,5)	-0.17331	0.24721	-0.70108	0.48325
Beta (1,6)	-0.035825	0.0241	-1.4865	0.13714
Beta (2,6)	-0.069572	0.040057	-1.7368	0.082418
Beta (1,7)	0.041654	0.070609	0.58992	0.55525

Beta (2,7)	-0.0095741	0.11736	-0.081577	0.93498
Beta (1,8)	-0.00090917	0.0028424	-0.31986	0.74907
Beta (2,8)	0.00012722	0.0047245	0.026928	0.97852

Table 45: Bitcoin-Ethereum VARX (2,1) estimation results.

Exactly like in the case of VARs, the coefficients must be interpreted ceteris paribus. The coefficients of the predictors express the percentage change in the response variable given one percent change in the predictors, since also the exogenous predictors are transformed into log-returns.

The same coefficients that were significant in the VAR (2,1) model are significant also including the predictors. Moreover, their estimated values are very close to those in the VAR.

Among the coefficients of crypto-predictors, only Beta (1,2) element is statistically significant. It is the coefficient linked to Nikkei 225 Index lagged log-returns. On average, ceteris paribus, if the Nikkei 225 returns at time t-1 increase by 1%, the Bitcoin returns decreases by 0.23%.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	-0.030252	0.070518	-0.42899	0.66793
Constant (2)	0.045255	0.10412	0.43463	0.66383
AR{1} (1,1)	-0.0033461	0.031315	-0.10685	0.91491
AR{1} (2,1)	0.042267	0.046238	0.91413	0.36065
AR{1} (1,2)	-0.013927	0.021177	-0.65766	0.51076
AR{1} (2,2)	-0.056288	0.031269	-1.8001	0.07184
Beta (1,1)	-0.16444	0.24661	-0.66683	0.50488
Beta (2,1)	-0.031638	0.36413	-0.086887	0.93076
Beta (1,2)	-0.23117	0.10142	-2.2794	0.022645
Beta (2,2)	-0.31595	0.14975	-2.1098	0.034876
Beta (1,3)	0.10253	0.16674	0.61491	0.53861
Beta (2,3)	0.47244	0.2462	1.9189	0.054993
Beta (1,4)	0.23166	0.27456	0.84376	0.3988
Beta (2,4)	-0.026645	0.40541	-0.065724	0.9476
Beta (1,5)	0.029728	0.1489	0.19965	0.84176

Beta (2,5)	0.076951	0.21987	0.34999	0.72635
Beta (1,6)	-0.036858	0.024138	-1.527	0.12677
Beta (2,6)	-0.027378	0.035642	-0.76814	0.44241
Beta (1,7)	0.033457	0.070646	0.47358	0.6358
Beta (2,7)	-0.039092	0.10431	-0.37475	0.70785
Beta (1,8)	-0.0008923	0.0028474	-0.31337	0.754
Beta (2,8)	-0.0042826	0.0042044	-1.0186	0.30839

Table 46: Bitcoin-Litecoin VARX (2,1) estimation results.

All autoregressive coefficients are not statistically significant, like in the VAR (2,1) model. Among the coefficients of predictors, Beta (1,2) and Beta (2,2) are significant. According to estimations, 1% variation in Nikkei 225 at time t-1 leads on average to -0.23% change in Bitcoin returns and -0.31% change in Litecoin returns.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	-0.027976	0.070326	-0.3978	0.69078
Constant (2)	0.10158	0.12797	0.79379	0.42732
AR{1} (1,1)	-0.0048459	0.026268	-0.18448	0.85364
AR{1} (2,1)	-0.049601	0.047799	-1.0377	0.29941
AR{1} (1,2)	-0.013627	0.014395	-0.94661	0.34384
AR{1} (2,2)	-0.054358	0.026194	-2.0752	0.037968
AR{2} (1,1)	-0.0071531	0.026329	-0.27168	0.78587
AR{2} (2,1)	-0.028335	0.047909	-0.59144	0.55423
AR{2} (1,2)	0.041792	0.014404	2.9014	0.003715
AR{2} (2,2)	0.10877	0.02621	4.15	3.3254e-05
Beta (1,1)	-0.17201	0.24588	-0.69957	0.4842
Beta (2,1)	0.26766	0.44741	0.59824	0.54968
Beta (1,2)	-0.22845	0.10125	-2.2563	0.024052
Beta (2,2)	-0.22663	0.18424	-1.2301	0.21867

Beta (1,3)	0.098608	0.16633	0.59285	0.55328
Beta (2,3)	0.52471	0.30266	1.7337	0.082975
Beta (1,4)	0.20733	0.27383	0.75712	0.44898
Beta (2,4)	0.28985	0.49827	0.5817	0.56077
Beta (1,5)	0.028797	0.14847	0.19396	0.84621
Beta (2,5)	0.075037	0.27016	0.27775	0.7812
Beta (1,6)	-0.039356	0.024068	-1.6352	0.10201
Beta (2,6)	-0.012298	0.043796	-0.28081	0.77886
Beta (1,7)	0.031087	0.070456	0.44123	0.65905
Beta (2,7)	-0.10197	0.1282	-0.7954	0.42638
Beta (1,8)	-0.00085871	0.0028385	-0.30252	0.76225
Beta (2,8)	0.0024321	0.005165	0.47087	0.63773

Table 47: Bitcoin-Ripple VARX (2,2) estimation results.

The significant autoregressive coefficients estimated with the VARX (2,2) are very close to those estimated with the VAR (2,2).

Again, the model suggests that Bitcoin returns are correlated to past returns of Nikkei 225, since Beta (1,2) coefficient is statistically significant. The parameter value is also very close to that previously estimated in the Bitcoin-Ethereum VARX (2,1) model.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0071847	0.11701	0.061402	0.95104
Constant (2)	0.041442	0.10417	0.39781	0.69077
AR{1} (1,1)	-0.090872	0.027226	-3.3376	0.00084493
AR{1} (2,1)	-0.030631	0.024239	-1.2637	0.20635
AR{1} (1,2)	0.051276	0.030547	1.6786	0.093234
AR{1} (2,2)	-0.02415	0.027196	-0.88799	0.37455
Beta (1,1)	-0.25396	0.40894	-0.62102	0.53459
Beta (2,1)	-0.027179	0.36407	-0.074652	0.94049

Beta (1,2)	-0.17458	0.16806	-1.0388	0.2989
Beta (2,2)	-0.31022	0.14963	-2.0733	0.038145
Beta (1,3)	0.14355	0.27648	0.5192	0.60362
Beta (2,3)	0.47523	0.24615	1.9307	0.05352
Beta (1,4)	0.74099	0.45636	1.6237	0.10444
Beta (2,4)	0.01824	0.40629	0.044895	0.96419
Beta (1,5)	-0.17311	0.24698	-0.70091	0.48336
Beta (2,5)	0.062062	0.21988	0.28225	0.77775
Beta (1,6)	-0.070185	0.040008	-1.7543	0.079384
Beta (2,6)	-0.027876	0.035619	-0.78263	0.43385
Beta (1,7)	-0.0039162	0.11723	-0.033406	0.97335
Beta (2,7)	-0.035141	0.10437	-0.3367	0.73634
Beta (1,8)	-2.1335e-05	0.0047214	-0.0045187	0.99639
Beta (2,8)	-0.0042962	0.0042034	-1.0221	0.30674

Table 48: Ethereum-Litecoin VARX (2,1) estimation results.

The coefficients in the VARX (2,1) for the Ethereum-Litecoin pair, confirm the finding of the VAR (2,1). The Ethereum returns at time t are significantly influenced by their own values at time t-1, since the p-value of the estimated parameter AR {1} (1,1) is lower than 0.05. The value of the parameter is also close to that previously estimated.

Beta (2,2) coefficient is also significant at 0.05 level. Again, 1% change in Nikkei returns at time t-1 leads on average to roughly -0.31% variation in Litecoin returns at time t. The value is very similar to the one estimated by Bitcoin Litecoin VARX (2,1) model.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.010286	0.11675	0.088106	0.92979
Constant (2)	0.10411	0.12765	0.81563	0.41471
AR{1} (1,1)	-0.075687	0.025971	-2.9143	0.0035648
AR{1} (2,1)	0.029839	0.028395	1.0509	0.29332
AR{1} (1,2)	0.012432	0.023681	0.52496	0.59961

AR{1} (2,2)	-0.080349	0.025891	-3.1033	0.0019138
AR{2} (1,1)	0.029645	0.025988	1.1407	0.25398
AR{2} (2,1)	0.0065865	0.028414	0.23181	0.81669
AR{2} (1,2)	0.0013176	0.023603	0.055823	0.95548
AR{2} (2,2)	0.10628	0.025806	4.1185	3.8133e-05
AR{3} (1,1)	0.078415	0.025975	3.0188	0.0025374
AR{3} (2,1)	0.012115	0.0284	0.42657	0.66969
AR{3} (1,2)	-0.0021727	0.023651	-0.091866	0.9268
AR{3} (2,2)	0.078318	0.025859	3.0287	0.0024562
Beta (1,1)	-0.31346	0.40836	-0.76761	0.44272
Beta (2,1)	0.24178	0.44648	0.54153	0.58814
Beta (1,2)	-0.13056	0.16851	-0.77481	0.43845
Beta (2,2)	-0.24262	0.18424	-1.3169	0.18787
Beta (1,3)	0.16794	0.27592	0.60866	0.54275
Beta (2,3)	0.53622	0.30167	1.7775	0.075485
Beta (1,4)	0.6837	0.4557	1.5003	0.13353
Beta (2,4)	0.20809	0.49823	0.41766	0.67619
Beta (1,5)	-0.13382	0.24669	-0.54247	0.5875
Beta (2,5)	0.12827	0.26971	0.47557	0.63438
Beta (1,6)	-0.068887	0.039957	-1.724	0.084704
Beta (2,6)	-0.0095763	0.043687	-0.2192	0.82649
Beta (1,7)	-0.0078899	0.11695	-0.067462	0.94621
Beta (2,7)	-0.10491	0.12787	-0.82046	0.41196
Beta (1,8)	0.00056222	0.0047198	0.11912	0.90518
Beta (2,8)	0.0024321	0.0051603	0.4713	0.63743

Table 49: Ethereum-Ripple VARX (2,3) estimation results.

Statistically significant parameters estimated in Ethereum Ripple VAR (2,3) are also significant in the VARX (2,3), and their values are similar. According to this model, and the previous ones, Ethereum and Ripple returns are not correlated to change in any predictor returns.

Next, VECX models with different lags are fitted to Litecoin Ripple log prices to define which one is better according to Likelihood Ratio criterium. The following table shows the results of selection criteria.

Litecoin-Ripple	AIC	BIC	Log Likelihood	Number of parameters
VECX (2,1,1)	-8935.99	-8795.63	4494.00	26
VECX (2,1,2)	-8949.15	-8787.19	4504.57	30
VECX (2,1,3)	-9026.91	-8843.36	4547.46	34
VECX (2,1,4)	-9023.44	-8818.28	4549.72	38
VECX (2,1,5)	-9020.8	-8794.06	4552.40	42
VECX (2,1,6)	-9023.62	-8775.28	4557.81	46
VECX (2,1,7)	-9024.59	-8754.65	4562.30	50
VECX (2,1,8)	-9029.65	-8738.11	4568.82	54

U-R	LR Critical value (α = 1%)
4	13.277

Table 50: Bitcoin-Ethereum VECX models selection criteria.

According to both BIC and Likelihood Ratio, the best fitting model is the VECX (2,1,3). The estimation results are reported in the next table.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.056733	0.10602	0.53512	0.59257
Constant (2)	0.015835	0.1279	0.12381	0.90146
Adjustment (1,1)	0.00076244	0.0014199	0.53695	0.5913
Adjustment (2,1)	-0.0045596	0.0017129	-2.6619	0.00777
Impact (1,1)	-0.0023749	0.0044229	-0.53695	0.5913
Impact (2,1)	0.014203	0.0053355	2.6619	0.00777
Impact (1,2)	0.0018139	0.0033782	0.53695	0.5913
Impact (2,2)	-0.010848	0.0040753	-2.6619	0.00777
ShortRun{1}(1,1)	-0.042196	0.027017	-1.5618	0.11833
ShortRun{1}(2,1)	0.065323	0.032592	2.0043	0.045043
ShortRun{1}(1,2)	0.010862	0.021923	0.49546	0.62027
ShortRun{1}(2,2)	-0.092037	0.026446	-3.4801	0.00050117

ShortRun{2}(1,1)	0.001078	0.02714	0.039719	0.96832
ShortRun{2}(2,1)	0.0072616	0.03274	0.22179	0.82447
ShortRun{2}(1,2)	0.0040636	0.021911	0.18546	0.85286
ShortRun{2}(2,2)	0.095072	0.026432	3.5969	0.00032206
ShortRun{3}(1,1)	-0.02094	0.026988	-0.7759	0.43781
ShortRun{3}(2,1)	0.23298	0.032556	7.1561	8.3014e-13
ShortRun{3}(1,2)	0.084747	0.021891	3.8713	0.00010824
ShortRun{3}(2,2)	0.0090172	0.026408	0.34146	0.73276
Beta (1,1)	-0.026371	0.36258	-0.072732	0.94202
Beta (2,1)	0.13284	0.4374	0.3037	0.76136
Beta (1,2)	-0.32896	0.14913	-2.2059	0.027389
Beta (2,2)	-0.23412	0.1799	-1.3014	0.19311
Beta (1,3)	0.50089	0.24534	2.0416	0.041188
Beta (2,3)	0.46958	0.29596	1.5866	0.1126
Beta (1,4)	-0.02754	0.40368	-0.068223	0.94561
Beta (2,4)	0.14707	0.48698	0.30202	0.76264
Beta (1,5)	0.09834	0.21892	0.44921	0.65328
Beta (2,5)	0.1693	0.26409	0.64106	0.52148
Beta (1,6)	-0.023973	0.035499	-0.67532	0.49947
Beta (2,6)	-0.014119	0.042824	-0.32971	0.74162
Beta (1,7)	-0.038651	0.10381	-0.37234	0.70964
Beta (2,7)	-0.089027	0.12522	-0.71093	0.47713
Beta (1,8)	-0.0042815	0.0041844	-1.0232	0.3062
Beta (2,8)	0.0027148	0.0050478	0.53783	0.5907

Cointegration Matrix (Estimated):

-3.1149
2.3791

Cointegration Constant:
-1.3545

Table 51: Litecoin-Ripple VECX (2,1,3) estimation results.

The estimated coefficients that are significant in the VEC (2,1,2) model are significant also in this case and have very close values.

Though, the VECX has one lag more than the VEC estimated in the previous section, and the lag 3 estimated parameters at position (1,2) and (2,1) in the short run equation are significant at 0.05 level. This means that there are significant evidences that values of Litecoin and Ripple at time t-3 are related to the values of the other currency in the model at time t.

The adjustment coefficient of Ripple returns is negative and significant like the one estimated by the VEC (2,1,2) model. While the correction towards long run equilibrium was estimated to be 0.005% each day by the VEC (2,1,2), this model predicts that about 0.004% of departure of Ripple price is corrected daily instead.

The VECX (2,1,3) suggests that Litecoin log-returns are correlated with both the Nikkei 225 and the STOXX 600 European index log-returns. The impact of the predictors on the cryptocurrency is represented by Beta (1,2) and Beta (1,3) coefficients, respectively.

In conclusion, some of the estimated coefficients of exogenous variables included in the model are statistically significant at 0.05 significance level, indicating that they are correlated with cryptocurrencies prices. In particular, the inclusion of Nikkei 225 in VAR and VEC models and STOXX 600 in VEC model could improve the fit to data and, therefore, models predictive power of future cryptocurrencies prices.

In order to obtain better forecasting performances, one would prefer that the models are as parsimonious as possible, in order to decrease the AIC criterion. The models' size is reduced, by cutting those exogenous predictors that appear to be not correlated to cryptocurrencies returns.

In the next tables, VARx and VECx are the models including only the predictors whose coefficients were statistically significant in the first VARX estimations.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0023546	0.00096865	2.4308	0.015066
Constant (2)	0.0033076	0.0016104	2.0539	0.039986
AR{1} (1,1)	0.014038	0.027756	0.50577	0.61302
AR{1} (2,1)	0.019214	0.046145	0.41639	0.67713
AR{1} (1,2)	-0.038158	0.016696	-2.2855	0.022284

AR{1} (2,2)	-0.073649	0.027757	-2.6534	0.0079689
Beta (1,1)	-0.208	0.094182	-2.2085	0.02721
Beta (2,1)	-0.1727	0.15658	-1.103	0.27004

Table 52: Bitcoin-Ethereum VARx (2,1) with only Nikkei 225 returns, estimation results.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0022961	0.00096976	2.3677	0.017901
Constant (2)	0.0018028	0.0014334	1.2577	0.20851
AR{1} (1,1)	-0.0011742	0.031348	-0.037457	0.97012
AR{1} (2,1)	0.04488	0.046336	0.96858	0.33275
AR{1} (1,2)	-0.015243	0.021209	-0.71871	0.47232
AR{1} (2,2)	-0.05856	0.031349	-1.868	0.061759
Beta (1,1)	-0.20508	0.094311	-2.1745	0.029667
Beta (2,1)	-0.16169	0.1394	-1.1599	0.24609

Table 53: Bitcoin-Litecoin VARx (2,1) with only Nikkei 225 returns, estimation results.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0022314	0.0009686	2.3038	0.021236
Constant (2)	0.0022494	0.0017643	1.275	0.20231
AR{1} (1,1)	-0.0034485	0.026299	-0.13113	0.89567
AR{1} (2,1)	-0.045858	0.047903	-0.95731	0.33841
AR{1} (1,2)	-0.014267	0.014421	-0.98928	0.32253
AR{1} (2,2)	-0.055959	0.026268	-2.1304	0.033143
AR{2} (1,1)	-0.007853	0.026351	-0.29802	0.76569
AR{2} (2,1)	-0.028966	0.047998	-0.60349	0.54618
AR{2} (1,2)	0.041669	0.014413	2.8911	0.0038384
AR{2} (2,2)	0.1086	0.026252	4.137	3.5187e-05
Beta (1,1)	-0.19964	0.094241	-2.1184	0.034144

Beta (2,1)	-0.093717	0.17166	-0.54596	0.5851

Table 54: Bitcoin-Ripple VARx (2,2) with only Nikkei 225 returns, estimation results.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0033038	0.0016081	2.0546	0.039922
Constant (2)	0.0019379	0.0014326	1.3528	0.17612
AR{1} (1,1)	-0.087159	0.0272	-3.2044	0.0013534
AR{1} (2,1)	-0.03029	0.024232	-1.25	0.21129
AR{1} (1,2)	0.050035	0.030594	1.6355	0.10195
AR{1} (2,2)	-0.025461	0.027255	-0.9342	0.3502
Beta (1,1)	-0.1735	0.15634	-1.1097	0.26711
Beta (2,1)	-0.15746	0.13928	-1.1306	0.25824

Table 55: Ethereum-Litecoin VARx (2,1) with only Nikkei 225 returns, estimation results.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.0029999	0.0016078	1.8659	0.062057
Constant (2)	0.0018246	0.0017592	1.0372	0.29965
AR{1} (1,1)	-0.072394	0.025922	-2.7928	0.0052254
AR{1} (2,1)	0.032084	0.028363	1.1312	0.25798
AR{1} (1,2)	0.011496	0.023707	0.48493	0.62773
AR{1} (2,2)	-0.081855	0.02594	-3.1555	0.0016021
AR{2} (1,1)	0.029409	0.025994	1.1314	0.2579
AR{2} (2,1)	0.0065627	0.028443	0.23073	0.81752
AR{2} (1,2)	0.0013827	0.023636	0.058501	0.95335
AR{2} (2,2)	0.10594	0.025862	4.0965	4.1952e-05
AR{3} (1,1)	0.080258	0.02591	3.0976	0.0019511
AR{3} (2,1)	0.013464	0.02835	0.47492	0.63484
AR{3} (1,2)	-0.0020746	0.023695	-0.087555	0.93023

AR{3} (2,2)	0.076922	0.025926	2.9669	0.003008
Beta (1,1)	-0.12747	0.15678	-0.81304	0.4162
Beta (2,1)	-0.10431	0.17155	-0.60806	0.54315

Table 56: Ethereum-Ripple VARx (2,3) with only Nikkei 225 returns, estimation results.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.013621	0.022459	0.60648	0.5442
Constant (2)	-0.070288	0.02711	-2.5927	0.0095235
Adjustment (1,1)	0.0007566	0.0014206	0.5326	0.59431
Adjustment (2,1)	-0.0045569	0.0017148	-2.6574	0.007875
Impact (1,1)	-0.0023567	0.0044249	-0.5326	0.59431
Impact (2,1)	0.014194	0.0053414	2.6574	0.007875
Impact (1,2)	0.0018001	0.0033798	0.5326	0.59431
Impact (2,2)	-0.010842	0.0040798	-2.6574	0.007875
ShortRun{1}(1,1)	-0.04293	0.027018	-1.5889	0.11207
ShortRun{1}(2,1)	0.066369	0.032614	2.035	0.041851
ShortRun{1}(1,2)	0.011095	0.02193	0.5059	0.61293
ShortRun{1}(2,2)	-0.092564	0.026473	-3.4966	0.00047123
ShortRun{2}(1,1)	-0.00043614	0.027125	-0.016079	0.98717
ShortRun{2}(2,1)	0.0043176	0.032743	0.13186	0.89509
ShortRun{2}(1,2)	0.0046137	0.021903	0.21065	0.83316
ShortRun{2}(2,2)	0.096663	0.026439	3.6561	0.00025608
ShortRun{3}(1,1)	-0.020116	0.026981	-0.74555	0.45594
ShortRun{3}(2,1)	0.23364	0.032569	7.1736	7.3043e-13
ShortRun{3}(1,2)	0.08467	0.02189	3.8679	0.00010977
ShortRun{3}(2,2)	0.0081788	0.026424	0.30952	0.75693
Beta (1,1)	-0.32307	0.148	-2.1829	0.029041

Beta (2,1)	-0.24964	0.17865	-1.3974	0.1623
Beta (1,2)	0.57077	0.19365	2.9474	0.0032044
Beta (2,2)	0.49644	0.23376	2.1237	0.033693

Cointegration Matrix (Estimated):

-3.1149
2.3791

Cointegration Constant:

15.4935

Table 57: Litecoin-Ripple VECx (2,1,3) with only Nikkei 225 and STOXX 600 log-returns, estimation results.

The following table gathers the AIC results for both the VAR and VEC models and the VARX and VECX estimated previously. The selection criteria are used to help determining which model is best for fitting each multivariate series of observations.

	Bitcoin Ethereum	Bitcoin Litecoin	Bitcoin Ripple	Ethereum Litecoin	Ethereum Ripple	Litecoin Ripple
VAR (2,1)	-10,619.4	-11,390.6		-9,283.86		
VAR (2,2)			-10,149.1			
VAR (2,3)					-8,457.99	
VEC (2,1,2)						-8,962.73
VARx (2,1)	-10,620.3	-11,391.3		-9,281.62		
VARx (2,2)			-10,149.7			
VARx (2,3)					-8,454.8	
VECx (2,1,3)						-9043.83
VARX (2,1)	-10,604.7	-11,380.5		-9,272.69		
VARX (2,2)			-10,137.6			
VARX (2,3)					-8,443.78	
VECX (2,1,3)						-9026.91

Table 58: AIC comparison of best VAR and VEC models.

As one could expect, eliminating the non-significant predictors coefficients produce a lower AIC for each model. The best model according to AIC is highlighted in grey. Results are in line with parameter estimation results, because the VARx and VECx models are preferred to VAR and VEC only when at least one predictor's coefficient is statistically significant.

Although AIC comparison is a first valuation to determine which model to choose, residual analysis should also be performed. This is investigated in the next section.

6. Residual Analysis

Although some useful information can be extracted by the VAR and VEC models estimations, they provide only little insights about the dynamics in cryptocurrencies' returns. Whenever an autoregressive model is estimated, it is always useful to check for autocorrelation in the residuals, in order to determine if the models captured all information in the data.

A first check in the analysis of residuals can be done just by looking at plots of autocorrelation functions produced by R software. For the sake of brevity only ACFs of residuals and squared residuals estimated from VAR (2,1) fitted to Bitcoin and Ethereum are reported here.

As one can observe from the plots, there is clear evidence of statistically significant correlation only at lag 10 for Bitcoin, even if also at lag 5 there might be some autocorrelation, the result is less clear in this case. For what it concerns the Ethereum residuals, the correlogram provides evidence of autocorrelation at lag 3 and an uncertain result at lag 7. Squared residuals correlograms show for both cryptocurrencies clear evidences of autocorrelation.

While the correlograms are a useful starting point in the analysis, they provide only partial information about the correlation features within each VAR or VEC models. In fact, the plots report only the ACF and PACF of the response variable for each equation in the model, but no information are provided about correlation of the response variable residuals and lagged residuals of the other cryptocurrency in the model (cross correlations).

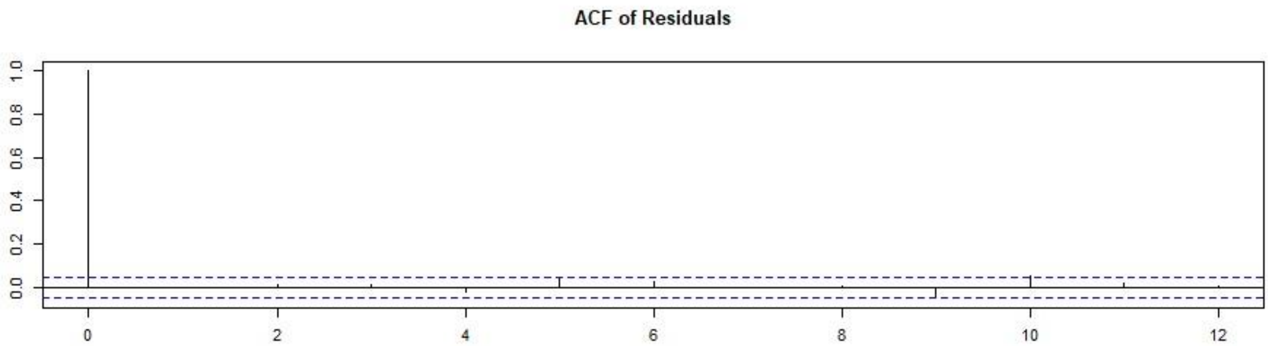

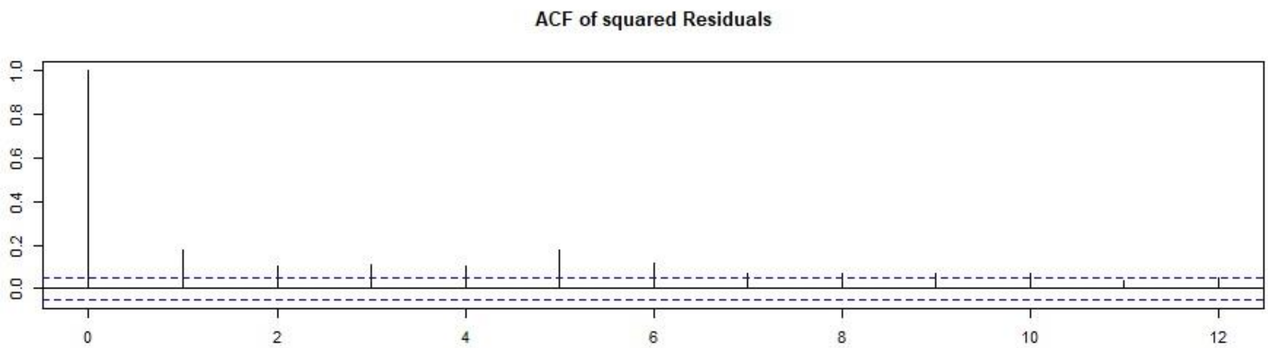

Figure 7: Correlograms of Bitcoin residuals from VAR (2,1) – BTC-ETH.

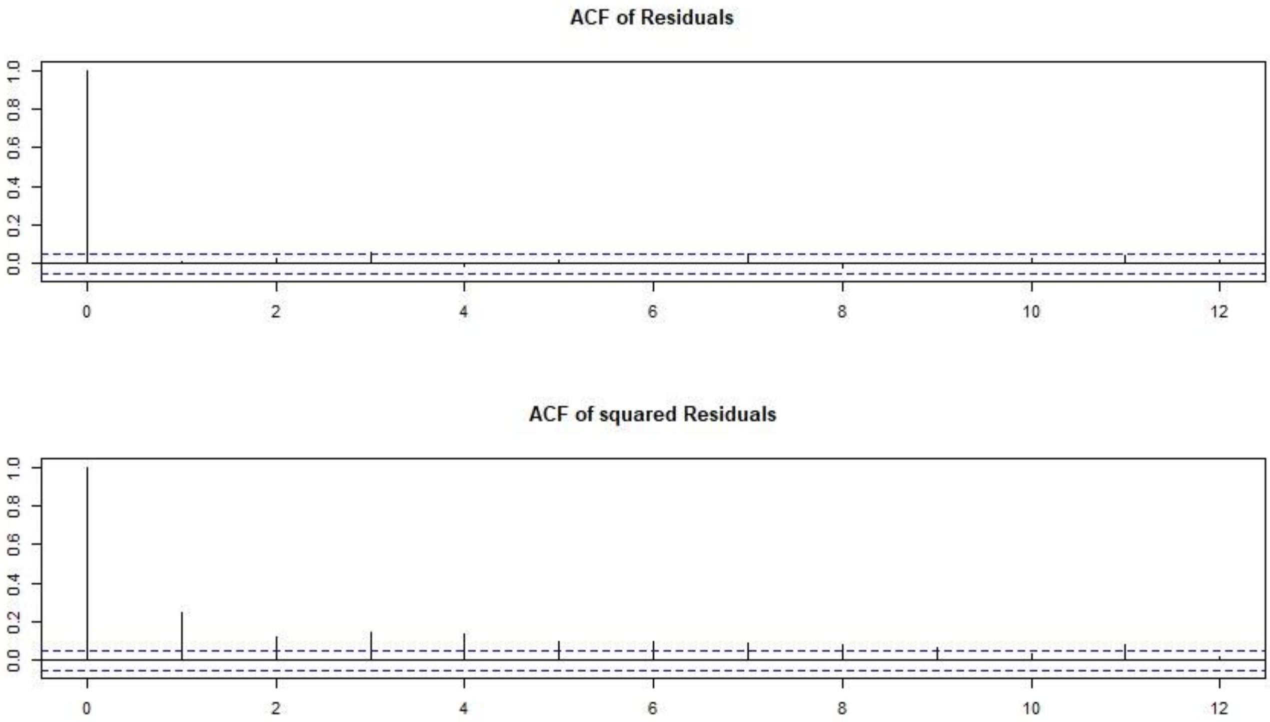

Figure 8: Correlograms of Ethereum residuals from VAR (2,1) – BTC-ETH.

To obtain a more precise result, it is better to run the multivariate Ljung-Box test for each series since, as showed in the previous chapter, the test results considers cross correlations of residuals in each model. The tests are run until lag 8, which roughly corresponds to the natural logarithm of the number of observations in a single series. This setting for selecting an appropriate number of lags is suggested in Tsay, "R. S. Analysis of Financial Time Series. 2nd Ed. Hoboken", NJ: John Wiley & Sons Inc., 2005, p.33. The results provided by R software are reported in the next table.

Cryptocurrencies residuals	p-value	Test result
Bitcoin-Ethereum	0.1913	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals
Bitcoin-Litecoin	0.05427	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals
Bitcoin-Ripple	0.02105	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Ethereum-Litecoin	0.01794	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Ethereum-Ripple	0.01005	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Litecoin-Ripple	< 2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals

Table 59: Portmanteau Ljung-Box test results on VARs and VECs residuals.

Cryptocurrencies residuals^2	p-value	Test result

Bitcoin-Ethereum	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Bitcoin-Litecoin	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Bitcoin-Ripple	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Ethereum-Litecoin	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Ethereum-Ripple	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Litecoin-Ripple	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals

Table 60: Portmanteau McLeod-Li test results on VARs and VECs squared residuals.

Test results indicate that there are no evidences of autocorrelation in residuals of Bitcoin-Ethereum and Bitcoin-Litecoin multivariate series, meaning that the fitted VARs and VEC in the other cases did not model properly the data.

The test on squared residuals strongly reject the null hypothesis of no autocorrelation for all models which means that the models are not capable of describing adequately the volatility process. Thus, it is not possible to exclude the presence of ARCH/GARCH effects in the multivariate residuals.

In the next tables are gathered test results for VARx and VECx models. Only the models estimated in the previous section, with a restricted number of predictors are tested. In particular, the VARx models with Nikkei 225 log-returns and the VECx with Nikkei 225 and STOXX 600 Index log-returns as exogenous regressors.

Cryptocurrencies residuals	p-value	Test result
Bitcoin-Ethereum	0.1649	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals
Bitcoin-Litecoin	0.04849	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Bitcoin-Ripple	0.02381	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Ethereum-Litecoin	0.01555	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Ethereum-Ripple	0.008987	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Litecoin-Ripple	0.0003026	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals

Table 61: Portmanteau Ljung-Box test results on VARXs and VECXs residuals.

Cryptocurrencies residuals^2	p-value	Test result
Bitcoin-Ethereum	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals

Bitcoin-Litecoin	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Bitcoin-Ripple	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Ethereum-Litecoin	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Ethereum-Ripple	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals
Litecoin-Ripple	<2.2e-16	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in squared residuals

Table 62: Portmanteau McLeod-Li test results on VARXs and VECXs squared residuals.

Although the inclusion of exogenous variables obviously changes the p-values of the tests, the result changes only for Bitcoin-Litecoin series in the test on residuals. Tests results on squared residuals do not change with the inclusion of exogenous variable.

In the case of Bitcoin-Ripple and Litecoin-Ripple pairs, the p-values of the test increase with the inclusion of exogenous variables, meaning that a VARX (2,2) and a VECX (2,1,3) fit a bit better to the sample data. These results suggest also that a VAR (2,1) might fit the data better than the VARx (2,1) both for Bitcoin-Ethereum and Bitcoin-Litecoin series. This is the opposite finding of AIC criteria in the previous section, but since the elimination of autocorrelation in residuals would be optimal, the VAR (2,1) is considered in the next steps even though there is only a tiny difference in p-values of the test.

To sum up, the best models for each multivariate time series, selected on their fitting performances to historical data are reported in the table.

Bitcoin Ethereum	VAR (2,1)
Bitcoin Litecoin	VAR (2,1)
Bitcoin Ripple	VARx (2,2)
Ethereum Litecoin	VAR (2,1)
Ethereum Ripple	VAR (2,3)
Litecoin Ripple	VECx (2,1,3)

Table 62: Best models fitting historical data.

Multivariate GARCHs can be used to model the errors of an autoregressive process when autocorrelation is left in residuals and squared residuals. This can be done with the "rmgarch" package of R software, which is used in this thesis to estimate multivariate DCC-GARCH models on the residuals of VAR, VARx and VECx models, which have just been selected on the basis of their fitting performances.

7. Forecasting

The multivariate models forecasting analysis reflects the steps of fitting analysis on historical data. The predictive log-Likelihood for each VAR, VARX, VEC and VECX is computed by filtering the out of sample data with fully specified models estimated on the full sample, in order to determine which is the best model to explain log-returns behaviour in the "new" timeseries. Both cryptocurrencies' log-returns and predictors' log-returns in the out of sample data are gathered for the period 08/02/2020 – 17/05/2020.

The aim of this analysis is to determine which models would have provided the best forecasts, so all models are tested. In fact, although the predictors included failed in describing the returns of cryptocurrencies in the historical data, they might be useful in modelling the out of sample data.

The following tables gather the predictive log-Likelihood, Akaike and Bayes Information criteria computed for each model.

The highlighted models on the first column are those chosen based on their fitting performances. The best results for predictive AIC and BIC are highlighted in grey.

Bitcoin-Ethereum	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-481.1487	-465.5177	246.5744	6
VAR (2,2)	-474.8493	-448.7976	247.4247	10
VAR (2,3)	-468.2226	-431.7502	248.1113	14
VAR (2,4)	-464.0467	-417.1536	250.0233	18
VAR (2,5)	-455.3877	-398.0739	249.6938	22
VAR (2,6)	-447.5350	-379.8006	249.7675	26
VAR (2,7)	-440.5602	-362.4051	250.2801	30
VAR (2,8)	-432.8295	-344.2537	250.4147	34
VARX (2,1)	-521.5408	-464.2270	282.7704	22
VARX (2,2)	-514.4635	-446.7290	283.2317	26
VARX (2,3)	-508.0292	-429.8741	284.0146	30
VARX (2,4)	-501.3795	-412.8037	284.6898	34
VARX (2,5)	-493.1674	-394.1710	284.5837	38
VARX (2,6)	-484.6876	-375.2704	284.3438	42
VARX (2,7)	-477.1855	-357.3477	284.5927	46
VARX (2,8)	-469.4009	-339.1424	284.7004	50
VARx (2,1)	-529.3428	-487.6601	280.6714	16

Table 63: Bitcoin-Ethereum VAR and VARX models selection criteria. Fit to out of sample data.

While the BIC confirm the VAR (2,1) as best model also on out of sample data, AIC indicates that all VARX models fitted the out of sample much better than their VAR counterparts for the Bitcoin-Ethereum series.

In particular, the VARX (2,1) including the full set of predictors minimizes the AIC, meaning that it would have provided the best forecasts for the out of sample data among all models. It is worth to investigate further, to determine which estimated predictors' coefficients caused this sharp decrease in the VARXs AIC criterion.

	Value	Standard Error	T Statistic	P-Value

Constant (1)	0.08655	0.039239	2.2057	0.027403
Constant (2)	0.13552	0.06344	2.1362	0.032659
AR{1} (1,1)	-0.0055095	0.027594	-0.19967	0.84174
AR{1} (2,1)	-0.0092328	0.044613	-0.20695	0.83605
AR{1} (1,2)	-0.043407	0.017075	-2.5421	0.01102
AR{1} (2,2)	-0.080457	0.027607	-2.9144	0.0035636
Beta (1,1)	0.47504	0.16303	2.9138	0.0035708
Beta (2,1)	0.655	0.26359	2.4849	0.012957
Beta (1,2)	-0.28135	0.096088	-2.928	0.0034111
Beta (2,2)	-0.23651	0.15535	-1.5224	0.12791
Beta (1,3)	0.51467	0.14266	3.6077	0.00030889
Beta (2,3)	0.51512	0.23065	2.2334	0.025522
Beta (1,4)	0.50584	0.20432	2.4757	0.013297
Beta (2,4)	0.86987	0.33034	2.6332	0.0084575
Beta (1,5)	-0.14602	0.11942	-1.2227	0.22144
Beta (2,5)	-0.29612	0.19308	-1.5337	0.12511
Beta (1,6)	0.0060969	0.019748	0.30873	0.75752
Beta (2,6)	-0.009661	0.031928	-0.30259	0.76221
Beta (1,7)	-0.081433	0.039503	-2.0614	0.039262
Beta (2,7)	-0.13054	0.063868	-2.0439	0.040959
Beta (1,8)	-0.0030185	0.0023225	-1.2997	0.19371
Beta (2,8)	-0.0021293	0.0037549	-0.56706	0.57067

Table 64: Bitcoin-Ethereum VARX (2,1) fitted to out of sample data.

The results of the selection criteria are clearly supported by the fact that many predictors coefficients are statistically significant. If in the historical data sample, only the Nikkei 225 past returns are linked to returns of cryptocurrencies, the VARX (2,1) suggests that the situation changed dramatically in the out of sample. In fact, the coefficients related to the Standard & Poor 500, Nikkei 225, STOXX 600, Gold, US 10 years bonds interest rates are all statistically significant at 0.05 level for at least one cryptocurrency in the model. In order to decrease further the AIC criteria, the uncorrelated predictors are cancelled from the model. This model is called VARx (2,1) in Table 63.

Bitcoin-Litecoin	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-525.5812	-509.9502	268.7906	6
VAR (2,2)	-519.5858	-493.5341	269.7929	10
VAR (2,3)	-511.9591	-475.4867	269.9795	14
VAR (2,4)	-507.5208	-460.6277	271.7604	18
VAR (2,5)	-500.3081	-442.9944	272.1541	22
VAR (2,6)	-492.2512	-424.5168	272.1256	26
VAR (2,7)	-485.2401	-407.0850	272.6201	30
VAR (2,8)	-477.4392	-388.8634	272.7196	34
VARX (2,1)	-561.3381	-504.0244	302.6691	22
VARX (2,2)	-554.5230	-486.7886	303.2615	26
VARX (2,3)	-547.1811	-469.0260	303.5906	30
VARX (2,4)	-541.0759	-452.5001	304.5380	34
VARX (2,5)	-534.8298	-435.8333	305.4149	38
VARX (2,6)	-525.9255	-416.5084	304.9627	42
VARX (2,7)	-518.3785	-398.5407	305.1893	46
VARX (2,8)	-510.6240	-380.3655	305.3120	50
VARx (2,1)	-569.0715	-527.3888	300.5358	16

Table 65: Bitcoin-Litecoin VAR and VARX models selection criteria. Fit to out of sample data.

The same results obtained for Bitcoin-Ethereum are replicated in this table. The AIC selects the VARX (2,1) as best model, while the BIC still select the VAR (2,1), without exogenous predictors' coefficients. The estimates for the VARX (2,1) are reported in the following table.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.089603	0.039291	2.2805	0.02258
Constant (2)	0.10724	0.056261	1.906	0.056647
AR{1} (1,1)	-0.026487	0.031184	-0.84938	0.39567
AR{1} (2,1)	0.0063391	0.044652	0.14197	0.88711
AR{1} (1,2)	-0.015683	0.021701	-0.72269	0.46987
AR{1} (2,2)	-0.058115	0.031074	-1.8703	0.061449
Beta (1,1)	0.48111	0.16331	2.946	0.0032192
Beta (2,1)	0.61968	0.23384	2.65	0.0080497

Beta (1,2)	-0.28164	0.096253	-2.9261	0.0034328
Beta (2,2)	-0.34643	0.13782	-2.5136	0.011951
Beta (1,3)	0.50935	0.1429	3.5644	0.00036469
Beta (2,3)	0.72219	0.20462	3.5295	0.00041635
Beta (1,4)	0.48628	0.20452	2.3776	0.017425
Beta (2,4)	0.35264	0.29286	1.2041	0.22853
Beta (1,5)	-0.1487	0.11962	-1.2431	0.21384
Beta (2,5)	-0.14111	0.17129	-0.82384	0.41003
Beta (1,6)	0.0059273	0.019791	0.2995	0.76456
Beta (2,6)	0.010651	0.028338	0.37585	0.70703
Beta (1,7)	-0.084591	0.039556	-2.1385	0.032477
Beta (2,7)	-0.10157	0.056641	-1.7933	0.072929
Beta (1,8)	-0.0029724	0.0023268	-1.2775	0.20143
Beta (2,8)	-0.0043305	0.0033317	-1.2998	0.19368

Table 66: Bitcoin-Litecoin VARX (2,1) fitted to out of sample data.

According to estimation results of Bitcoin-Litecoin VARX (2,1), Litecoin returns are correlated with less predictors. Significant coefficients in the Litecoin equation are only those of Standard and Poor 500, Nikkei 225 and STOXX 600. AIC is minimized again by the reduced version of the VARX (2,1), the VARx (2,1), according to results in Table 65.

Bitcoin-Ripple	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-473.2106	-457.5796	242.6053	6
VAR (2,2)	-465.9616	-439.9099	242.9808	10
VAR (2,3)	-456.7796	-420.3073	242.3898	14
VAR (2,4)	-450.8654	-403.9724	243.4327	18
VAR (2,5)	-442.8794	-385.5656	243.4397	22
VAR (2,6)	-434.5533	-366.8189	243.2766	26
VAR (2,7)	-427.3266	-349.1715	243.6633	30
VAR (2,8)	-419.8758	-331.3001	243.9379	34
VARx (2,2)	-425.8563	-295.5978	262.9281	14
VARX (2,1)	-508.8765	-451.5627	276.4382	22
VARX (2,2)	-501.1599	-433.4254	276.5799	26
VARX (2,3)	-492.1376	-413.9825	276.0688	30
VARX (2,4)	-484.5270	-395.9512	276.2635	34

VARX (2,5)	-476.9483	-377.9518	276.4741	38
VARX (2,6)	-467.6483	-358.2311	275.8241	42
VARX (2,7)	-460.0000	-340.1622	276.0000	46
VARX (2,8)	-452.7051	-322.4466	276.3526	50
VARx (2,1)	-516.3343	-474.6516	274.1672	16

Table 67: Bitcoin-Ripple VAR and VARX models selection criteria. Fit to out of sample data.

Other than VAR and VARX models, for Bitcoin-Ripple series also the VARx (2,2) containing only the Nikkei 225 as predictor, is reported in the table. This is the best model selected on fitting performances to historical data, according to residual analysis and AIC selection criterion results. As one can see, it is not a good model for forecasting, as it performs worse than both a VAR (2,2) and a VARX (2,2). The best model, the one minimizing the AIC, is the VARX (2,1). Like for other crypto-currencies series, the selected model does not minimize also the BIC, which instead indicates a VAR (2,1) as the best alternative.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.089363	0.039281	2.2749	0.02291
Constant (2)	0.12221	0.068302	1.7893	0.073569
AR{1} (1,1)	-0.029718	0.025821	-1.1509	0.24975
AR{1} (2,1)	-0.071992	0.044897	-1.6035	0.10882
AR{1} (1,2)	-0.016855	0.014822	-1.1372	0.25547
AR{1} (2,2)	-0.060367	0.025773	-2.3423	0.019166
Beta (1,1)	0.47735	0.16328	2.9235	0.003461
Beta (2,1)	0.58886	0.28391	2.0741	0.038068
Beta (1,2)	-0.27813	0.096276	-2.8889	0.0038658
Beta (2,2)	-0.27108	0.1674	-1.6193	0.10538
Beta (1,3)	0.50744	0.14284	3.5524	0.00038168
Beta (2,3)	0.67397	0.24837	2.7135	0.0066569
Beta (1,4)	0.48454	0.20448	2.3696	0.017807
Beta (2,4)	0.44635	0.35555	1.2554	0.20934
Beta (1,5)	-0.14908	0.1196	-1.2465	0.21257
Beta (2,5)	-0.054143	0.20795	-0.26036	0.79458

Beta (1,6)	0.0052848	0.019776	0.26723	0.78929
Beta (2,6)	0.010127	0.034387	0.2945	0.76838
Beta (1,7)	-0.084308	0.039546	-2.1319	0.033014
Beta (2,7)	-0.1196	0.068762	-1.7393	0.08198
Beta (1,8)	-0.0029962	0.0023259	-1.2882	0.19768
Beta (2,8)	-0.00066885	0.0040443	-0.16538	0.86864

Table 68: Bitcoin-Ripple VARX (2,1) fitted to out of sample data.

According to estimation results of the Bitcoin-Ripple VARX (2,1), Ripple returns seems correlated only with Standard and Poor 500 and STOXX 600 indices returns. A VARx (2,1) is estimated also in this case. Information criterion results gathered in table 67 indicate again it would have been the best model for forecasting.

Ethereum-Litecoin	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-489.4816	-473.8506	250.7408	6
VAR (2,2)	-482.0987	-456.0470	251.0494	10
VAR (2,3)	-473.3119	-436.8395	250.6559	14
VAR (2,4)	-469.2402	-422.3471	252.6201	18
VAR (2,5)	-461.2159	-403.9022	252.6080	22
VAR (2,6)	-452.3414	-384.6070	252.1707	26
VAR (2,7)	-443.5406	-365.3855	251.7703	30
VAR (2,8)	-435.7482	-347.1724	251.8741	34
VARX (2,1)	-515.1476	-457.8338	279.5738	22
VARX (2,2)	-507.5971	-439.8627	279.7986	26
VARX (2,3)	-499.1134	-420.9583	279.5567	30
VARX (2,4)	-492.3025	-403.7268	280.1513	34
VARX (2,5)	-484.7839	-385.7874	280.3919	38
VARX (2,6)	-475.3443	-365.9271	279.6721	42
VARX (2,7)	-466.0546	-346.2168	279.0273	46
VARX (2,8)	-458.2215	-327.9630	279.1107	50
VARx (2,1)	-524.9789	-483.2962	278.4895	16

Table 69: Ethereum-Litecoin VAR and VARX models selection criteria. Fit to out of sample data.

Results for the Ethereum-Litecoin series are very close to the previous ones. VARX (2,1) is the best performer in a first analysis. A better version of it, cutting the non-significant predictors' parameters is represented as VARx (2,1) in the table.

	Value	Standard Error	T Statistic	P-Value

Constant (1)	0.1335	0.063412	2.1053	0.035263
Constant (2)	0.1043	0.05624	1.8545	0.063669
AR{1} (1,1)	-0.10034	0.02688	-3.733	0.00018919
AR{1} (2,1)	-0.040762	0.02384	-1.7098	0.087298
AR{1} (1,2)	0.041096	0.030229	1.3595	0.17399
AR{1} (2,2)	-0.033953	0.02681	-1.2664	0.20535
Beta (1,1)	0.66832	0.26288	2.5423	0.011012
Beta (2,1)	0.60264	0.23314	2.5849	0.0097417
Beta (1,2)	-0.24743	0.15492	-1.5972	0.11023
Beta (2,2)	-0.3397	0.1374	-2.4724	0.01342
Beta (1,3)	0.50644	0.23053	2.1968	0.028033
Beta (2,3)	0.73	0.20446	3.5704	0.00035643
Beta (1,4)	0.87704	0.33016	2.6564	0.0078973
Beta (2,4)	0.37141	0.29281	1.2684	0.20465
Beta (1,5)	-0.30236	0.1929	-1.5674	0.11702
Beta (2,5)	-0.13486	0.17108	-0.78828	0.43053
Beta (1,6)	-0.0093701	0.031876	-0.29395	0.76879
Beta (2,6)	0.0098541	0.02827	0.34856	0.72742
Beta (1,7)	-0.12844	0.063842	-2.0119	0.044229
Beta (2,7)	-0.098481	0.056621	-1.7393	0.08198
Beta (1,8)	-0.0022278	0.0037535	-0.59354	0.55282
Beta (2,8)	-0.0043834	0.003329	-1.3168	0.18792

Table 70: Ethereum-Litecoin VARX (2,1) fitted to out of sample data.

Ethereum-Ripple	AIC	BIC	Log Likelihood	Number of parameters
VAR (2,1)	-475.9238	-460.2928	243.9619	6
VAR (2,2)	-469.3897	-443.3380	244.6949	10
VAR (2,3)	-458.1239	-421.6515	243.0620	14
VAR (2,4)	-453.3490	-406.4559	244.6745	18

VAR (2,5)	-445.4275	-388.1138	244.7138	22
VAR (2,6)	-437.0012	-369.2668	244.5006	26
VAR (2,7)	-427.7418	-349.5867	243.8709	30
VAR (2,8)	-419.9755	-331.3997	243.9878	34
VARX (2,1)	-492.1881	-434.8743	268.0940	22
VARX (2,2)	-485.2126	-417.4782	268.6063	26
VARX (2,3)	-474.4981	-396.3430	267.2491	30
VARX (2,4)	-467.6016	-379.0258	267.8008	34
VARX (2,5)	-460.0661	-361.0696	268.0330	38
VARX (2,6)	-452.5588	-343.1417	268.2794	42
VARX (2,7)	-443.4847	-323.6469	267.7423	46
VARX (2,8)	-435.9904	-305.7319	267.9952	50

VARx (2,1)	-504.1812	-467.7088	266.0906	14

Table 71: Ethereum-Ripple VAR and VARX models selection criteria. Fit to out of sample data.

In the case of Ethereum-Ripple the VARX (2,1) is the best model according to predictive AIC. Since both the Ethereum and Ripple returns seem to be not correlated with Nikkei 225 past returns, but they are still correlated with Standard & Poor 500, STOXX 600, Gold and US 10 years bond interest rates, the VARx (2,1) in this case has only 4 exogenous variables.

	Value	Standard Error	T Statistic	P-Value
Constant (1)	0.135	0.06343	2.1283	0.033309
Constant (2)	0.12305	0.068366	1.7998	0.07189
AR{1} (1,1)	-0.087924	0.025336	-3.4703	0.00051989
AR{1} (2,1)	0.01726	0.027308	0.63204	0.52736
AR{1} (1,2)	0.012404	0.023503	0.52775	0.59768
AR{1} (2,2)	-0.081745	0.025332	-3.2269	0.0012513
Beta (1,1)	0.66296	0.26302	2.5206	0.011716
Beta (2,1)	0.63143	0.28349	2.2274	0.025923
Beta (1,2)	-0.24315	0.15506	-1.5681	0.11686
Beta (2,2)	-0.29327	0.16713	-1.7548	0.079301
Beta (1,3)	0.51378	0.23057	2.2283	0.02586
Beta (2,3)	0.65473	0.24852	2.6346	0.0084241
Beta (1,4)	0.87393	0.33033	2.6457	0.0081529

Beta (2,4)	0.43418	0.35603	1.2195	0.22266
Beta (1,5)	-0.29807	0.19296	-1.5447	0.12241
Beta (2,5)	-0.071573	0.20798	-0.34414	0.73074
Beta (1,6)	-0.0089725	0.031898	-0.28128	0.77849
Beta (2,6)	0.012458	0.034381	0.36237	0.71708
Beta (1,7)	-0.13003	0.063858	-2.0362	0.041728
Beta (2,7)	-0.12058	0.068828	-1.7519	0.079791
Beta (1,8)	-0.0021414	0.0037546	-0.57034	0.56845
Beta (2,8)	-0.000688	0.0040468	-0.17001	0.865

Table 72: Ethereum-Ripple VARX (2,1) fitted to out of sample data.

Litecoin-Ripple	AIC	BIC	Log Likelihood	Number of parameters
VEC (2,1,1)	-506.4390	-480.3873	263.2195	10
VEC (2,1,2)	-500.0019	-463.5295	264.0010	14
VEC (2,1,3)	-486.6498	-439.7567	261.3249	18
VEC (2,1,4)	-481.0252	-423.7114	262.5126	22
VEC (2,1,5)	-472.9194	-405.1850	262.4597	26
VEC (2,1,6)	-463.6963	-385.5412	261.8481	30
VEC (2,1,7)	-457.5520	-368.9762	262.7760	34
VEC (2,1,8)	-450.5138	-351.5173	263.2569	38
VECx (2,1,3)	-497.8608	-440.5471	270.9304	22
VECX (2,1,1)	-462.2999	-394.5654	257.1499	26
VECX (2,1,2)	-456.1483	-377.9932	258.0742	30
VECX (2,1,3)	-435.2284	-346.6526	251.6142	34
VECX (2,1,4)	-428.9186	-329.9221	252.4593	38
VECX (2,1,5)	-420.1268	-310.7096	252.0634	42
VECX (2,1,6)	-410.0368	-290.1990	251.0184	46
VECX (2,1,7)	-404.9371	-274.6785	252.4685	50
VECX (2,1,8)	-397.6137	-256.9345	252.8068	54

Table 73: Litecoin-Ripple VEC and VECX models selection criteria. Fit to out of sample data.

In the case of Litecoin-Ripple, the simple VEC (2,1,1) performed best in the out of sample data.

Different sets of predictors were included into several VEC models with different lags, but despite some predictors' parameters were statistically significant, these models are not able to overperform the VEC (2,1,1).

	Value	Standard Error	T Statistic	P-Value

Constant (1)	0.019711	0.022741	0.86676	0.38607
Constant (2)	-0.082641	0.027375	-3.0189	0.0025372
Adjustment (1,1)	0.0011363	0.0014113	0.80512	0.42075
Adjustment (2,1)	-0.0052613	0.001699	-3.0968	0.0019562
Impact (1,1)	-0.0036144	0.0044893	-0.80512	0.42075
Impact (2,1)	0.016736	0.0054041	3.0968	0.0019562
Impact (1,2)	0.0027426	0.0034064	0.80512	0.42075
Impact (2,2)	-0.012699	0.0041005	-3.0968	0.0019562
ShortRun{1}(1,1)	-0.065185	0.026714	-2.4401	0.014682
ShortRun{1}(2,1)	0.048859	0.032158	1.5194	0.12867
ShortRun{1}(1,2)	0.013992	0.022001	0.63597	0.5248
ShortRun{1}(2,2)	-0.096189	0.026485	-3.6319	0.00028138

Cointegration Matrix (Estimated):

-3.1809
2.4136

Cointegration Constant:

15.7804

Table 74: Litecoin-Ripple VEC (2,1,1) fitted to out of sample data.

Results obtained by fitting the models to out of sample data differ from those obtained by estimating autoregressive models on past log-returns. An analyst choosing the models for forecasting relying on fitting performances to historical data, would have never selected the best models to produce forecasts.

It is evident from estimations produced that during the out of sample period the correlation between predictors and cryptocurrencies increased dramatically. The reason for this is quite intuitive. The coronavirus emergency impact on the global markets led many financial assets to drop in value simultaneously, increasing their correlations.

To sum up, the best models for each multivariate time series, selected on their performance in filtering the out of sample data are collected in the following table.

	Model	Predictors
Bitcoin Ethereum	VARx (2,1)	S&P 500, Nikkei 225, STOXX, Gold, US 10Y bond rate

Bitcoin Litecoin	VARx (2,1)	S&P 500, Nikkei 225, STOXX, Gold, US 10Y bond rate
Bitcoin Ripple	VARx (2,1)	S&P 500, Nikkei 225, STOXX, Gold, US 10Y bond rate
Ethereum Litecoin	VARx (2,1)	S&P 500, Nikkei 225, STOXX, Gold, US 10Y bond rate
Ethereum Ripple	VARx (2,1)	S&P 500, STOXX, Gold, US 10Y bond rate
Litecoin Ripple	VEC (2,1,1)	

Table 75: Best models estimated on full data, filtering the out of sample data.

Anyway, this is just one part of the forecasting valuation for multivariate models. It must be completed with the inclusion of DCC GARCHs modelling the residuals that are left after the VARx and VEC filters are applied to data. Forecasting performances of conditional volatility models are tested in the next Chapter.

VII DCC GARCH Models: Implementation and results

1. Models Selection

Multivariate DCC GARCH models are fitted to residuals of best VAR, VARx and VECx models, estimated on historical data. This step is carried out by using the "rmgarch" package available for R software. The best DCC GARCH is selected with Akaike Information Criterion instead of Likelihood Ratio, because for these models, the log-Likelihood function is not estimated in a single step and Likelihood Ratio results might not be reliable. The fitted models are:

- *DCC (1,1) – N, GARCH (1,1)(1,1) – N*: represents a multivariate DCC model composed of two univariate GARCH (1,1) with normally distributed innovations for the residual series and a DCC equation which has one lag of symmetric innovation and one lag of correlation. The standardized innovations in the DCC are modelled by a multivariate Normal distribution;
- *DCC (1,1) – N, GARCH (1,2)(1,2) – N*: this multivariate DCC model is composed of two univariate GARCH (1,2) with normally distributed innovations for the residual series and a DCC equation with one lag of symmetric innovations and one lag of conditional covariance matrix. The standardized innovations are again multivariate Normal;
- *DCC (1,1) – N, GARCH-GJR (1,1)(1,1) – N*: while the DCC equation is the same as before, the univariate GARCH-GJRs allow for large impact of negative residuals values on the conditional variance;
- *DCC (1,1) – N, GARCH (1,1)(1,1) – T*: this is similar to the first model. The only difference is that innovations in the univariate GARCHs follow a Student-t distribution;
- *DCC (1,1) – T, GARCH (1,1)(1,1) – N*: it is the first of a series of models in which the standardized innovations in the DCC equation are modelled by a multivariate Student T distribution rather than a multivariate Normal. The other features are the same as in the first model;
- *DCC (1,1) – T, GARCH (1,2)(1,2) – N*: the model is specified like the second one, with the only difference that standardized innovations in the DCC equation are multivariate Student T distributed;

- *DCC (1,1) – T, GARCH-GJR (1,1)(1,1) – N*: the model is specified like the third one but, again, the standardized innovations in the DCC equation are multivariate T distributed;
- *DCC (1,1) – T, GARCH (1,1)(1,1) – T*: in this model both the innovations of the univariate GARCH processes and the standardized innovations in the DCC equation follow a Student-t distribution, an univariate distribution in the GARCH processes, a multivariate distribution in the DCC;
- *DCC (1,1) –* T, *GARCH (1,1)(1,2) – T:* the two univariate models within the DCC-GARCH do not need to be identical. In this case the first GARCH has only one lag of conditional volatility while the second one has two lags;
- *DCC (1,1) – T, GARCH (1,1)(1,3) – T*: the first series is modelled by a GARCH(1,1) and the second univariate GARCH has three lags of conditional volatility;
- *DCC (1,1) – T, GARCH (1,1)(1,4) – T*: the second univariate GARCH has four lags of conditional volatility;
- *DCC (1,1) – T, GARCH (1,1)(2,3) – T*: the second univariate GARCH has two lags of squared residuals and three lags of conditional volatility;
- *DCC (1,1) – T, GARCH (1,2)(1,1) – T*: In this case the first univariate GARCH model is defined with two lags of conditional volatility, while the second series is modelled by a GARCH (1,1);
- *DCC (1,1) – T, GARCH (1,2)(1,3) – T*: the first model has two lags of conditional volatility like in the previous case, the second univariate GARCH has three of them;
- *DCC (1,1) – T, GARCH (1,2)(1,4) – T*: the second model has four lags of conditional volatility in this case;
- *DCC (1,1) – T, GARCH (1,2)(2,3) – T*: the second univariate GARCH has two lags of squared residuals and three lags of conditional volatility;
- *DCC (1,1) – T, GARCH-GJR (1,1)(1,1) – T*: while the innovations in the model are again T distributed, the univariate GARCHs model have a GJR feature like in third and seventh cases;
- *DCC (1,1) – T, GARCH (1,2)(1,2) – T*: both the univariate GARCH components have two lags of conditional variance. Innovations of GARCHs and DCC are T distributed, like in previous cases.

The inclusion of a third and fourth lag of the conditional variance, and the inclusion of a second lag of squared residuals might seem quite specific. This is because the other specifications were tested first, and they resulted to be not appropriate to model the univariate Ripple residuals timeseries. The inclusion of further lags is an attempt of obtaining estimation improvements in that direction.

The best models according to Akaike and Bayes information criteria are highlighted in the following table.

		VAR (2,1) BTC-ETH	VAR (2,1) BTC-LTC	VARx (2,2) BTC-XRP	VAR (2,1) ETH-LTC	VAR (2,3) ETH-XRP	VECx (2,1,3) LTC-XRP
DCC (1,1)-N GARCH (1,1)(1,1)-N 9 parameters	AIC	-7.3819	-7.5403	-7.2208	-6.4910	-6.4543	-6.3535
	BIC	-7.3522	-7.5105	-7.1911	-6.4612	-6.4246	-6.3238
DCC (1,1)-N	AIC	-7.3842	-7.5460	-7.2251	-6.4935	-6.4574	-6.3537

GARCH (1,2)(1,2)-N 11 parameters	BIC	-7.3479	-7.5097	-7.1887	-6.4571	-6.4211	-6.3174
DCC (1,1)-N GARCH-GJR (1,1)(1,1)-N 11 parameters	AIC	-7.3801	-7.5325	-7.2327	-6.4892	-6.4535	-6.3693
	BIC	-7.3438	-7.4962	-7.1964	-6.4529	-6.4171	-6.3330
DCC (1,1)-N GARCH (1,1)(1,1)-T 11 parameters	AIC	-7.3110	-7.4155	-7.1629	-6.4898	-6.3488	-6.2886
	BIC	-7.2747	-7.3792	-7.1266	-6.4585	-6.3125	-6.2523
DCC (1,1)-T GARCH (1,1)(1,1)-N 10 parameters	AIC	-7.7688	-8.1212	-7.7792	-6.9994	-6.9485	-6.9648
	BIC	-7.7357	-8.0881	-7.7462	-6.9664	-6.9154	-6.9317
DCC (1,1)-T GARCH (1,2)(1,2)-N 12 parameters	AIC	-7.7729	-8.1251	-7.7862	-7.0010	-6.9504	-6.9681
	BIC	-7.7333	-8.0854	-7.7466	-6.9613	-6.9108	-6.9285
DCC (1,1)-T GARCH-GJR (1,1)(1,1)-N 12 parameters	AIC	-7.7645	-8.1089	-7.7748	-6.9909	-6.9440	-6.9594
	BIC	-7.7248	-8.0692	-7.7351	-6.9513	-6.9044	-6.9198
DCC (1,1)-T GARCH (1,1)(1,1)-T 12 parameters	AIC	-7.7905	-8.1828	-7.7941	-7.0330	-6.9795	-6.9815
	BIC	-7.7509	-8.1432	-7.7545	-6.9934	-6.9398	-6.9419
DCC (1,1)-T GARCH (1,1)(1,2)-T 13 Parameters	AIC	-7.7886	-8.1834	-7.7973	-7.0320	-6.9783	-6.9883
	BIC	-7.7456	-8.1404	-7.7543	-6.9891	-6.9354	-6.9453
DCC (1,1)-T GARCH (1,1)(1,3)-T 14 Parameters	AIC	-7.7845	-8.1817	-7.7963	-7.0256	-6.9716	-6.9864
	BIC	-7.7382	-8.1354	-7.7501	-6.9794	-6.9253	-6.9401
DCC (1,1)-T GARCH (1,1)(1,4)-T 15 Parameters	AIC	-7.7843	-8.1806	-7.7979	-7.0248	-6.9686	-6.9895
	BIC	-7.7348	-8.1311	-7.7483	-6.9753	-6.9190	-6.9400
DCC (1,1)-T GARCH (1,1)(2,3)-T 15 Parameters	AIC	-7.7848	-8.1805	-7.7951	-7.0244	-6.9704	-6.9832
	BIC	-7.7352	-8.1309	-7.7455	-6.9748	-6.9208	-6.9337
DCC (1,1)-T GARCH (1,2)(1,1)-T 13 Parameters	AIC	-7.7927	-8.1858	-7.7955	-7.0342	-6.9797	-6.9799
	BIC	-7.7497	-8.1428	-7.7525	-6.9913	-6.9368	-6.9369
DCC (1,1)-T GARCH (1,2)(1,3)-T 15 Parameters	AIC	-7.7896	-8.1871	-7.7993	-7.0300	-6.9751	-6.9859
	BIC	-7.7400	-8.1376	-7.7497	-6.9804	-6.9255	-6.9364
DCC (1,1)-T GARCH (1,2)(1,4)-T 16 Parameters	AIC	-7.7898	-8.1860	-7.8008	-7.0292	-6.9717	-6.9890
	BIC	-7.7370	-8.1331	-7.7480	-6.9763	-6.9189	-6.9362
DCC (1,1)-T GARCH (1,2)(2,3)-T 16 Parameters	AIC	-7.7895	-8.1859	-7.7981	-7.0288	-6.9738	-6.9827
	BIC	-7.7366	-8.1330	-7.7452	-6.9759	-6.9210	-6.9299
DCC (1,1)-T	AIC	-7.7852	-8.1804	-7.7903	-7.0334	-6.9793	-6.9816

GARCH-GJR (1,1)(1,1)-T 14 parameters	BIC	-7.7389	-8.1341	-7.7440	-6.9871	-6.9331	-6.9353
DCC (1,1)-T GARCH (1,2)(1,2)-T 14 parameters	AIC	-7.7931	-8.1879	-7.8001	-7.0351	-6.9806	-6.9875
	BIC	-7.7468	-8.1416	-7.7539	-6.9888	-6.9343	-6.9412

Table 76: Akaike and Bayes Information Criteria results for DCC GARCHs models.

The best models, according to AIC and BIC, are highlighted in grey. While both results are reported, the best fitting models indicated by AIC are chosen for further analysis.

According to Akaike Information Criterion results, the assumption of a multivariate Student's t distribution for the DCC model leads to a better fit to VARs and VEC residuals than the assumption of multivariate normal distribution. This means that the probability of observing extreme values in the innovations series is larger than under the multivariate normal distribution assumption.

In general, the addition of lags of conditional volatility in univariate GARCHs produces the best fit, while the addition of lags of squared lagged residuals does not lead to significant improvements.

2. Models Estimations

The R software package "rmgarch" is used in this exercise in order to estimate the parameters of selected models. Results are reported in the following tables.

	Value	Standard Error	T Statistic	P-Value
Constant (BTC)	0.000022	0.000015	1.4586	0.144678
Alpha1 (BTC)	0.182880	0.027123	6.7425	0.000000
Beta1 (BTC)	0.478944	0.155875	3.0726	0.002122
Beta2 (BTC)	0.337176	0.134760	2.5021	0.012348
Shape (BTC)	3.332591	0.117533	28.3545	0.000000
Constant (ETH)	0.000279	0.000102	2.7394	0.006155
Alpha1 (ETH)	0.294075	0.046306	6.3506	0.000000
Beta1 (ETH)	0.406453	0.141732	2.8678	0.004134
Beta2 (ETH)	0.298471	0.121566	2.4552	0.014080
Shape (ETH)	3.298277	0.069908	47.1799	0.000000
A1 (DCC)	0.084411	0.017263	4.8898	0.000001
B1 (DCC)	0.911871	0.018625	48.9604	0.000000
Shape (DCC)	4.000000	0.398130	10.0470	0.000000

Table 77: Bitcoin-Ethereum DCC (1,1) – T, GARCH (1,2) (1,2) – T estimation results.

It is worth to recall that both the univariate GARCHs and the DCC component must satisfy the conditions of positiveness and stationarity in a DCC GARCH model. Since all parameters are positive and the sum of the coefficients within each component of the model is less than one it is possible to state that this DCC GARCH model fulfils the requirements.

P-values above 0.05 level indicate that the constant in the univariate Bitcoin GARCH is not statistically significant. All other estimated parameters in the univariate GARCHs are significant. The shape parameters are quite explicative. The two innovation distributions have respectively 3.33 and 3.30 degrees of freedom, meaning that marginal distributions' tails are quite large with respect to tails of the standard Normal distribution.

As it is highlighted in the previous section, the Akaike Information Criterion always favour models assuming a multivariate Student t distribution for the DCC equation. In fact, the degrees of freedom estimated for the multivariate Student t distribution assumed for the standardized innovations in the DCC equation are 4, meaning that the probability of observing extreme values is high.

A1 and B1 parameters in the DCC equation control for the dynamic dependencies of conditional correlations. The former parameter is linked to the lagged values of conditional covariance matrix of standardized marginal innovations at time t-1, while the latter defines the impact of previous squared standardized innovations on conditional covariance matrix at time t. Since B1 is quite large, the lagged squared standardized innovations are strongly explicative in the definition of the conditional covariance matrix at time t.

The final output of a DCC model are the conditional correlations among each pair of cryptocurrencies. In the next Figure the estimated conditional correlation among Bitcoin and Ethereum is plotted as a timeseries.

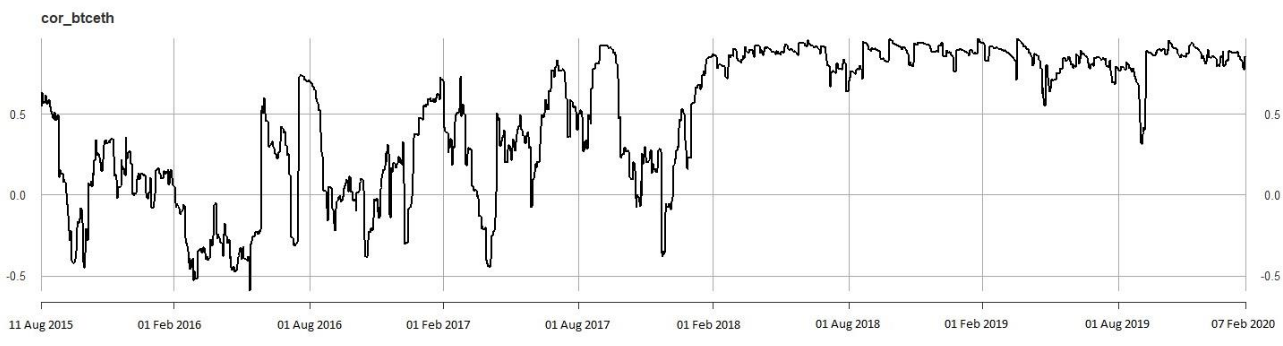

Figure 9: Estimated Bitcoin-Ethereum correlation.

This plot shows how the value of estimated correlation between the two timeseries changes over time. The first thing one can notice is a break in the timeseries in concomitance with the cryptocurrency bubbles of December 2017. While before that period the correlation among the two timeseries was unstable, after the break the correlation is high and characterized by small volatility.

	Value	Standard Error	T Statistic	P-Value

Constant (BTC)	0.000021	0.000014	1.41769	0.156283
Alpha1 (BTC)	0.184610	0.028671	6.43898	0.000000
Beta1 (BTC)	0.455562	0.168354	2.70597	0.006810
Beta2 (BTC)	0.358828	0.143521	2.50019	0.012413
Shape (BTC)	3.316010	0.060091	55.18325	0.000000
Constant (LTC)	0.000092	0.000206	0.44475	0.656504
Alpha1 (LTC)	0.169831	0.149734	1.13422	0.256701
Beta1 (LTC)	0.516952	0.152110	3.39854	0.000677
Beta2 (LTC)	0.312217	0.141735	2.20282	0.027608
Shape (LTC)	3.046509	0.043844	69.48577	0.000000
A1 (DCC)	0.073497	0.013288	5.53107	0.000000
B1 (DCC)	0.912616	0.017225	52.98295	0.000000
Shape (DCC)	4.000000	0.733021	5.45687	0.000000

Table 78: Bitcoin-Litecoin DCC (1,1) – T, GARCH (1,2)(1,2) – T estimation results.

The estimated DCC (1,1) – T, GARCH (1,2)(1,2) – T satisfy the stationarity and positiveness conditions required by DCC-GARCHs. The two univariate GARCHs' constants and Alpha1 coefficient in the LTC equation are not significant.

The estimated Alpha1 parameter in Bitcoin GARCH equation is roughly 0.18, a similar value to the one estimated in the previous model.

The estimated beta parameters are quite large, like in the case of Bitcoin-Ethereum model. The conditional volatility values at time t are strongly affected by their own previous values.

B1 in the DCC equation is large, meaning that the lagged squared standardized innovations are strongly explicative of the conditional covariance matrix at time t. The estimated degrees of freedom for the standardized innovations in the DCC equation are 4, indicating again that extreme values are likely to be observed.

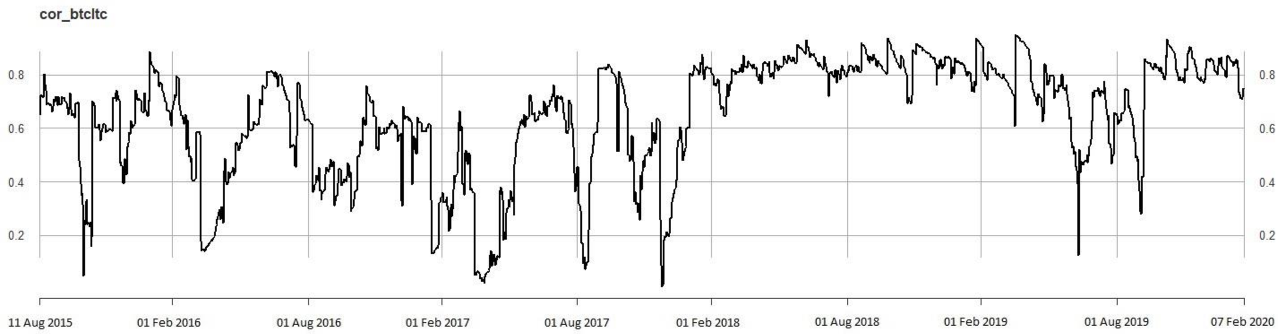

Figure 10: Estimated Bitcoin-Litecoin correlation.

It is possible to recognize again a change in the behaviour of the correlation timeseries in December 2017. After this period the correlation among the two series stabilized on high values, about 0.8. In the case of Bitcoin-Litecoin, though, the estimated correlation in the period after the cryptocurrencies bubble burst is less stable than the correlation between Bitcoin and Ethereum, as a drop in the correlation is observed in the summer/autumn 2019.

	Value	Standard Error	T Statistic	P-Value
Constant (BTC)	0.000026	0.000018	1.438923	0.150172
Alpha1 (BTC)	0.186963	0.031290	5.975231	0.000000
Beta1 (BTC)	0.422238	0.146785	2.876578	0.004020
Beta2 (BTC)	0.389800	0.126222	3.088198	0.002014
Shape (BTC)	3.293467	0.090461	36.407609	0.000000
Constant (XRP)	0.000252	0.000118	2.140320	0.032329
Alpha1 (XRP)	0.365324	0.064290	5.682454	0.000000
Beta1 (XRP)	0.342873	0.246088	1.393294	0.163531
Beta2 (XRP)	0.129869	0.129862	1.000052	0.317285
Beta3 (XRP)	0.000007	0.248042	0.000026	0.999979
Beta4 (XRP)	0.160927	0.163780	0.982580	0.325814
Shape (XRP)	3.089364	0.142120	21.737707	0.000000
A1 (DCC)	0.072187	0.013972	5.166607	0.000000
B1 (DCC)	0.924164	0.015399	60.013186	0.000000
Shape (DCC)	4.000000	0.723430	5.529217	0.000000

Table 79: Bitcoin-Ripple DCC (1,1) – T, GARCH (1,2) (1,4) – T estimation results.

The constant term in the Bitcoin GARCH and all Beta parameters in the Ripple univariate GARCH model are not significant at 0.05 level.

While Beta parameters are not significant in the univariate GARCH, alpha parameter is significant and quite large in comparison to previous estimated models. This means that, differently from other residual timeseries, lagged values of squared Ripple residuals are more impactful than their lagged conditional variance according to GARCH estimations.

The estimated degrees of freedom are low both in the two univariate GARCHs and in the DCC equation, suggesting fat tailed distributions for the innovations of each model's component.

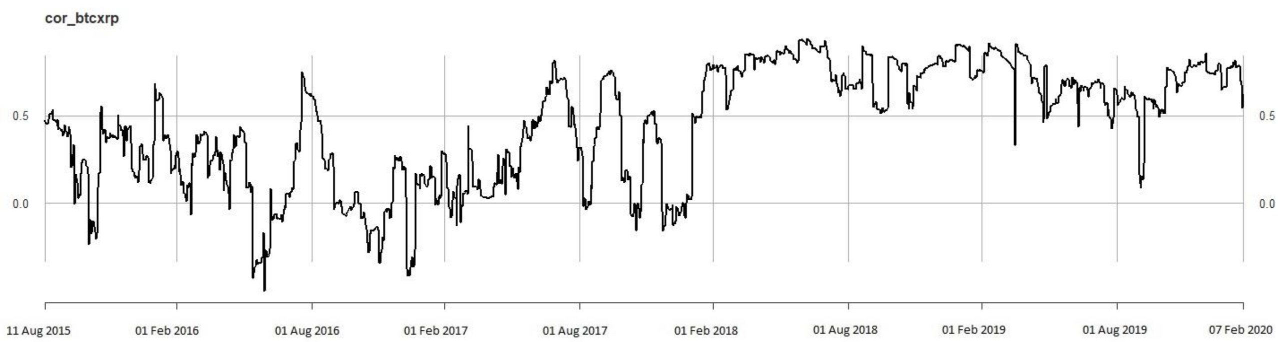

Figure 11: Estimated Bitcoin-Ripple correlation.

The correlation timeseries has the same behaviour as those of Bitcoin-Ethereum and Bitcoin-Litecoin pairs. The same break in the timeseries is evident at the end of 2017. While the correlation before the break between the two series were highly volatile, it became more stable after that period. The Bitcoin-Ripple correlation after the break is lower than correlation among Bitcoin and other cryptocurrencies, on average.

	Value	Standard Error	T Statistic	P-Value
Constant (ETH)	0.000270	0.000099	2.72012	0.006526
Alpha1 (ETH)	0.290399	0.046312	6.27044	0.000000
Beta1 (ETH)	0.416781	0.141024	2.95540	0.003123
Beta2 (ETH)	0.291819	0.119736	2.43719	0.014802
Shape (ETH)	3.317472	0.292871	11.32742	0.000000
Constant (LTC)	0.000094	0.000178	0.52823	0.597337
Alpha1 (LTC)	0.165585	0.125172	1.32285	0.185885
Beta1 (LTC)	0.533812	0.146335	3.64788	0.000264
Beta2 (LTC)	0.299603	0.131212	2.28335	0.022410

Shape (LTC)	3.021692	0.152812	19.77396	0.000000
A1 (DCC)	0.043510	0.010041	4.33318	0.000015
B1 (DCC)	0.955371	0.010180	93.85047	0.000000
Shape (DCC)	4.000000	1.053986	3.79512	0.000148

Table 80: Ethereum Litecoin DCC (1,1) – T, GARCH (1,2) (1,2) – T estimation results.

The constant and Alpha1 parameters in the LTC univariate GARCH are not statistically significant. This is the same result estimated in the DCC GARCH model applied to Bitcoin-Litecoin.

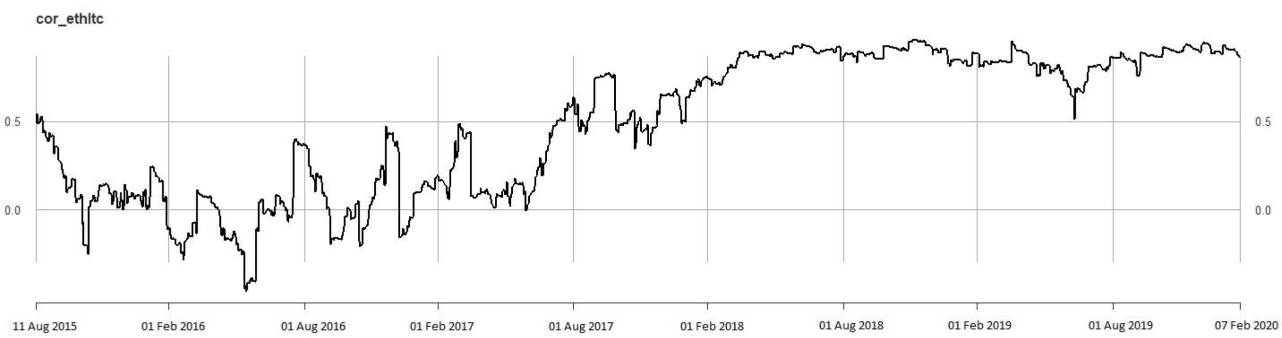

Figure 12: Estimated Ethereum-Litecoin correlation.

The Ethereum and Litecoin conditional correlation timeseries plot shows a very similar pattern to that of correlations among other cryptocurrencies' pairs.

	Value	Standard Error	T Statistic	P-Value
Constant (ETH)	0.000274	0.000099	2.76890	0.005625
Alpha1 (ETH)	0.291356	0.045845	6.35525	0.000000
Beta1 (ETH)	0.371693	0.140498	2.64554	0.008156
Beta2 (ETH)	0.335951	0.120273	2.79323	0.005219
Shape (ETH)	3.302318	0.245452	13.45405	0.000000
Constant (XRP)	0.000076	0.000403	0.18902	0.850076
Alpha1 (XRP)	0.154491	0.294892	0.52389	0.600356
Beta1 (XRP)	0.485535	0.211816	2.29225	0.021891
Beta2 (XRP)	0.358974	0.252127	1.42379	0.154509
Shape (XRP)	3.106983	0.357236	8.69728	0.000000
A1 (DCC)	0.042349	0.009860	4.29523	0.000017

B1 (DCC)	0.956176	0.010164	94.07573	0.000000
Shape (DCC)	4.000000	0.856024	4.67277	0.000003

Table 81: Ethereum-Ripple DCC (1,1) – T, GARCH (1,2) (1,2) – T estimation results.

Alpha1 and Beta2 in the Ripple univariate GARCH are not significant. The degrees of freedom of the Student t distributions assumed for the innovations are again quite low, 3.30 and 3.10 for Ethereum and Ripple, respectively. The DCC equation estimated parameters are very close to those estimated in previous cases and the degrees of freedom are still 4.

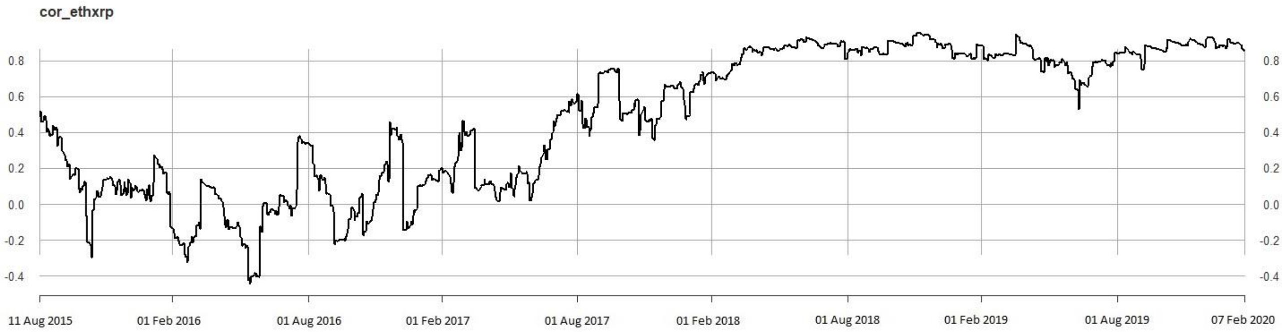

Figure 13: Estimated Ethereum-Ripple correlation.

The estimated conditional correlation timeseries shows a break in the data in the same period of 2017 and follows a pattern which is similar to that of other cryptocurrencies pairs.

	Value	Standard Error	T Statistic	P-Value
Constant (LTC)	0.000089	0.000098	0.911615	0.361971
Alpha1 (LTC)	0.141869	0.072156	1.966141	0.049282
Beta1 (LTC)	0.857131	0.092560	9.260269	0.000000
Shape (LTC)	3.066720	0.246301	12.451084	0.000000
Constant (XRP)	0.000297	0.000143	2.069454	0.038503
Alpha1 (XRP)	0.370129	0.083675	4.423398	0.000010
Beta1 (XRP)	0.262852	0.298868	0.879494	0.379134
Beta2 (XRP)	0.223212	0.110960	2.011639	0.044258
Beta3 (XRP)	0.000002	0.477681	0.000005	0.999996
Beta4 (XRP)	0.142804	0.278082	0.513533	0.607578
Shape (XRP)	3.250142	0.353550	9.192891	0.000000

A1 (DCC)	0.041987	0.009565	4.389720	0.000011
B1 (DCC)	0.956618	0.010172	94.044812	0.000000
Shape (DCC)	4.000000	1.329010	3.009759	0.002615

Table 82: Litecoin Ripple DCC (1,1) – T, GARCH (1,1) (1,4) – T estimation results.

AIC favours the DCC (1,1) – T, GARCH (1,1) (1,4) – T to model the Litecoin-Ripple VECx (2,1,3) residuals. In this model, Constant (LTC), Beta1 (XRP), Beta3 (XRP) and Beta4 (XRP) coefficients are not statistically significant at 5% significant level. DCC GARCH equation estimations are close to those in previous models and the degree of freedom of the multivariate Student t distribution of residuals are 4, again.

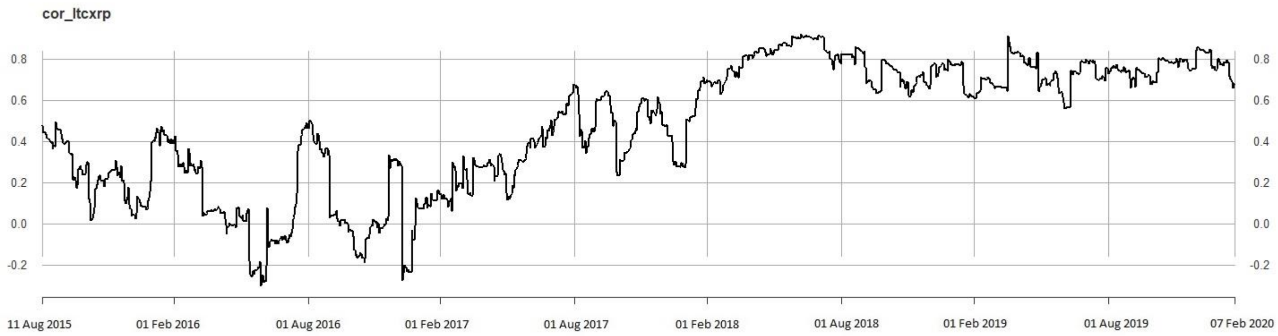

Figure 14: Estimated Litecoin-Ripple correlation.

The break in the conditional correlation timeseries is clear cut also for Litecoin-Ripple pair.

3. Residual Analysis

In order to determine whether DCC GARCH estimations are reliable and to assess models' goodness of fit, it is necessary to test for autocorrelation in their residuals.

The multivariate portmanteau Ljung-Box test and McLeod-Li test are run respectively on DCC GARCHs standardized residuals and squared standardized residuals to determine the presence of autocorrelation in these series. The test is run again considering 8 lags and the rejection of the null hypothesis of no autocorrelation requires the estimated p-value to be lower than 5% significance level. Test results are reported in the following tables.

Cryptocurrencies residuals	p-value	Test result
Bitcoin-Ethereum	3.8665e-04	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Bitcoin-Litecoin	0.2701	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals
Bitcoin-Ripple	0.0423	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Ethereum-Litecoin	0.0468	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals

Ethereum-Ripple	2.3063e-05	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals
Litecoin-Ripple	1.7167e-05	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals

Table 83: Portmanteau Ljung-Box test results on DCC GARCHs residuals.

Cryptocurrencies residuals^2	p-value	Test result
Bitcoin-Ethereum	0.5909	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals^2
Bitcoin-Litecoin	0.9710	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals^2
Bitcoin-Ripple	0.6220	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals^2
Ethereum-Litecoin	0.9981	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals^2
Ethereum-Ripple	0.9979	p-value > 0.05 critical value, do not reject the null hypothesis of no autocorrelation in residuals^2
Litecoin-Ripple	4.6634e-08	p-value < 0.05 critical value, reject the null hypothesis of no autocorrelation in residuals^2

Table 84: Portmanteau McLeod-Li test results on DCC GARCHs squared residuals.

As one may expect, Multivariate Ljung-Box autocorrelation test results show that the DCC GARCHs do not address the presence of autocorrelation in residuals of conditional mean models, but they address the autocorrelation in their squared residuals.

In fact, McLeod-Li tests on squared residuals indicate the absence of autocorrelation in all multivariate series, except for Litecoin-Ripple one. The inclusion of multivariate GARCHs seems to be not helpful in modelling the volatility process of this cryptocurrencies' returns pair.

The absence of autocorrelation is not the only desirable feature that the standardized residuals should have. In fact, also for DCC GARCH it is recommended to check whether the residuals follow the distribution assumed by the models. Firstly, it is checked whether the marginal distributions of each model residuals are those assumed.

The same univariate Kolmogorov-Smirnov test applied to analyse univariate models' residuals is fitted to marginal residuals of univariate GARCHs in the DCC GARCH models. Next, the distribution analysis of multivariate residuals is carried out by running the Baringhaus-Franz test. This test is necessary to define whether residuals are distributed as the multivariate distribution assumed by each model.

Cryptocurrency	Model residuals	Distribution	p-value	Test result
Bitcoin	Btc-Eth DCC (1,1)-T GARCH (1,2) (1,2)-T	Student-t Df = 3.3325	1.161e-07	p-value < 0.05 critical value, residuals are not Student t distributed

	Btc-Ltc DCC (1,1)-T GARCH (1,2) (1,2)-T	Student-t Df = 3.3160	6.41e-10	p-value < 0.05 critical value, residuals are not Student t distributed
	Btc-Xrp DCC (1,1)-T GARCH (1,2) (1,4)-T	Student-t Df = 3.2934	6.095e-09	p-value < 0.05 critical value, residuals are not Student t distributed
Ethereum	Btc-Eth DCC (1,1)-T GARCH (1,2) (1,2)-T	Student-t Df = 3.2982	8.173e-12	p-value < 0.05 critical value, residuals are not Student t distributed
	Eth-Ltc DCC (1,1)-T GARCH (1,2) (1,2)-T	Student-t Df = 3.3174	9.485e-08	p-value < 0.05 critical value, residuals are not Student t distributed
	Eth-Xrp DCC (1,1)-T GARCH (1,2) (1,2)-T	Student-t Df = 3.3023	3.578e-13	p-value < 0.05 critical value, residuals are not Student t distributed
Litecoin	Btc-Ltc DCC (1,1)-T GARCH (1,2) (1,2)-T	Student-t Df = 3.0465	1.052e-11	p-value < 0.05 critical value, residuals are not Student t distributed
	Eth-Ltc DCC (1,1)-T GARCH (1,2) (1,2)-T	Student-t Df = 3.0216	6.34e-12	p-value < 0.05 critical value, residuals are not Student t distributed
	Ltc-Xrp DCC (1,1)-T GARCH (1,1) (1,4)-T	Student-t Df = 3.0667	1.212e-13	p-value < 0.05 critical value, residuals are not Student t distributed
Ripple	Btc-Xrp DCC (1,1)-T GARCH (1,2) (1,4)-T	Student-t Df = 3.0893	2.937e-12	p-value < 0.05 critical value, residuals are not Student t distributed
	Eth-Xrp DCC (1,1)-T GARCH (1,2) (1,2)-T	Student-t Df = 3.1069	1.986e-10	p-value < 0.05 critical value, residuals are not Student t distributed
	Ltc-Xrp DCC (1,1)-T GARCH (1,1) (1,4)-T	Student-t Df = 3.0667	1.232e-10	p-value < 0.05 critical value, residuals are not Student t distributed

Table 85: Univariate Kolmogorov-Smirnov test on marginal DCC GARCHs residuals.

The Kolmogorov-Smirnov test is run for each univariate standardized residual series estimated with the DCC GARCHs. This means that for each cryptocurrency, three different univariate series are estimated in the multivariate GARCHs. The degree of freedom for the Student t distributions used as comparison are taken from the DCC GARCHs estimated shape parameters. The models estimating the univariate series are reported in the second column of the table.

The test indicates that the univariate residuals series are never distributed like assumed by the models. This indicates that the estimated models do not appropriately fit the data. To get further information about the distribution of residuals one could look at the Q-Q Plot comparing the quantiles of the univariate theoretical distribution assumed for the residuals and the quantiles of their empirical distribution. In order to cut the number of plots reported in the thesis, only few examples are reported. The Q-Q plots which are not reported share very similar results with the following ones.

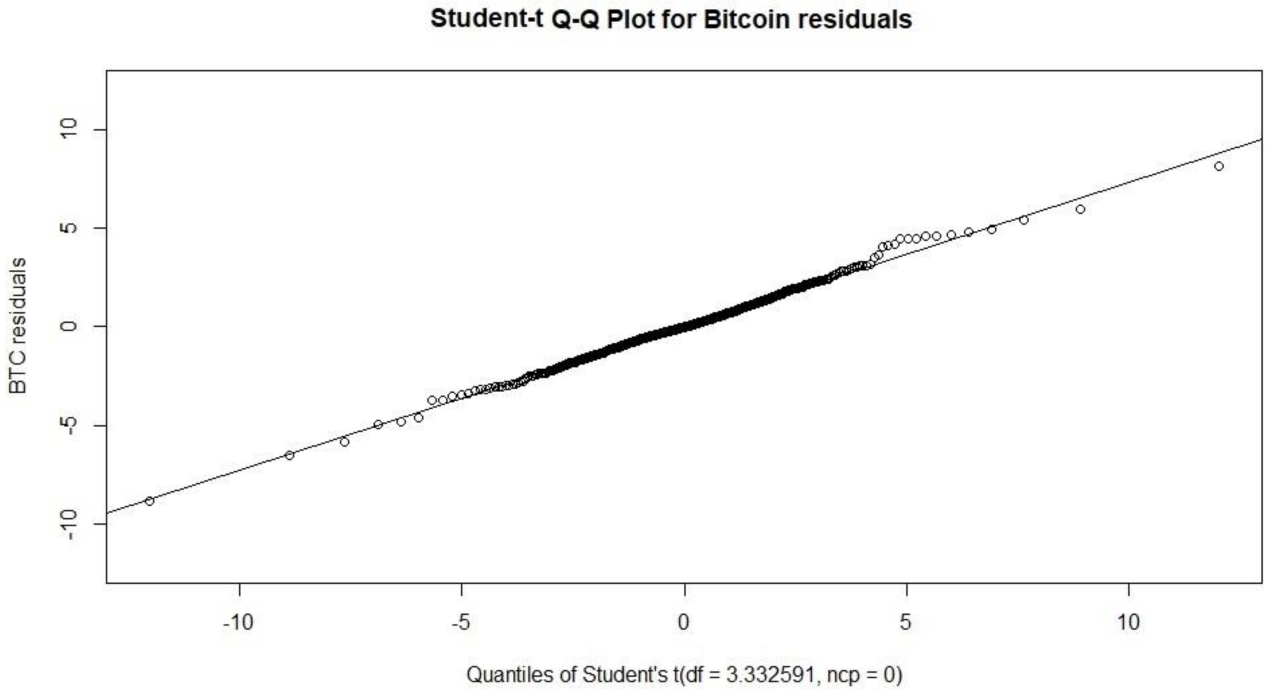

Figure 15: Marginal Bitcoin residuals from Btc-Eth DCC (1,1)-T, GARCH (1,2) (1,2)-T compared to their theoretical distribution.

The figure reports the Q-Q Plot for the Bitcoin residuals of DCC (1,1)-T, GARCH (1,2) (1,2)-T, modelling Bitcoin-Ethereum VAR (2,1) residuals. The sample distribution is compared to a Student-t distribution with 3.3325 degrees of freedom. One could observe that the standardized residuals distribution has a fatter right tail than the Student-t. This means that the assumed distribution underestimates the probability of large values in the residuals.

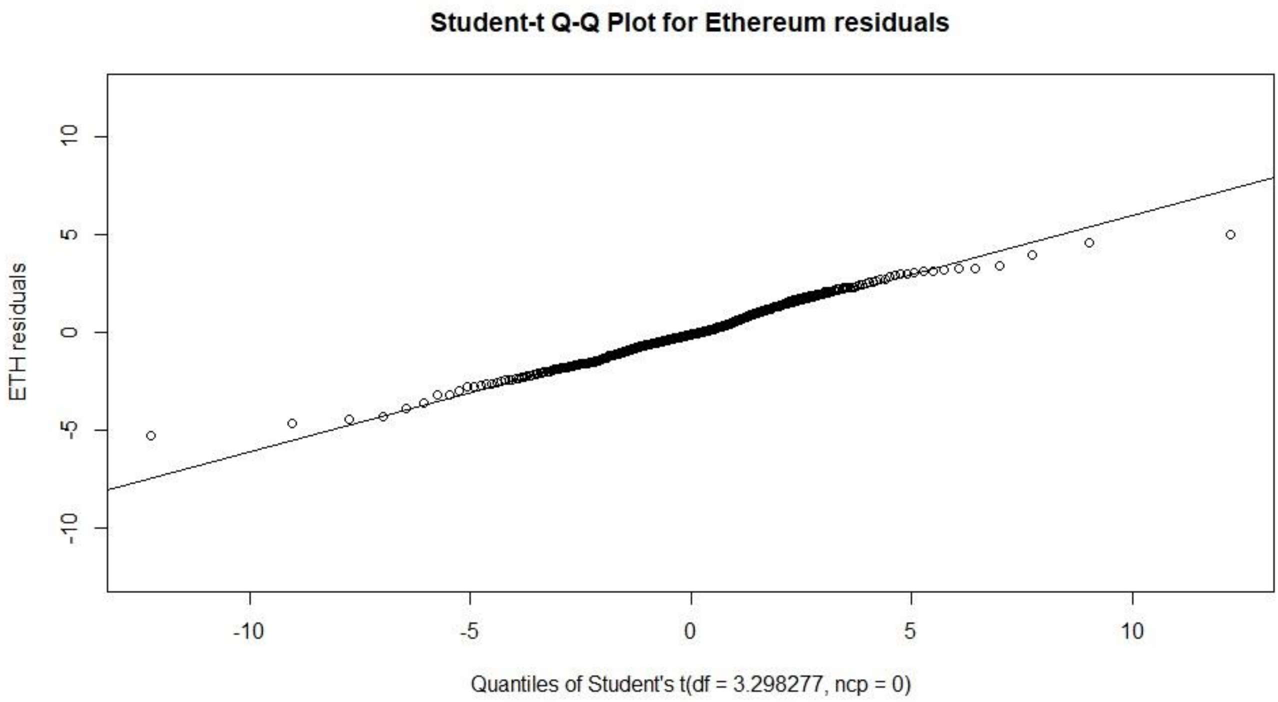

Figure 16: Marginal Ethereum residuals from Btc-Eth DCC (1,1)-T, GARCH (1,2) (1,2)-T compared to their theoretical distribution.

It is clear that the sample distribution of Ethereum residuals is light tailed in comparison to the estimated Student-t in the model. This imply that the assumed Student-t over-estimates the probability of observing extreme values in the Ethereum residuals distribution.

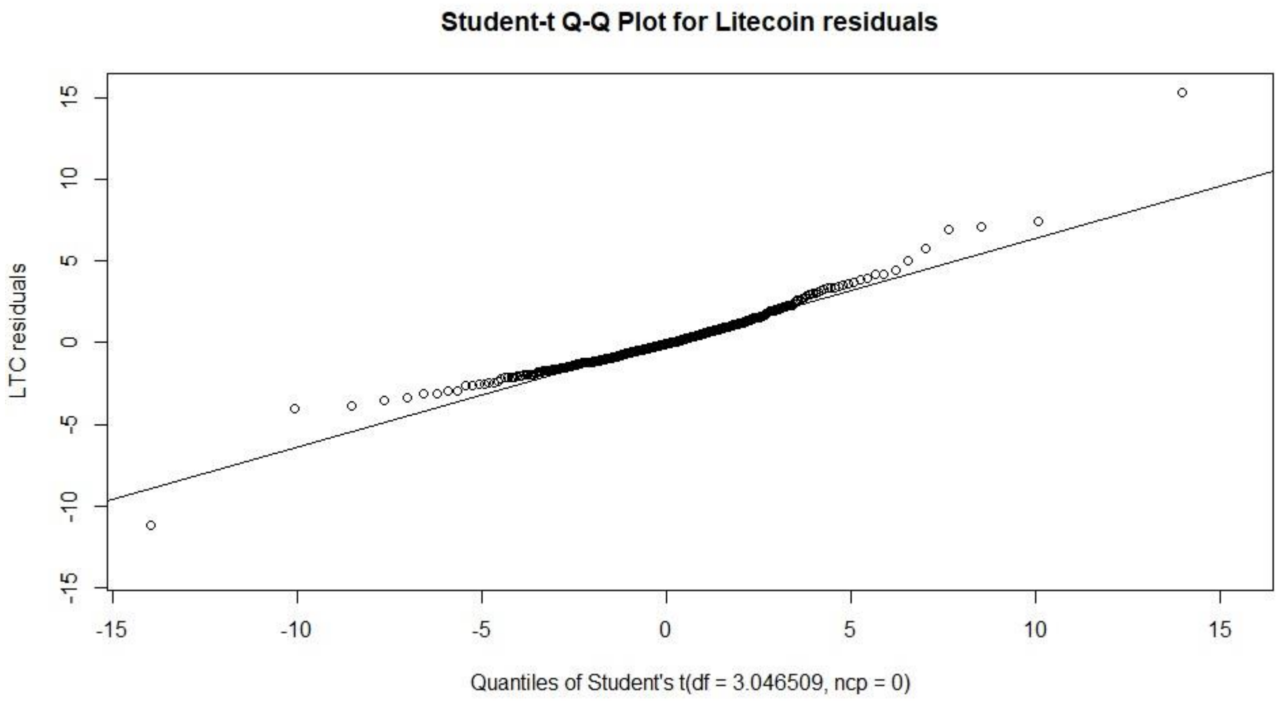

Figure 17: Marginal Litecoin residuals from Btc-Ltc DCC (1,1)-T, GARCH (1,2) (1,2)-T compared to their theoretical distribution.

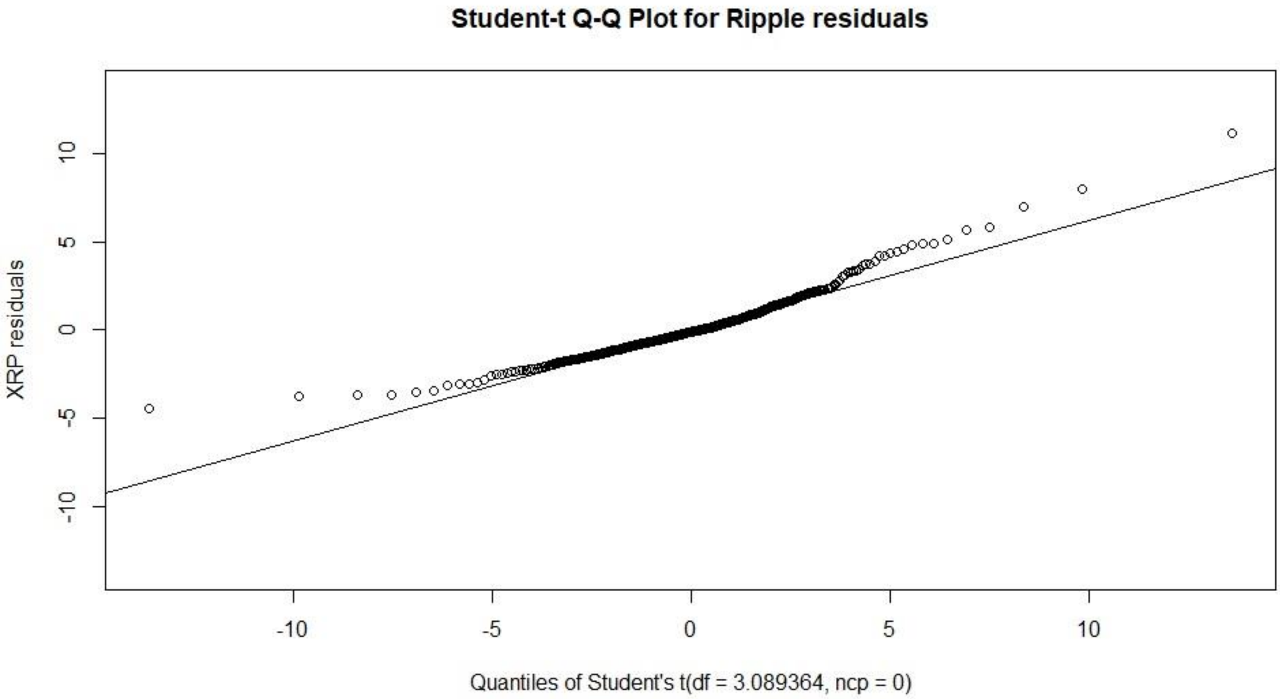

Figure 18: Marginal Ripple residuals from Btc-Xrp DCC (1,1)-T, GARCH (1,2) (1,4)-T compared to their theoretical distribution.

Figures 17 and 18 show similar results. The assumed Student-t distributions have fatter left tail and lighter right tail than the empirical distributions for both the Litecoin and Ripple residuals series. The following table reports the results gathered from the multivariate in-sample Baringhaus-Franz test.

Cryptocurrency	Model residuals	Distribution	p-value	Test result
Bitcoin-Ethereum	DCC (1,1)-T GARCH (1,2) (1,2)-T	Multivariate Student-t Df = 4.0	0.000	p-value < 0.05 critical value, residuals are not multivariate Student t distributed
Bitcoin-Litecoin	DCC (1,1)-T GARCH (1,2) (1,2)-T	Multivariate Student-t Df = 4.0	0.000	p-value < 0.05 critical value, residuals are not multivariate Student t distributed
Bitcoin-Ripple	DCC (1,1)-T GARCH (1,2) (1,4)-T	Multivariate Student-t Df = 4.0	0.000	p-value < 0.05 critical value, residuals are not multivariate Student t distributed
Ethereum-Litecoin	DCC (1,1)-T GARCH (1,2) (1,2)-T	Multivariate Student-t Df = 4.0	0.000	p-value < 0.05 critical value, residuals are not multivariate Student t distributed
Ethereum-Ripple	DCC (1,1)-T GARCH (1,2) (1,2)-T	Multivariate Student-t Df = 4.0	0.000	p-value < 0.05 critical value, residuals are not multivariate Student t distributed
Litecoin-Ripple	DCC (1,1)-T GARCH (1,1) (1,4)-T	Multivariate Student-t Df = 4.0	0.000	p-value < 0.05 critical value, residuals are not multivariate Student t distributed

Table 86: Multivariate Baringhaus-Franz test on DCC GARCHs residuals.

The degree of freedom of the multivariate Student t are 4.0 in each test, consistently with estimation results obtained from the DCC GARCHs. The covariance matrix assumed for each theoretical distribution is extracted from DCC estimations, too.

Results are quite clear cut also in this case. The null hypothesis that residuals are distributed as multivariate Student t is rejected at all significance level for each DCC GARCH model, indicating that the assumption for the innovations multivariate distribution is never correct.

4. Forecasting

Multivariate DCC GARCHs are added to the best performing conditional mean models selected for forecasting in the previous Chapter. Like for VAR and VEC models the predictive log-Likelihood and predictive Information Criteria are computed for DCC GARCHs to determine which model filters better the data. In this case the residuals of VARs and VEC models are filtered.

The best multivariate GARCH models are selected according to predictive AIC criterion. Not only the log-Likelihoods are computed in two separated steps, but they are also the Likelihoods of models estimated on all data available, obtained filtering the out of sample data. A direct comparison through the Likelihood Ratio would not be correct.

		VARx (2,1) BTC-ETH	VARx (2,1) BTC-LTC	VARx (2,1) BTC-XRP	VARx (2,1) ETH-LTC	VARx (2,1) ETH-XRP	VEC (2,1,1) LTC-XRP
DCC (1,1)-N GARCH (1,1)(1,1)-N 9 parameters	AIC	-5.32981	-5.43720	-4.58273	-5.44965	-4.49400	-6.90983
	BIC	-5.09535	-5.20274	-4.34827	-5.21518	-4.25953	-6.67537
DCC (1,1)-N GARCH (1,2)(1,2)-N 11 parameters	AIC	-5.18727	-5.30753	-4.63904	-5.25179	-5.18727	-6.99789
	BIC	-4.90070	-5.02096	-4.35247	-4.96522	-4.90070	-6.71132
DCC (1,1)-N GARCH-GJR (1,1)(1,1)-N 11 parameters	AIC	-5.04424	-5.03983	-4.20340	-5.66795	-5.04424	-6.90779
	BIC	-4.75767	-4.75326	-3.91683	-5.38138	-4.75767	-6.62123
DCC (1,1)-N GARCH (1,1)(1,1)-T 11 parameters	AIC	-4.72177	-5.41345	-4.55023	-5.43080	-4.72177	-6.91375
	BIC	-4.43520	-5.12688	-4.26366	-5.14424	-4.43520	-6.62718
DCC (1,1)-T GARCH (1,1)(1,1)-N 10 parameters	AIC	-5.15871	-5.37596	-4.95013	-5.38244	-5.15871	-7.79028
	BIC	-4.84609	-5.06334	-4.63751	-5.06982	-4.84609	-7.47766
DCC (1,1)-T GARCH (1,2)(1,2)-N 12 parameters	AIC	-5.42140	-5.69634	-4.94523	-5.72676	-5.42140	-7.81549
	BIC	-5.10878	-5.38372	-4.63261	-5.41414	-5.10878	-7.50287
DCC (1,1)-T GARCH-GJR (1,1)(1,1)-N 12 parameters	AIC	-5.37421	-5.44878	-5.03733	-5.68717	-5.37421	-7.83871
	BIC	-5.06159	-5.13616	-4.72471	-5.37455	-5.06159	-7.52609
DCC (1,1)-T	AIC	-4.63530	-5.52304	-4.75893	-5.45947	-5.12060	-7.78349

GARCH (1,1)(1,1)-T 12 parameters	BIC	-4.32268	-5.21042	-4.44631	-5.14685	-4.80798	-7.47087
DCC (1,1)-T GARCH (1,1)(1,2)-T 13 Parameters	AIC	-5.09989	-5.51256	-4.95096	-5.44628	-5.09989	-7.79512
	BIC	-4.76121	-5.17388	-4.61229	-5.10761	-4.76121	-7.45645
DCC (1,1)-T GARCH (1,1)(1,3)-T 14 Parameters	AIC	-5.18892	-5.62104	-4.93303	-5.53890	-5.18892	-7.78376
	BIC	-4.82419	-5.25632	-4.56831	-5.17417	-4.82419	-7.41904
DCC (1,1)-T GARCH (1,1)(1,4)-T 15 Parameters	AIC	-5.05966	-5.60107	-4.91905	-5.51929	-5.05966	-7.77516
	BIC	-4.66889	-5.21029	-4.52827	-5.12852	-4.66889	-7.38439
DCC (1,1)-T GARCH (1,1)(2,3)-T 15 Parameters	AIC	-5.22430	-5.60103	-4.88219	-5.51882	-5.22430	-7.76376
	BIC	-4.83352	-5.21025	-4.49142	-5.12805	-4.83352	-7.37298
DCC (1,1)-T GARCH (1,2)(1,1)-T 13 Parameters	AIC	-5.09511	-5.49978	-4.94864	-5.43874	-5.09511	-7.74716
	BIC	-4.75644	-5.16110	-4.60997	-5.10007	-4.75644	-7.40849
DCC (1,1)-T GARCH (1,2)(1,3)-T 15 Parameters	AIC	-5.17181	-5.60004	-4.91104	-5.51719	-5.17181	-7.76299
	BIC	-4.78103	-5.20926	-4.52026	-5.12641	-4.78103	-7.37221
DCC (1,1)-T GARCH (1,2)(1,4)-T 16 Parameters	AIC	-5.03418	-5.58006	-4.89705	-5.49759	-5.03418	-7.75855
	BIC	-4.61735	-5.16323	-4.48022	-5.08076	-4.61735	-7.34172
DCC (1,1)-T GARCH (1,2)(2,3)-T 16 Parameters	AIC	-5.20640	-5.58002	-4.86016	-5.49711	-5.20640	-7.74299
	BIC	-4.78957	-5.1632	-4.44333	-5.08028	-4.78957	-7.32616
DCC (1,1)-T GARCH-GJR (1,1)(1,1)-T 14 parameters	AIC	-5.34234	-5.42710	-4.99980	-5.55275	-5.34234	-7.74280
	BIC	-4.97761	-5.06238	-4.63508	-5.18803	-4.97761	-7.37808
DCC (1,1)-T GARCH (1,2)(1,2)-T 14 parameters	AIC	-5.07443	-5.15998	-4.92895	-5.42557	-5.07443	-7.76972
	BIC	-4.70971	-4.79525	-4.56423	-5.06084	-4.70971	-7.40499

Table 87: Predictive Akaike and Bayes Information Criteria results for DCC GARCHs models.

Predictive BICs are reported together with predictive AICs but only the latter is considered for model selection. According to both AIC and BIC, the assumption of a multivariate Student's t distribution for the innovations in the DCC equation produces much better fit than the assumption of a multivariate normal distribution. This means that the probability of observing extreme values in the residuals series is larger than under the multivariate normal distribution assumption.

Interestingly, the best fitting model for Bitcoin-Ripple and Litecoin-Ripple is the DCC (1,1)-T GARCH-GJR (1,1)(1,1)-N. The selection of this model suggests that volatility of residuals series increases after negative values are observed. The explanation of this results might be intuitive, because the filtered residuals come from a period of distress in the market, after the coronavirus pandemic burst.

Conclusions

In the attempt of modelling the cryptocurrencies returns, this thesis considered two different approaches: univariate volatility models and multivariate volatility models. The steps carried out in both kind of analysis are similar. Firstly, the set of models considered, and the methodology used to compare estimation results are introduced formally. In the second step the models are estimated and compared to determine the best ones. After this, their fit to data is evaluated with different tests on models' residuals. Finally, the forecasting performances of selected models, estimated on historical data, were compared with those of other models discarded in the first analysis, in order to determine whether the models which better describe past cryptocurrencies' returns movements, would have also produced the best forecasts.

The estimations of univariate models are obtained through Bayesian estimation technique explained and applied by professors Chan and Grant, in their 2015 paper "Modeling Energy Price Dynamics: GARCH versus Stochastic Volatility". Multivariate estimations instead are obtained with more common Maximum Likelihood method.

The estimations of univariate models indicate that the best fit to past data is obtained by assuming fat tailed distributions for the innovations. In fact, among GARCH models, the GARCH-t is preferred for all cryptocurrencies, and the SV-t model showed better performances than other Stochastic Volatility models for Bitcoin, Ethereum and Litecoin. In general, the SV family models performed better than their GARCH counterparts because in SV the log-volatility is assumed to be a random variable. This makes SV models more robust to changes in the time series with respect to GARCH family models.

The residual analysis on univariate models showed that generally GARCHs capture adequately the volatility dynamics in the data, but they do not model properly the autocorrelation. Results from autocorrelation tests are better for the SV models, since both residuals and squared residuals seem to exhibit no serial correlation for all cryptocurrencies. The Kolmogorov-Smirnov tests indicate that, with the only exception of models estimated on Ethereum series, the errors follow a Student-t distribution, as assumed in both the best fitting GARCHs and SVs, according to marginal Likelihoods. It must be taken into account that for SV models, the test statistics of Ljung-Box and McLeod-Li tests, and the degree of freedom estimated for the Student-t distributions have very large standard deviation and the results might be inaccurate for this family of models.

While univariate models are powerful tools, often a multivariate study is more convenient, as it provides clues on the long run relations among assets. The most interesting result is the finding of cointegration among Litecoin's returns and Ripple's returns, even if the fitted VEC model estimates meaningful and significant parameters only for Ripple adjustment coefficient.

The main predictors of cryptocurrencies hypothesized by Catania et al. in 2019 were included in the multivariate models. Among all the variables considered (S&P 500, Nikkei 225, STOXX 600 indexes, US bonds interest rates, Gold and Silver, VIX index), the Nikkei 225 seems to be correlated with cryptocurrencies returns, at least for the period considered, since its coefficients in the VAR and VEC models is always found to be statistically significant. Despite this, often the models without predictors performed better, and are preferred on the basis of their fit to historical data.

Residuals analysis of best multivariate VAR, VARx and VECx show that only the models selected for the couples Bitcoin-Ethereum and Bitcoin-Litecoin can explain the behaviour of cryptocurrencies' returns. For all the remaining couples of timeseries, the residual analysis indicates that there is autocorrelation left in the residuals. Moreover, the VARs and VEC alone are unable to model the volatilities of the series since the McLeod-Li tests show that autocorrelation is left in the squared residuals.

Multivariate GARCH models are fitted to VARs and VEC residuals. Ljung-Box results show that the inclusion of the DCC GARCH does not eliminate the autocorrelation in cryptocurrencies residuals, but as one could expect, it addresses the problem of autocorrelation in squared residuals. The selected DCC GARCH eliminated the autocorrelation in squared residuals of Bitcoin-Ethereum, Bitcoin-Litecoin, Bitcoin-Ripple, Ethereum-Litecoin and Ethereum-Ripple series. Both Kolmogorov-Smirnov tests on marginal residuals series and Baringhaus-Franz test of multivariate distributions show the inadequacy of the innovations' distributions assumptions in these models.

The pseudo forecasting procedures adopted for both univariate and multivariate models allowed to compare the performances of analysed models on out of sample data. Results for univariate models suggest that SV-t is the best performer in both fitting historical data and filtering the out of sample data for Bitcoin and Ethereum series, but that SV (2) would have produced better forecasts than the selected SV-t and SV-MA for Litecoin and Ripple respectively.

For what concerns the multivariate analysis, Predictive AIC results suggest that the best models to filter the "new" sample was never the one selected on historical data, for all cryptocurrencies. This result is strongly influenced by the impact of coronavirus emergency on global markets. In the out of sample data both the cryptocurrencies and the predictors considered dropped in value dramatically and the models selected to fit historical data might not have been able to forecast such large drops.

There is room for several improvements over the analysis implemented in this thesis. The inclusion of univariate conditional mean models could improve the fit to cryptocurrencies' returns series. Moreover, although many flavours of GARCH and SV models are used, there exist many others that might be tested. For what concerns the multivariate models, VARX and VECX could include other predictors. If the correlation of new predictors with cryptocurrencies would be higher, they could lead to better forecasts for the future.

In this thesis only the DCC GARCH is analysed, but there exist other multivariate volatility models that might be better for forecasting. For example, a multivariate SV model could be tested. Moreover, the DCC GARCH analysis could be deepen by considering other univariate GARCH models or by assuming different distributions for the innovations of univariate GARCH components or DCC equation.

References

- Williams Mark T., World Bank Conference Washington DC, 11 November 2014
- Moore Heidi, "Confused about Bitcoin? It's 'the Harlem Shake of currency", The Guardian, 1 March 2014 https://www.theguardian.com/commentisfree/2013/apr/03/bitcoin-currency-bubble-crash-not-rocking-financial-markets
- Catania, Grassi, Ravazzolo, "Forecasting Cryptocurrencies financial timeseries, Predicting the Volatility of Cryptocurrency Time–Series", 2018 https://biopen.bi.no/bi-xmlui/bitstream/handle/11250/2489408/WP_CAMP_5_2018.pdf?sequence=1&isAllowed=y
- R. Bothe, L. Rossini, "Comparing the forecasting of cryptocurrencies by Bayesian time-varying volatility models", 2019 https://www.researchgate.net/publication/335903053_Comparing_the_Forecasting_of_Cryptocurrencies_by_Bayesian_Time-Varying_Volatility_Models
- C. Conrad, A. Custovic, E. Ghysels, "Long- and Short-Term Cryptocurrency Volatility Components: A GARCH-MIDAS Analysis", J. Risk Financial Manag., 2018 https://www.mdpi.com/1911-8074/11/2/23
- T. Walther, T. Klein, E. Bouri "Exogenous drivers of Bitcoin and Cryptocurrency volatility – A mixed data sampling approach to forecasting", Elsevier B.V., 2019 https://www.researchgate.net/publication/335722974_Exogenous_Drivers_of_Bitcoin_and_Cryptocurrency_Volatility_A_Mixed_Data_Sampling_Approach_to_Forecasting
- Werner Kristjanpoller, Marcel C. Minutolo, "A hybrid volatility forecasting framework integrating GARCH, artificial neural network, technical analysis and principal components analysis", Elsevier Ltd. 2018 https://www.researchgate.net/publication/325108806_A_hybrid_volatility_forecasting_framework_integrating_GARCH_Artificial_Neural_network_Technical_Analysis_and_Principal_Components_Analysis
- Christopher Sims, "Macroeconomics and Reality ", Econometrica 1980 https://www.jstor.org/stable/1912017?seq=1
- Jon Matonis, "How Cryptocurrencies Could Upend Banks' Monetary Role", American Banker, 26 May 2013
- M. Szmigiera, Cryptocurrency market capitalization 2013-2019, Statista, 20 January 2020 https://www.statista.com/statistics/730876/cryptocurrency-maket-value/
- M. Szmigiera, "Traders: biggest problems of cryptoexchanges 2018 | Statistic", Statista, 3 September 2018. https://www.statista.com/statistics/596462/biggest-problems-cryptoexchanges/
- M. Szmigiera, "Expectations of traders from cryptocurrency exchanges 2018 | Statistic", Statista, 3 September 2018 https://www.statista.com/statistics/864797/expectations-of-traders-cryptocurrency-exchanges/
- cryptodatadownload.com https://www.cryptodatadownload.com/data/
- Rick Bagshaw, "Top 10 cryptocurrencies by market capitalisation", Coin Rivet, 8 October 2019

https://coinrivet.com/it/top-10-cryptocurrencies-by-market-capitalisation/

- Usman W. Chohan, “Cryptocurrencies: A Brief Thematic Review”, Social Science Research Network, 28 August 2017 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3024330
- "Ether Supply Growth Chart", etherscan.io, 9 January 2020 https://etherscan.io/chart/ethersupplygrowth
- Robert McMillan, "Ex-Googler Gives the World a Better Bitcoin", Wired https://www.wired.com/2013/08/litecoin/
- Ian Steadman, "Wary of Bitcoin? A guide to some other cryptocurrencies", Ars Technica, 11 May 2013 https://www.wired.co.uk/article/alternative-cryptocurrencies-guide
- Matt Scully, "Alternative Money Mover Ripple Labs Enters U.S. Banking System", American Banker, 24 September 2014 https://www.americanbanker.com/news/alternative-money-mover-ripple-labs-enters-us-banking-system
- Tom Simonite, "Big-name investors back effort to build a better Bitcoin", MIT Technology Review, 11 April 2013 https://www.technologyreview.com/2013/04/11/179054/big-name-investors-back-effort-to-build-a-better-bitcoin/
- Chan and Eisenstat, “Marginal Likelihood Estimation with the Cross-Entropy Method”, September 2012 https://joshuachan.org/papers/Chan-Eisenstat%202012.pdf
- Chan and Grant paper “Modeling Energy Price Dynamics: GARCH versus Stochastic Volatility”, November 2015. https://joshuachan.org/papers/energy_GARCH_SV.pdf
- Eliezer Bose, Marilyn Hravnak and Susan M. Sereika, “Vector Autoregressive (VAR) Models and Granger Causality in Time Series Analysis in Nursing Research: Dynamic Changes Among Vital Signs Prior to Cardiorespiratory Instability Events as an Example“, Copyright © 2017 Wolters Kluwer Health, Inc. All rights reserved https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5161241/
- Maddala, "Introduction to Econometrics (3d edition), ch 6.7, and 13. 5 p 528, 2001 and Hayashi, ch. 2.10 "Testing For serial correlation", 2000
- G. Ljung; G. E. Box, “On a Measure of a Lack of Fit in Time Series Models”, Biometrika, 1978 https://www.researchgate.net/publication/246995234_On_a_Measure_of_Lack_of_Fit_in_Time_Series_Models
- “Augmented Dickey–Fuller test”, Wikipedia.org , 27 February 2020 https://en.wikipedia.org/wiki/Augmented_Dickey%E2%80%93Fuller_test
- S. Laurent, L. Bauwens, J. Rombouts, “Multivariate GARCH models: a survey”, Journal of Applied Econometrics, 2006 https://onlinelibrary.wiley.com/doi/epdf/10.1002/jae.842
- Wikipedia, the Free Encyclopedia, “Test di Kolmogorov-Smirnov”, 24 February 2018 https://it.wikipedia.org/wiki/Test_di_Kolmogorov-Smirnov#:~:text=Il%20test%20di%20Kolmogorov%2DSmirnov,oppure%20per%20confrontare%20due%20campioni.
- MathWorks Inc, “varm”, 1994-2020 The MathWorks, Inc.

https://it.mathworks.com/help/econ/varm.html

- Wikipedia, the Free Encyclopedia, “Cointegration”, 17 December 2019 https://en.wikipedia.org/wiki/Cointegration#Engle%E2%80%93Granger_two-step_method
- Munich Personal RePEc Archive, “Stationarity and cointegration tests: Comparison of Engle - Granger and Johansen methodologies”, 1998 https://mpra.ub.uni-muenchen.de/75967/1/MPRA_paper_75967.pdf
- Hänninen, R. (2012). "The Law of One Price in United Kingdom Soft Sawnwood Imports – A Cointegration Approach". Modern Time Series Analysis in Forest Products Markets. Springer. p. 66. ISBN 978-94-011-4772-9
- Wikipedia, the Free Encyclopedia, “Johansen test”, 15 April 2019 https://en.wikipedia.org/wiki/Johansen_test
- © 1994-2020 The MathWorks, Inc, “jcitest documentation”, 2011 https://it.mathworks.com/help/econ/jcitest.html
- © 1994-2020 The MathWorks, Inc, “vecm documentation”, 2017 https://it.mathworks.com/help/econ/vecm.html
- Muglia, Santabarbara and Grassi, “Is Bitcoin a Relevant Predictor of Standard & Poor’s 500?” Journal of Risk and Financial Management. 12. 93. 10.3390/jrfm12020093 https://www.mdpi.com/1911-8074/12/2/93
- Stavroyiannis and Babalos, “Dynamic Properties of the Bitcoin and the US Market”, SSRN Electronic Journal. 10.2139/ssrn.2966998 https://www.researchgate.net/publication/316938705_Dynamic_properties_of_the_Bitcoin_and_the_US_market
- Massimo Guidolin, “Modelling and Forecasting Conditional Covariances: DCC and Multivariate GARCH”, Dept. of Finance, Bocconi University file:///C:/Users/tommy/Downloads/Lecture_6_part_II_(Multivariate_GARCH_only)20130527002724%20(3).pdf
- Elisabeth Orskaug, “Multivariate DCC-GARCH Model -With Various Error Distributions”, NTNU Norwegian University of Science and Technology, June 2009 https://core.ac.uk/download/pdf/52106361.pdf
- Alexios Ghalanos, “The rmgarch models: Background and properties”, Version 1.3-0, September 12, 2019 https://cran.r-project.org/web/packages/rmgarch/vignettes/The_rmgarch_models.pdf
- Dedalo Invest Finance, “Modelli GARCH univariati e multivariati“, ©2014-2020 Dedalo Invest Finance Srl https://www.dedaloinvest.com/education/didattica-investimenti/garch#Modelli%20GARCH%20multivariati
- Jon Danielsson, “Financial Risk Forecasting”, John Wiley and Sons Ltd, 2011 edition
- Tsay, “R. S. Analysis of Financial Time Series. 2nd Ed. Hoboken”, NJ: John Wiley & Sons Inc., 2005, p.33