ERASMUS UNIVERSITEIT ROTTERDAM

Master Thesis Quantitative Finance

Predicting Cryptocurrency Returns in a Time-Varying Volatility Context

---

Tim Stolte (451143)

Supervisor: Dr. Wendun Wang

Second assessor: Dr. Rasmus Lönn

AMDAX Supervisor: Marcel Burger, MSc.

December 31, 2021

Abstract

This paper attempts to predict cryptocurrency returns using multiple Bayesian stochastic volatility models and Markov regime-switching models, both with autoregressive components and exogenous predictor variables. The methods are used to explicitly model the apparent time-varying nature that is present in the volatilities of the series. The results show that accounting for co-movement in the residual volatility innovations enhances the prediction performance of the Bayesian models, but none of these models are able to significantly outperform benchmark econometric ARX and VARX models. However, significant improvement in predicting power is found for Markov regime-switching models that account for volatility changes across regimes.

Keywords: *Cryptocurrencies, Bayesian stochastic volatility, Markov regime-switching*

Contents

1 Introduction

Over the past few years, cryptoassets have developed rapidly and established themselves as a means for the innovative investor to gain exposure to a domain that satisfies many demands. Among these are the search for diversification possibilities and the pursuit of technological advancements. The assets may also provide practitioners with a combination of unusually high returns and low correlations with traditional investment opportunities such as stocks or commodities. Moreover, the fundamental basics of the assets are, in fact, completely different from those of any other asset that has ever been available on the market. We thus argue that pricing the assets requires a different approach than is used for traditional investments.

Greer (1997) defines an asset class as *"a set of assets that bear some fundamental economic similarities to each other, and that have characteristics that make them distinct from other assets that are not part of that class"*. Cryptoassets differ from other assets based on multiple aspects. Most prominently, the manner in which the assets are governed is unique. All cryptoassets are birthed and maintained by developers that use them to fuel a blockchain-based application. Secondly, miners support the network and thus also provide maintenance of the asset itself. Companies that contribute to the availability of the asset to the public, such as centralised exchanges or digital asset managers, form the third and final element of governance. Most other asset classes are characterised by financial institutions taking care of most (if not all) of these tasks. Another significant difference is the supply of the asset. If we, for example, take a look at the currency asset class, central banks hold the power to adjust the money supply. This enables them to increase or decrease the value of the currency that is held by the public. We find similar features for other asset classes, where companies can issue or buy back stocks and bonds, thus controlling their market value and affecting portfolios of investors. In contrast, most cryptoassets are subject to a predetermined supply schedule, as the amount of native units at a certain point in time is common knowledge and cannot be altered by a single entity. Putting these characteristics together, we find that according to the definition of Greer (1997), we are dealing with a new asset class.

Krueckeberg and Scholz (2019) draw the same conclusion, taking the ever-increasing liquidity of this particular market into account. They state that instability of the market is the only thing that is possibly holding back the universal acknowledgement of the asset class. While agreeing with this statement, it is our expectation that, as the volume increases, the volatility of the crypto market will decrease as it has done over time (Bariviera et al., 2017). It is not unthinkable that it will continue to decrease and eventually reach a level similar or even lower than that of the stock market. Until that time, it is even more important to attempt explaining and predicting their returns with care.

Another important topic is whether this newly defined asset class is worth considering in terms of allocation strategies. Dyhrberg (2016) is among the first to publish an academic view on the possibilities that Bitcoin could serve in hedging risk in the US market. She compares the digital coin to gold, regarding its potential to both provide diversification and protect against inflation. These hedging abilities are also documented on a more global spectrum and for coins other than Bitcoin (Bouri et al., 2020). An analysis into the dynamic relationships by cryptocurrencies and other rates and indices, as done by Corbet, Meegan, Larkin, Lucey and Yarovaya (2018), yields a similar conclusion, further highlighting the uniqueness of the asset class. Lastly, Platanakis and Urquhart (2020) find that the addition of Bitcoin increases the portfolio performance considerably across multiple portfolio construction techniques and different levels of risk aversion.

However, before one can optimally profit from the added benefits of crypto investments, inferences need to be made on the behaviour of their returns. In other words, we need to identify driving factors and explanatory patterns in order to determine when to adjust allocations. Existing literature shows that, thus far, little is known on this subject. Bianchi (2020) finds that macroeconomic factors do not explain dynamics in the prices of cryptocurrencies, nor is there evidence of volatility spillovers with respect to other asset classes. Li and Yi (2019) attempt to uncover evidence of a factor structure in the crypto market, but conclude that risk premia are insignificant and that idiosyncratic noise is too large. Evidence of a non-linear Granger causality between Bitcoin's trading volume and returns is discovered by Balcilar et al. (2017), which is particularly useful for short-term trading strategies.

Recognising the lack of knowledge about cryptoassets in portfolio management, we seek to fill the gap by attempting to explain cryptocurrency returns and selecting the tools that may allow us to do so. In this paper, we opt for a volatility-based approach. Lahmiri et al. (2018) and Zhang et al. (2018) show that Bitcoin and other digital coins display strong volatility clustering and long-range volatility memory. We try to take advantage of these features by simultaneously forming the hypothesis that prices behave differently in high-volatility regimes than they do in low-volatility regimes. This is in line with the asymmetric volatility documented by Baur and Dimpfl (2018), who attribute their findings to the 'fear of missing out' (FOMO) that is especially present among crypto market participants. We also relate the importance of the public FOMO by the positive correlation between Bitcoin-related Twitter and Google activity and the Bitcoin price (Georgoula et al., 2015; Matta et al., 2015).

In our methodology, we make use of two types of time-varying volatility models that are similar to ones that show signs of outperformance in predicting either cryptoassets or traditional assets. The first model is a Bayesian vector autoregressive model with exogenous input and stochastic volatility as proposed by Bohte and Rossini (2019) in the search for cryptocurrency forecasting power. They also introduce a model that uses a t-distribution to allow for fat tails in the residuals, which shows promising results. Nevertheless, in their research, they fail to compare these methods with more traditional econometric models. We fit both models to our data (allowing for both correlated and uncorrelated residual variances) and expect that comparisons with benchmark ARX and VARX models based on various performance measures provide more knowledge of their value. The second type of model that we consider is a Markov switching autoregressive (MS-ARX) model. This approach uses time-varying volatility in a more straightforward manner and assumes that a predetermined set of parameters can vary across multiple estimated states. It has been proven useful by existing literature for predicting stock returns (Schaller and Norden, 1997) and foreign exchange returns (Dueker and Neely, 2007). Markov switching models have also been extensively used in cryptocurrency volatility forecasting (Ma et al., 2020; Poyser, 2018), often in combination with GARCH models (Tan et al., 2021; Caporale and Zekokh, 2019). Although useful, these GARCH applications are not focused on return modelling and are thus out of scope for this paper. We also assess the performance of the MS-ARX model by comparing it with our benchmark models.

One advantage of our methods compared to more standard models for time-varying volatility is that we are able to identify the dynamics of the underlying volatility process while simultaneously gaining knowledge of the relation between our predictor variables and the dependent returns. If we look at for example GARCH models in this particular context, we would first have to obtain a residual series, only to perform the estimation of the GARCH model afterwards. In this way, allowing for time-variation of the model volatility will not affect the forecasting of the return series, which is our ultimate goal. Chung et al. (2012) do propose a way to incorporate a GARCH specification in a Bayesian VAR model. Nevertheless, Bohte and Rossini (2019) conclude that this does not yield additional forecasting power. Clark and Ravazzolo (2015) even state that VAR models with a GARCH specification perform significantly worse than their counterparts. Fernández-Villaverde and Rubio-Ramírez (2010) relate this shortcoming to the fact that GARCH does not distinguish between level and volatility shocks, making it less flexible than a stochastic volatility approach. We therefore do not consider any GARCH equations in our research.

Between our multiple VARX-SV type models and the MS-ARX approach, we also come across theoretical differences. The main advantage of the former, especially the versions with correlated residual variance innovations, is that these are able to account for any co-movement between the cryptocurrency returns. The MS-ARX model estimates all of its parameters for each equation seperately. On the other hand, the MS-ARX model has the ability to capture any level shifts in the returns. Financial markets are typically subject to bull and bear markets, which respectively relate to rising and falling prices. This would indicate that return series contain level shifts across different regimes. The resulting time-varying feature in the first moment of the returns may give the MS-ARX an edge over other models that assume a constant mean over the full sample period.

Turning to the question of which predictor variables to choose, we note that one largely overlooked advantage in cryptocurrency research is the amount of information that can be extracted from the underlying blockchain network. Every transaction that is carried out using a specific coin is executed on the blockchain that corresponds with that particular coin. It is thus possible to trace back each occasion on which coins traded hands. This paves the way for a wide range of on-chain metrics, containing for example the number of active accounts and the total value that is transmitted on the network. One of such metrics is the hashrate: the total amount of computing power that is generated by the miners on the blockchain at any point in time. The higher the hashrate, the greater the security of the blockchain and the more resistant it is to attacks, which could possibly drive up the price. Hayes (2017) and Aoyagi and Hattori (2019), among others, apply and claim to prove this hypothesis for Bitcoin. Nevertheless, Fantazzini and Kolodin (2020) refute these findings and conclude that there is no evidence of Granger-causality going from the hashrate to the price. Shanaev et al. (2019) also state that positive causality between prices and hashrates are false due to serial correlation and endogeneity. This shows that, contrary to other metrics, the hashrate is studied rather extensively in existing literature. In our methodology, we hope to satisfy the need for other use cases of on-chain data. Bohte and Rossini (2019) use macroeconomic variables, such as indexes, commodity prices and bond rates. These series can be suitable for pricing traditional assets, but show little to no correlation with prices of cryptoassets. Another contribution of this paper to existing research, therefore, is the revision of the set of exogenous variables in predicting cryptocurrency returns and the role that on-chain data can play in this context.

Our results show that our predictor variables have predictive value to some extent, as most of them are statistically significant in one or more of our econometric models. Moreover, all of our models show highly time-varying standard deviation in the residuals. This highlights the relevance of our stochastic volatility and regime switching methods. In terms of performance, we find that the Markov switching approach has the most predictive power, showing significant superiority over all other models based on each of our performance measures. The stochastic volatility models appear to be unable to outperform the benchmark ARX and VARX models for all series. Nevertheless, our results show a small but significant improvement when modelling correlated stochastic volatility innovations.

The remainder of this paper is structured as follows. Section 2 describes the cryptocurrencies and introduces both economic variables and on-chain metrics that are used in our research, after which Section 3 discusses the methods used to predict the cryptocurrency returns. Section 4 presents our findings with respect to our models' time-varying residual variances and their prediction performance. Finally, Section 5 contains some concluding remarks.

2 Data

In our research, we focus on a set of three cryptocurrencies that are available to the public: Bitcoin (BTC), Ether (ETH) and Litecoin (LTC), of which we extract daily prices from CoinMarketCap[1]. We choose these coins based on both their market capitalisation and age. Some other coins are larger than Litecoin in market capitalisation, but have a far shorter lifespan. We believe that the dynamics of these particular coins are mostly driven by public sentiment and that their 'adolescence' would make them unsuitable for correct interpretation by econometric models. It can be argued that our selection forms the establishment of the crypto asset class and has been largely stripped down to their fundamental basics over the course of their existence.

It is important to note that Bitcoin had and still has the largest market capitalisation of all cryptocurrencies ever. Until early 2017, more than 95% of the total crypto market value was invested in Bitcoin.[2] This percentage has decreased since then, but never fell below 38%, which it reached in early 2018. Still, the second largest coin Ether only captured 24% of the market at that time. Bitcoin's market capitalisation has since recovered and takes up around half of the total market value and equals almost three times the market value of runner-up Ether at the end of our sample period.

Figure 1 shows the prices of each cryptocurrency over a period ranging from 28 April 2013 until 31 August 2021 on a logarithmic scale. Note that Ether was birthed after the start of this

[1] Cryptocurrency prices, charts and market capitalizations. CoinMarketCap. (2013). Retrieved September 1, 2021, from https://coinmarketcap.com/

[2] Crypto market cap and defi market cap charts. TradingView. (2018). Retrieved September 1, 2021, from https://www.tradingview.com/markets/cryptocurrencies/global-charts/

period. Not only do we observe high correlation across the series, but we also see sharp peaks in at the end of the year in 2013 and 2017, almost exactly four years apart. These speculative bubbles have been analysed to great extent in the past (Cheah and Fry, 2015; Corbet, Lucey and Yarovaya, 2018). We attribute these bubbles to the so-called 'halving effect' (Meynkhard, 2019).

(a) BTC

(b) ETH

(c) LTC

Figure 1: Prices of the three cryptocurrencies, with the Bitcoin halving dates in red.

By construction, every 210,000 blocks in the Bitcoin blockchain, the mining reward is halved. This means that miners only receive half the amount of Bitcoin that they would have received up to that point. After each halving, the inflation rate automatically decreases because of the reduced rate of issuance of new coins. This reduces the supply and thus increases the price. Since the crypto market was still emerging in the past years, this change in inflation was not efficiently priced into the value of the asset. Hence, we find that such sudden increases lead to speculative bubbles. The Bitcoin dominance in the market causes this rally to be present in the prices of other cryptocurrencies as well. In the original Bitcoin white paper (Nakamoto, 2019), it is assumed that a block would be added to the blockchain every ten minutes. The amount of 210,000 blocks is chosen by Bitcoin's creator(s) such that halvings would take place approximately every four years, which lines up with the forming of bubbles.

Because of the impact that the Bitcoin halving has on the prices of all cryptoassets, we find that it is mindful to take this into consideration when selecting a sample period for our research. Hence, we select a sample period ranging from 1 April 2018 until 31 August 2021, resulting in a total number of 1249 daily observations. This is a period in which all of our coins have mostly recovered from the bubble in late 2017 and are most probably least affected by inflated public sentiment. We do note that the most recent Bitcoin halving of March 2020 may have given rise to another speculative bubble, seeing the increasing trend at the end of the sample period and the crash on 19 May 2021. Nevertheless, the returns before and after this crash show relatively low volatility compared to the ones around previous bubbles. We therefore choose to also consider these last months.

2.1 Cryptocurrency returns

Now that we have defined our sample period, we can dive deeper into the statistics and dynamics of our selection of coins in this particular period. We start by transforming the daily prices of each coin into log returns:

$$y_{it} = \log \frac{P_{i,t}}{P_{i,t-1}},$$

where $P_{i,t}$ denotes the price of coin i at time t. These returns are displayed in Figure 2. We immediately observe some volatility clustering, looking at the small absolute returns in for example the periods of late 2019 and mid 2020 and large positive and negative peaks around the beginning of 2021. However, the most notable observation corresponds to the one on 12 March 2020, where all three reached their lowest return in the sample period. On this day, the price of Bitcoin on one of the largest crypto derivative exchanges plunged almost 10% lower than on other exchanges, resulting in a crash in the total market.[3] This supports the statement made earlier about the large dominance of Bitcoin and also reflects the 'growing pains' of the emerging asset class. The large crash on 19 May that ended the period of increasing price movements in the first few months of 2021 is also visible in the return series, especially for Litecoin.

From the summary statistics of the log returns in Table 1, we conclude that in our sample period, the three coins show similar characteristics to traditional assets. The average log returns are very close to zero, the kurtoses are high, indicating heavy tails, and all series are negatively skewed. It is also clear that Bitcoin's volatility is noticeably lower than that of the other coins. Moreover, all three coins show cross-correlations of well over 80% with respect to each other

[3] Bambrough, B. (Ed.). (2020, March 19). Here's what caused Bitcoin's 'extreme' price plunge. Forbes. Retrieved December 28, 2021, from https://www.forbes.com/sites/billybambrough/2020/03/19/major-bitcoin-exchange-bitmex-has-a-serious-problem/?sh=24c15d754f7d

over the full sample period.

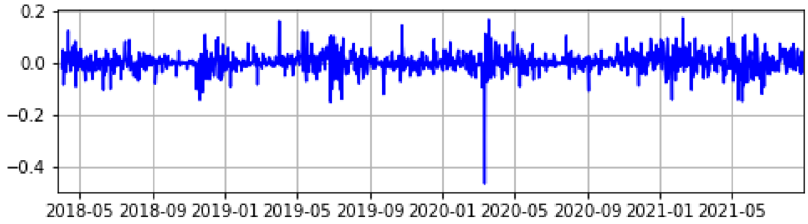

(a) BTC

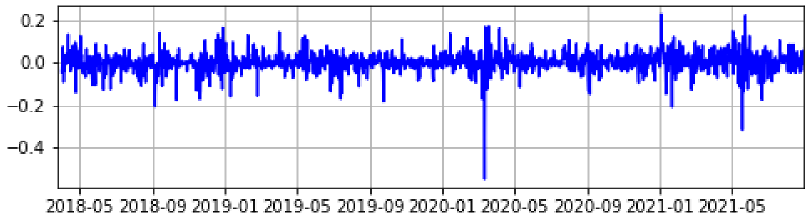

(b) ETH

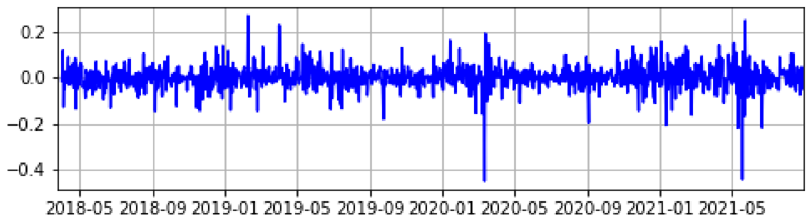

(c) LTC

Figure 2: Log returns of the three cryptocurrencies from 1 April 2018 until 31 August 2021.

Table 1: Descriptive statistics of the log returns of the three cryptocurrencies

	BTC	ETH	LTC
Mean	0.002	0.002	0.000
St Dev	0.039	0.052	0.053
Minimum	-0.465	-0.551	-0.449
Maximum	0.172	0.231	0.269
Skewness	-1.311	-1.255	-0.846
Kurtosis	18.502	13.066	9.594

This table shows the mean, standard deviation, minimum, maximum, skewness and kurtosis of the log returns of the three cryptocurrencies in the sample period 1 April 2018 - 31 August 2021.

Regarding any time-variation in the first three moments of our return series, we refer to the rolling window estimates in Figure A1 in the Appendix. We observe very strong variation in the mean and standard deviation. The latter observation supports our choice of methodology and the former specifically hints at the possible advantage of the regime-switching mean in the MS-ARX model. The skewness is less time-varying, but we do notice a significant period of exceptionally low skewness in between March and June of 2020. We leave the possible implications of this feature on the forecasting power of certain methods for further research.

2.2 Economic variables

Similar to existing research in cryptocurrencies (Bianchi, 2020; Catania et al., 2019), we exploit the possible predictive power of multiple economic variables. Even though we take a reserved stance on the use of such variables due to low correlations, we deem it useful to include them to gain insights in the states of the overall economy and/or investor behaviour. Since we define crypto as a new asset class, there must be an incentive for practitioners to allocate some capital to it. We hope to find explanations for such incentives in the variables that we introduce in this subsection.

To proxy for the state of the global economy, we consider the S&P 500 index. We aim to explain investor behaviour by including the 1-year and 10-year US bond rates, the price of gold futures and the DXY index. The latter is an index that tracks the value of the US dollar relative to other large currencies. Bitcoin is known to be used as a means to hedge inflation risk, similar to the way gold is used in traditional portfolios. This relationship motivates the need for such variables. We obtain daily prices from Investing.com and transform them into log returns. Contrary to the crypto market, traditional markets are not open 24/7. We therefore fill the missing days with the last valid value in the series. In the end, we also standardise each of the series. Figure A2 in the Appendix contains plots of the five transformed variables.

2.3 On-chain metrics

An important part of our research is the use of data that is saved on the underlying blockchains of cryptocurrencies. While it is possible to access this data directly, there are many providers that gather this data and present it as open source to the public. We extract our on-chain from Glassnode[4].

In particular, we use three on-chain metrics. Note that the series are unique for each coin, resulting in a total of nine on-chain metric series. We briefly describe them and discuss the manner in which we expect them to impact the returns below:

1. *Fee Ratio Multiple (FRM)*

   This metric is defined as the ratio between the total miner revenue and total transaction fees. The former consists of both block rewards and transaction fees. One can argue that a high FRM is caused by relatively low transaction fees, indicating that the blockchain is not as actively used as it should be. This signals potential insecurity of the blockchain,

[4] On-chain market intelligence. Glassnode. (2018). Retrieved September 1, 2021, from http://www.glassnode.com/

which in turn has a negative effect on the token value.

2. *Market Value to Realised Value (MVRV)*
   The MVRV ratio is the ratio between the market capitalisation and the realised capitalisation. The latter is an alternative to the market capitalisation and uses the price of each individual coin at the time it last moved on the blockchain instead of the most recent price for all coins. This ratio can be used to identify whether the asset is overvalued or undervalued. For instance, the more the market capitalisation exceeds the realised capitalisation, the higher the price is relative to what current investors have entered their positions for, the more profit can be realised and vice versa. Signs of overvaluation would generally cause an increase in selling pressure in most markets. However, especially since this particular metric is highly popularised by crypto investors, it is possible that the MVRV ratio can be used as a measure that captures investor FOMO.

3. *Network Value to Transactions (NVT)*
   This is the ratio of the market capitalisation and the total volume that is transmitted on the network on a single day. We relate this ratio to the price-to-earnings ratio that is widely used in traditional finance and it can also be seen as an expression of inverse monetary velocity. We expect this metric to positively affect the price.

Before we can use these metrics, we need to remove any deterministic components from them. Both the FRM and the NVT ratio appear to contain weekly seasonality for all three series. We find that trading volume is significantly lower during weekends, causing both metrics to sharply increase on Saturdays and Sundays, as seen in Table 2. As a result, the null-hypothesis of unit root presence is not rejected by means of an Augmented Dickey-Fuller test with a significance level of 5% for all six series. We remove the unit roots by taking seventh-order differences and are left with stationary series. The Augmented Dickey-Fuller test also fails to reject the null-hypothesis for the MVRV ratio for all coins. Taking the first differences solves this problem.

One important issue to note is that using the on-chain metrics as is would result in endogeneity within an econometric model. Table 3 shows that the MVRV ratio deltas are highly cross-correlated with the log returns of the assets. We also find that the returns explain a significant part of the metric. In order to solve this problem we rely on an instrumental variable approach, also known as two-stage least squares (2SLS). This method requires us to choose an instrumental variable that is highly correlated with the endogenous variable conditionally on all other predictor variables and only affects the returns through the predictor variables. We use the realised capitalisation as such a variable, since it shows both of the aforementioned prop-

Table 2: Means of two on-chain metrics on each day of the week

	BTC		ETH		LTC	
	FRM	NVT	FRM	NVT	FRM	NVT
Monday	48.498	28.582	23.311	47.971	637.340	29.197
Tuesday	45.093	27.680	22.925	48.433	616.057	29.935
Wednesday	44.068	27.833	22.363	48.232	634.998	29.951
Thursday	44.225	27.480	21.874	47.523	635.746	29.534
Friday	43.649	27.321	22.223	47.992	638.280	30.304
Saturday	59.873	40.798	27.246	71.934	715.825	39.379
Sunday	68.850	44.151	28.558	75.041	742.669	41.952

This table shows the means of the raw FRM and NVT ratio values on each day of the week for the three cryptocurrencies in the sample period 1 April 2018 - 31 August 2021.

erties. We apply standardisation to the on-chain metrics as done previously with the economic variables.

Table 3: Signs of endogeneity for the first differences of the MVRV ratio

	BTC	ETH	LTC
Correlation coefficient	0.513	0.528	0.536
β	13.227	10.220	10.049
p-value	0.000	0.000	0.000
R^2	0.263	0.279	0.287

This table shows the correlation coefficients of the MVRV ratio and the log returns of all three cryptocurrencies. It also contains the coefficient, corresponding p-value and the coefficient of determination (R^2) of the model $\Delta MVRV_t = \alpha + \beta \times \log r_t + \varepsilon_t$ obtained by OLS estimation. All calculations are performed using the sample period 1 April 2018 - 31 August 2021.

The first step in our 2SLS procedure calls for regression of the endogenous variable on the instrumental variable and all predictor variables in our model. In these regressions we document significant effects of the realised caps on the MVRV ratios keeping all other regressors constant, which supports our choice of instrumental variables. Even though there are signs of non-stationarity for the realised cap deltas for BTC and ETH (p-values of 0.024 and 0.015 respectively), we find that the residuals of the three regressions are stationary at a 1% significance level. This means that either the realised cap deltas are actually stationary or the two series cointegrate. The fitted values, which we will use in our models as an exogenous proxy for the MVRV ratios, are also found to be stationary at a 1% significance level.

3 Methodology

There are many ways in which we can incorporate our set of variables in econometric models in orders to predict cryptocurrency returns. We adopt methods that approach the problem from a volatility point of view. In this section, we describe our proposed models in more detail.

3.1 Stochastic Volatility VARX Model

Given the n series of cryptocurrency log returns $y_t = (y_{1,t}, \dots, y_{n,t})'$ and m series of explanatory variables $x_t = (x_{1,t}, \dots, x_{m,t})'$, we define the SV-VARX model as follows[5]:

$$y_t = \alpha + \sum_{j=1}^{p} B_j y_{t-j} + \Gamma x_{t-1} + \eta_t, \tag{1}$$

$$\eta_t = A^{-1} \Lambda_t^{0.5} \varepsilon_t, \qquad \varepsilon_t \sim \text{i.i.d. } N(0, I_n), \tag{2}$$

$$\Lambda_t = \text{diag}(\lambda_{1,t}, \dots, \lambda_{n,t}), \tag{3}$$

$$\log \lambda_{i,t} = \log \lambda_{i,t-1} + \nu_{i,t}, \qquad i = 1, \dots, n, \tag{4}$$

$$\nu_{i,t} \sim \text{i.i.d. } N(0, \phi_i^2), \qquad i = 1, \dots, n, \tag{5}$$

where A is a upper triangular matrix with ones on the diagonal and $\frac{n(n-1)}{2}$ non-zero coefficients below it. We denote the VARX parameters to be estimated as the $(n \times (1 + np + m))$ matrix $\Theta = (\alpha, B, \Gamma)'$, with $B = (B_1, \dots, B_p)$. Note that we only consider past values of every predictor variable, such that we keep in mind the application possibilities of this and every other of our models for real-world investing strategies.

Our prior specification, inspired by Clark and Ravazzolo (2015) and Cogley and Sargent (2005), implies a Minnesota prior for the B parameters and a flat prior for the Γ parameters. The remaining priors are the following[6]:

$$(a_{2,1}, \dots, a_{n,n-1}) \sim N\left(0, 1000^2 \times I_{\frac{n(n-1)}{2}}\right), \tag{6}$$

$$\log \lambda_{i,0} \sim N\left(\log \hat{\lambda}_{i,0,OLS}, 4\right), \tag{7}$$

$$\phi_i^2 \sim IG\left(\frac{n+1}{2}, \frac{0.01^2}{2}\right). \tag{8}$$

The initial log variance of each equation $\log \hat{\lambda}_{i,0,OLS}$ is obtained by taking the natural logarithm of the residual variances from ARX(p) models.

We estimate the model parameters by means of a Metropolis-within-Gibbs sampler, which is often used for stochastic volatility models (Jacquier et al., 1999). The first simulation run starts by initialization according to Equations 6-8. Then we apply MCMC algorithm that consists of

[5] Note that we can also capture the log variance equations in vector notation. We then have $\log \lambda_t = \log \lambda_{t-1} + \nu_t, \nu_t \sim$ i.i.d. $N(0, \Phi)$ with Φ restricted to be a diagonal matrix.

[6] In order to avoid inconsistencies that often arise in the parametrisation of the Inverse Gamma distribution in Bayesian research, we explicitly use the notation $IG(\alpha, \beta)$ with α and β as the shape and scale parameters respectively. They relate to the degrees of freedom ν and scaling parameter τ^2 of the often used scaled inverse chi-squared distribution as $\alpha = \frac{d}{2}$ and $\beta = \frac{d\tau^2}{2}$. Also note that this implies that if a univariate stochastic variable follows an Inverse Wishart distribution $x \sim IW(\Psi, \nu)$, then $x \sim IG(\frac{\nu}{2}, \frac{\Psi}{2})$.

the following steps:

1. Draw from the full conditional posterior distribution

$$
\begin{aligned}
p(\text{vec}(\Theta)|.) &\sim N\left(\mu_\Theta, \Omega_\Theta\right), \\
\Omega_\Theta &= \left(\sum_{t=1}^{T} \left(\Sigma_t^{-1} \otimes X_t X_t' \right) + V^{-1} \right)^{-1}, \\
\mu_\Theta &= \Omega_\Theta \left(\text{vec}\left(\sum_{t=1}^{T} \Sigma_t^{-1} y_t X_t' \right) + V^{-1}\text{vec}(P) \right).
\end{aligned}
$$

   Here, $X_t = (\iota, y_{t-i}, \ldots, y_{t-p}, x_t)'$ and the residual covariance matrix $\Sigma_t = Var\,[\eta_t] = A^{-1}\Lambda_t(A^{-1})'$. vec$(P)$ denotes the prior mean and V the diagonal prior covariance matrix of vec(Θ):

$$
V^{(i,i)} = \begin{cases}
\frac{\pi_1}{l} \text{ if the } i\text{-th regressor is the } l\text{-th lag of the dependent variable in that equation} \\
\frac{\pi_2 s_j^2}{l s_k^2} \quad \text{if the } i\text{-th regressor in equation } j \text{ is the } l\text{-th lag of an independent} \\
\quad\quad \text{variable in the } k\text{-th equation} \\
\frac{1}{\pi_3^2} \quad \text{if the } i\text{-th regressor is a on-chain metric of an independent variable in} \\
\quad \text{another equation} \\
\pi_3 s_j^2 \text{ if the } i\text{-th regressor in equation } j \text{ is other}
\end{cases}
$$

   where s_j^2 are the residuals variances of univariate ARX(p) models. We set the hyperparameters $\pi_1 = 0.05$, $\pi_2 = 0.005$, and $\pi_3 = 100,000$ following common settings as described in Kadiyala and Karlsson (1997). The prior variance setting for on-chain metrics corresponding to the independent variable of another equation reflects our attempt to enforce parameter restrictions. P is a vector of zeros in our case.

2. Conditional on Θ and Λ_t, use $\varepsilon_t \sim$ i.i.d. $N(0, I_n)$ in the system of equations $\Lambda_t^{-0.5} A\hat{\eta}_t = \varepsilon_t$ to extract the lower diagonal elements of A by means of transformed regressions as in Cogley and Sargent (2005):

$$
\frac{\hat{\eta}_{i,t}}{\sqrt{\lambda_{i,t}}} = -\sum_{j=1}^{i-1} a_{i,j} \frac{\hat{\eta}_{j,t}}{\sqrt{\lambda_{i,t}}} + \varepsilon_{i,t}, \qquad i = 2, \ldots, n,
$$

   The posterior distribution from which we draw the $a_{i,j}$ coefficients uses the prior set in Equation (6).

3. Draw the stochastic variances Λ_t using a Metropolis step. Due to our univariate specifica-

tion of the log variance innovations, we can consider each $\lambda_{i,t}$ independently in this step. Jacquier et al. (2004) give the corresponding conditional posterior:

$$
\begin{aligned}
p(\lambda_{i,t}|\lambda_{i,t-1}, \lambda_{i,t+1}, \eta_{i,t}, \phi_i^2) &\propto p(\eta_{i,t}|\lambda_{i,t})p(\lambda_{i,t}|\lambda_{i,t-1})p(\lambda_{i,t+1}|\lambda_{i,t}) \\
&\propto \lambda_{i,t}^{-1.5} \exp\left(-\frac{\eta_{i,t}^2}{2\lambda_{i,t}}\right) \exp\left(-\frac{(\log\lambda_{i,t} - \mu_{i,t})^2}{\phi_i^2}\right),
\end{aligned}
$$

with[7] $\mu_{it} = \frac{1}{2}\left(\log\lambda_{i,t-1} + \log\lambda_{i,t+1}\right)$. Since this density function is of unknown form, we rely on a Metropolis sampler. Our candidate-generating density function is

$$
q(\lambda_{i,t}|\lambda_{i,t-1}, \lambda_{i,t+1}, \phi_i^2) \propto \lambda_{i,t}^{-1} \exp\left(-\frac{(\log\lambda_{i,t} - \mu_{i,t})^2}{\phi_i^2}\right),
$$

which assumes that $\lambda_{i,t}$ follows a log-normal distribution with parameters $\mu_{i,t}$ and $\frac{1}{2}\phi_i^2$. We sample from the known density $q(\lambda_{i,t}|.)$ once in each iteration and accept or reject the draw based on the density $p(\lambda_{i,t}|.)$. The acceptance probability of the m-th draw is then given by

$$
\begin{aligned}
\alpha_m &= \min\left(\frac{p(\lambda_{i,t}^m|.)q(\lambda_{i,t}^{m-1}|.)}{p(\lambda_{i,t}^{m-1}|.)q(\lambda_{i,t}^m|.)}, 1\right) \\
&= \min\left(\frac{(\lambda_{i,t}^m)^{-0.5}\exp\left(-\frac{\eta_{i,t}^2}{2\lambda_{i,t}^m}\right)}{(\lambda_{i,t}^{m-1})^{-0.5}\exp\left(-\frac{\eta_{i,t}^2}{2\lambda_{i,t}^{m-1}}\right)}, 1\right) \\
&= \min\left(\sqrt{\frac{\lambda_{i,t}^{m-1}}{\lambda_{i,t}^m}} \exp\left(\frac{1}{2}\left(\frac{\eta_{i,t}^2}{\lambda_{i,t}^{m-1}} - \frac{\eta_{i,t}^2}{\lambda_{i,t}^m}\right)\right), 1\right).
\end{aligned}
$$

4. For each equation, draw ϕ_i^2 from $IG\left(\frac{n+1+T}{2}, \frac{0.01^2+\sum_{t=1}^T \nu_{i,t}^2}{2}\right)$, with $\nu_{i,t} = \log\hat{\lambda}_{i,t} - \log\hat{\lambda}_{i,t-1}$ conditional on the results of the previous step.

It is also possible to generalise the specification of the log variance innovations ν_t such that we model any volatility co-movement across our returns series, resulting in a correlated stochastic volatility VARX model (CSV-VARX). We then assume that the n-dimensional time series ν_t follows a multivariate normal distribution with covariance matrix Φ, with prior specification $\Phi \sim IW(n+1, 0.01^2 \times I_n)$. Consequently, the last step requires drawing from $IW(n+1+T, 0.01^2 \times I_n + \nu_t\nu_t')$. The stochastic variances $\lambda_{i,t}$ for each equation i are to be drawn conditional on the variances of the other equations (denoted with subscript $-i$) using a Metropolis sampler.

[7] The edge cases are different. The prior in Equation 7 comes into play for $\mu_{i,1} = \frac{1}{2}\left(\log\lambda_{i,0} + \log\lambda_{i,2}\right)$ and at the end of the sample we apply $\mu_{i,T} = \log\lambda_{i,T-1}$.

The conditional posterior now reads:

$$
\begin{aligned}
p(\lambda_{i,t}|\lambda_{i,t-1},\lambda_{i,t+1},\lambda_{-i,t},\eta_{i,t},\Phi_i^2) &\propto p(\eta_{i,t}|\lambda_{i,t})p(\lambda_{i,t}|\lambda_{i,t-1})p(\lambda_{i,t+1}|\lambda_{i,t})p(\lambda_{-i,t}|\lambda_{i,t})\\
&\propto \lambda_{i,t}^{-1.5}\exp\left(-\frac{\eta_{i,t}^2}{2\lambda_{i,t}}\right)\exp\left(-\frac{(\log\lambda_{i,t}-\mu_{i,t})^2}{\Phi_{i,i}}\right)\\
&\quad\prod_{j=1,j\neq i}^{n}\exp\left(-\frac{\left(\log\lambda_{j,t}-\left(\mu_{j,t}+\frac{\Phi_{i,j}}{\Phi_{i,i}}(\log\lambda_{i,t}-\mu_{i,t})\right)\right)^2}{\Phi_{j,j}-\frac{\Phi_{i,j}^2}{\Phi_{i,i}}}\right)\\
&\propto \lambda_{i,t}^{-1.5}\exp\left(-\frac{\eta_{i,t}^2}{2\lambda_{i,t}}\right)\exp\left(-\frac{(\log\lambda_{i,t}-\mu_{i,t})^2}{\Phi_{i,i}}\right)\\
&\quad\exp\left(-\sum_{j=1,j\neq i}^{n}\frac{\left(\log\lambda_{j,t}-\left(\mu_{j,t}+\frac{\Phi_{i,j}}{\Phi_{i,i}}(\log\lambda_{i,t}-\mu_{i,t})\right)\right)^2}{\Phi_{j,j}-\frac{\Phi_{i,j}^2}{\Phi_{i,i}}}\right),
\end{aligned}
$$

in which the product density function $p(\lambda_{-i,t}|\lambda_{i,t})$ is derived from the conditional distribution of bivariate lognormal distributions. The acceptance probability is derived as before.

We initialize both stochastic volatility models using the first 25% of our full sample period. We then use 75% of the remaining period to fit the models to our data, after which we apply our performance measures to the remaining data points.

3.2 Stochastic Volatility VARX Model with fat tails

Allowing for fat tails requires modelling the residuals η_t differently than before. The SVt-VARX model is as follows:

$$y_t = \alpha + \sum_{i=1}^{p} B_j y_{t-j} + \Gamma x_{t-1} + \eta_t, \tag{9}$$

$$\eta_t = A^{-1}Q_t^{0.5}\Lambda_t^{0.5}\varepsilon_t, \qquad \varepsilon_t \sim \text{i.i.d. } N(0, I_n), \tag{10}$$

$$\Lambda_t = \text{diag}(\lambda_{1,t},\dots,\lambda_{n,t}), \tag{11}$$

$$\log\lambda_{i,t} = \log\lambda_{i,t-1} + \nu_{i,t}, \qquad i = 1,\dots,n, \tag{12}$$

$$\nu_{i,t} \sim \text{i.i.d. } N(0,\phi_i^2), \qquad i = 1,\dots,n, \tag{13}$$

$$Q_t = \text{diag}(q_{1,t},\dots,q_{n,t}), \tag{14}$$

$$\frac{d_i}{q_{i,t}} \sim \text{i.i.d. } \chi^2_{d_i}. \tag{15}$$

Note that only the last two equations differ from the specification in the previous subsection. This particular modelling of Q_t implies that η_t follow a t-distribution when conditioned on Λ_t.

In terms of simulation, we adjust our previously described MCMC algorithm to account for the newly introduced parameters. This includes the use of an alternative residual covariance matrix in the first step: $\Sigma_t = A^{-1} Q_t^{0.5} \Lambda_t Q_t^{0.5} (A^{-1})'$. The regressions in the second step change to

$$\lambda_{i,t}^{-0.5} q_{i,t}^{-0.5} \hat{\eta}_{i,t} = -\sum_{j=1}^{i-1} a_{i,j} \lambda_{i,t}^{-0.5} q_{i,t}^{-0.5} \hat{\eta}_{j,t} + \varepsilon_{i,t}, \qquad i = 2, \ldots, n. \tag{16}$$

Lastly, we add an additional step after the fourth step to draw q_{it} from independent distributions using the orthogonal residuals $\xi_t = A\hat{\eta}_t$:

$$p(q_{i,t} | \eta_t, \Lambda_t, d_i) \sim IG\left(\frac{d_i + 1}{2}, \frac{d_i + \frac{\xi_{i,t}^2}{\lambda_{i,t}}}{2} \right). \tag{17}$$

Seeing that the model in Clark and Ravazzolo (2015) in which the degrees of freedom d_i are estimated does not yield significantly better results than the model in which d_i is fixed, we do not account for this parameter uncertainty in our specification. This enhances simplicity and sharply reduces computation time. We tried multiple settings and found that fixing d_i at 10 yields the best results.

As described for the stochastic volatility model in the previous subsection, we also use a version of the fat-tailed model with correlated log variances (CSVt-VARX).

3.3 Markov Switching ARX Model

The Markov Switching ARX (MS-ARX) model also assumes that the log returns depend on its lagged values and any exogenous variables:

$$y_{i,t} = \alpha_{i,s} + \sum_{j=1}^{p} \beta_{i,j,s} y_{t-j} + \Gamma_{i,s} x_{t-1} + \varepsilon_{i,s,t}, \qquad \varepsilon_{i,s,t} \sim \text{i.i.d. } N(0, \sigma_{i,s}^2), \qquad i = 1, \ldots, n, \tag{18}$$

but the main distinction between it and the models in the previous subsections is that its model parameters differ across multiple states $s \in \{0, 1\}$. Since the residual variance is also subject to this distinction, estimation of the model allows us to identify low-volatility and high-volatility regimes. Within these regimes, the ways in which the explanatory variables affect the log returns can also vary in sign, magnitude and statistical significance.

We find that expanding the number of states to three does not enhance the performance of the model. Restricting the autoregressive and/or exogenous coefficients such that they do not differ across states leads to a reduction in both computation time and the number of parameters to be estimated, but also impairs explanatory power. We also refrain from applying a multivariate extension and using a MS-VARX model. Using two states, this would add another $2 \times p \times n \times (n-1)$ parameters to a model that already contains twice the number of parameters of a regular ARX model and can be harmful for the out-of-sample predictions. Too many parameters can also impair our ability to interpret the impact of some of the regressors.

3.4 Estimation

We initialise our stochastic volatility models using the first 25% of our full sample period. We then use 75% of the remaining period to fit the models to our data, after which we apply our performance measures to the remaining data points. This yields 234 predicted returns. In total, we run 10,000 simulations, 2,000 of which are burn-in simulations. With the aim of reducing both storage space and correlation between draws, we use a thinning value equal to 10. This means that we only collect information from every 10th draw of the Markov chain. After running the full algorithm, we take posterior means.

The estimation sample of the benchmark ARX and VARX models and the MS-ARX model are constructed such that the hold-out sample period lines up with that of the stochastic volatility models. We fit the benchmark models using OLS and the MS-ARX model by means of a maximum likelihood estimation that applies the expectation maximisation algorithm. The posterior means of the stochastic volatility models as well as the estimated parameters and corresponding standard errors of the ARX, VARX and MS-ARX models can be found in the Appendix.

4 Results

Before we make use of the methods proposed in Section 3, we run estimations of the ARX and VARX models. These estimations can form a benchmark to which we can measure the capabilities of our models. Furthermore, we extract the ideal lag order p using the Akaike Information Criterion. Lag order $p = 1$ yields the best results according to the VARX model. Nevertheless, we expect that some kind of mean reversion may be present in the log returns. Higher lag orders may add some sort of momentum factor to the model without capturing the negative effects of this mean reversion. We therefore set $p = 2$, since we also exercise caution by keeping in mind the number of parameters in our models.

For the ARX and VARX models, we find only a few statistically significant parameters and most of them correspond to the BTC series. Most notably, in both models, the FRM, MVRV ratio and DXY index appear to affect the returns positively. The latter result is odd, as we would expect a rise in relative USD value to correspond to a decline in any USD denominated currency pair. However, one can also argue that buying power increases as the DXY rises and that this would then cause investors to search for store of values such as Bitcoin. The S&P 500 index has a negative effect on the BTC returns. This supports the statement we made earlier about investors fleeing to cryptocurrencies when traditional markets experience a downfall.

The MS-ARX model, which distinguishes between low-volatility and high-volatility regimes as indicated by the σ estimates, yields more significant results. One striking feature is that the effect of gold futures returns on BTC returns differs across these two regimes. During times of low volatility this effect is negative and it is positive when volatility is high. This is in line with our expectations, as both assets are known to be used as hedge instruments against losses in uncertain times. We also find that for LTC, all of our predictor except from the DXY index contain significant explanatory value at a 10% significance level. Furthermore, we note that our choice of lag order yields the desired results. The first lag order parameters are negative in the low-volatility regime for all three series, while the second lag order parameters are positive. Lastly, it appears that the BTC and ETH returns are significantly positive during times of low volatility.

4.1 Time-varying volatilities

Since our methodology is heavily focused on accounting for time variation in the asset volatilities, it is useful to look at the behaviour of the model residuals over the course of our sample period. We ignore our ARX model for now and only look at our benchmark VARX model. In the estimation of this model we do not make any other assumptions than the standard OLS assumptions, which means that the volatilities of interest are best approximated by centered rolling window residual standard deviation estimates. We do the same for the MS-ARX residuals. The SV-VARX and CSV-VARX residual standard deviations are explicitly modeled and can be extracted by taking the posterior median of the square root of the diagonal elements of $\Sigma_t = A^{-1}\Lambda_t(A^{-1})'$. We discuss the results for SVt-VARX, as well as for our correlated models, later on in this subsection.

There are four conclusions to be drawn from Figure 3. First, we notice the effect that the residual variance modelling of SV-VARX and CSV-VARX has on the standard deviations when compared to the other models. The residual series of all models contain a large negative spike on 12 March 2020, as could be expected from the return series outlier described in Section 2.1. The rolling window estimates instantly recognise this as a structural change in volatility and

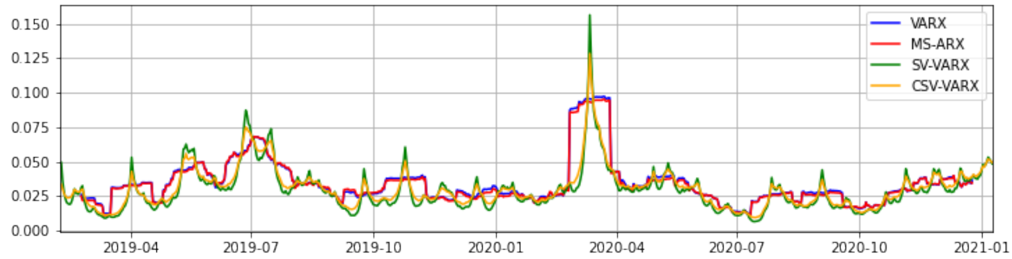

(a) BTC

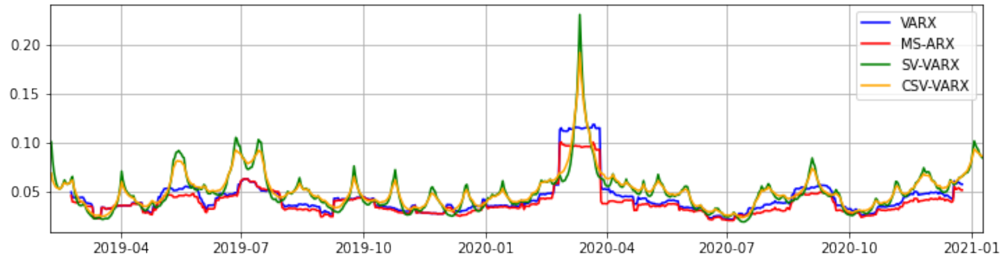

(b) ETH

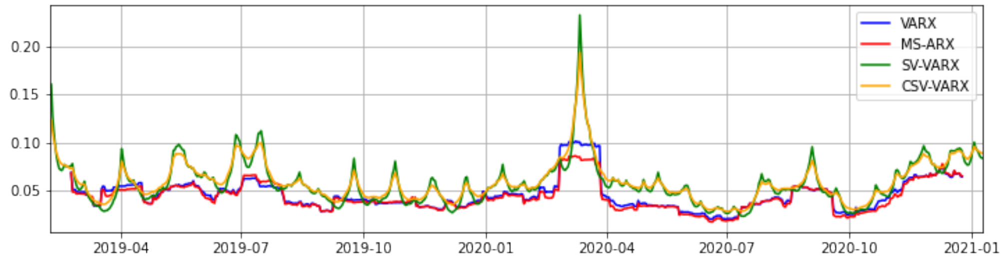

(c) LTC

Figure 3: 31-day centered rolling window standard deviation estimates of the fitted residuals of the VARX and MS-ARX models and posterior medians of the square roots of the diagonal elements of Σ_t for the SV-VARX and CSV-VARX models in the sample period 8 February 2019 - 9 January 2021.

return to the initial level as soon as the observation drops out of the rolling window. In contrast, the SV residual standard deviations are initially less affected by the aberrant observation and do not consider the change in volatility to last as long. This may benefit the SV models, as it can use this parameter uncertainty to its advantage. Secondly, we observe that the MS-ARX standard deviation estimates are generally lower than those of the VARX model. This can be explained by the dominance of the low-volatility regime that the MS-ARX model identifies. During approximately 80% of the sample period, the three series are in a low-volatility regime, which means that large shocks in the remaining 20% return series can be accounted for with another set of parameters. The advantage of the MS-ARX is that even if these observations in the high-volatility regime are hard to fit, it can still provide a good fit for the larger part of the sample period. This may imply that the residuals are generally smaller and therefore are characterised by a lower standard deviation. Our third conclusion is that the correlated log variance innovations of CSV-VARX result in a smoother residual standard deviation series when compared to those of SV-VARX. This can be explained by lower idiosyncratic variances of the innovations, i.e. the diagonal elements of Φ, which results in lower tops and higher bottoms in the time series. Finally, we notice that the correlation across our three cryptocurrencies are exceptionally high and range from 76% to 95%. This highlights the main shortcoming of uncorrelated stochastic volatility model specifications, which assume no correlation between the log variance innovations.

The residual standard deviations of the fat-tailed models are obtained by taking the posterior median of the square roots of the diagonal elements of $\Sigma_t = A^{-1}Q_t^{0.5}\Lambda_t Q_t^{0.5}(A^{-1})'$. Displayed in Figure A3 in the Appendix, we find that these time series are characterised by large shocks and are more volatile than their counterparts.

4.2 Out-of-sample performance

The day-ahead point forecasts that are obtained after fitting the models to our data run from 10 January 2021 until 31 August 2021, resulting in 234 observations that we can extract performance information from. In particular, we consider three measures:

$$\begin{aligned}\text{Root mean square error (RMSE)} &:= \sqrt{\frac{1}{T}\sum_{t=1}^{T}(\hat{y}_t - y_t)^2},\\ \text{Hit rate} &:= \frac{1}{T}\big(\mathbb{1}[\hat{y}_t > 0 \wedge y_t > 0] + \mathbb{1}[\hat{y}_t < 0 \wedge y_t < 0]\big),\\ \text{Continuous ranked probability score } (\mathrm{CRSP_t}) &:= \mathbb{E}\,|Y_t - y_t| - 0.5\,\mathbb{E}\,|Y_t - Y_t'|,\end{aligned}$$

with Y_t and Y_t' denoting independent random draws from the posterior density function. Note that this means that the CRSP only applies to our Bayesian stochastic volatility models. Statistically significant differences in both the RMSE and CRSP are assessed using a modified Diebold-Mariano test (Harvey et al., 1997). Besides traditional ARX and VARX model, we compare our methods with a random walk prediction, which simply sets the expected value of at time t equal to the value at time $t-1$, and a zero return prediction for each observation in our out-of-sample period.

The performance measures are displayed in Table 4. Comparing the stochastic volatility models with other methods using the RMSE, we notice that only for the BTC return series their predictions are significantly more accurate than those of the benchmark ARX and VARX models. They fail to outperform the zero return prediction for all three series. We do observe that the

Table 4: RMSE, hit rate and CRSP for the proposed models and the three cryptocurrencies

	RMSE			Hit Rate			CRSP		
	BTC	ETH	LTC	BTC	ETH	LTC	BTC	ETH	LTC
Random Walk	0.0672	0.0914	0.0999	0.4418	0.4658	0.5085	-	-	-
Zero Prediction	0.0448	0.0607	0.0671	-	-	-	-	-	-
ARX	0.0455	0.0604	0.0669	0.5085	0.5342	0.5085	-	-	-
VARX	0.0457	0.0614**	0.0678**	0.5299	0.5385	0.5043	-	-	-
SV-VARX	0.0450*	0.0616**	0.0687**	0.5641	0.5385	0.4872	1.0000	1.0000	1.0000
CSV-VARX	0.0450*	0.0615**	0.0686**	0.5641	0.5342	0.4786	0.9986**	0.9990**	0.9981**
SVt-VARX	0.0450*	0.0616**	0.0687**	0.5684	0.5342	0.4829	1.0000	1.0001**	1.0000
CSVt-VARX	0.0449*	0.0615**	0.0686**	0.5641	0.5342	0.4786	0.9986**	0.9989**	0.9978**
MS-ARX	0.0447*	0.0592**	0.0661*	0.5513	0.6709	0.5641	-	-	-

This table shows the RMSE, hit rate and relative CRSP for our five proposed models, benchmark ARX and VARX models and random walk and zero return predictions across the three cryptocurrencies in the sample period 10 January 2021 - 31 August 2021. For the sake of clarity, the CRSP is displayed relative to the SV-VARX model. Statistically significant differences in RMSE between our proposed models and the ARX model and in relative CRSP between the SV-VARX model and the other stochastic volatility models, as identified by the modified Diebold-Mariano test using a 10% and 5% significance level, are denoted by * and ** respectively.

MS-ARX model appears to have more predictive power than all other methods. The modified Diebold-Mariano tests indicate that its accuracy is overall superior for BTC and LTC at a 10% significance level and for ETH at a 5% significance level. The latter result is also supported by the relatively impressive hit rate, which shows that a correct sign is predicted for 67.09% of the observations. We mainly attribute this superiority to the previously mentioned advantage of allowing for a time-varying mean return, but also to the shortcoming of the straightforward (V)ARX specifications of the other methods.

We find that across our four stochastic volatility models, the measure values change rarely. This points out that, even though the dynamics in the residual variances are different to some extent, no large disparities are present in the coefficients in the VARX equation of the models. Nevertheless, based on the RMSE we observe that the correlated models significantly outperform their uncorrelated counterparts for all series. A similar result can be found regarding the regular and fat-tailed models, but we note that the difference is not as extreme. Density forecast performances, as measured by the CRSP, display inferiority of the fat-tailed models to a certain extent. The SVt-VARX density forecasts are not more accurate than those of the SV-VARX model at a 10% significance level for the BTC and LTC series and are even found to be significantly worse for the ETH series. There is no difference in density forecasting precision between the two correlated models. Again, we find that correlated innovations improve accuracy.

5 Conclusion

In this paper, we introduce multiple approaches that attempt to predict the log returns of three cryptocurrencies under the assumption that volatilities are highly time-varying. Four of our methods use a Bayesian stochastic volatility approach and model the log variance of the residuals of a VARX specification to follow a random walk process. Within this group of models, we differentiate between both uncorrelated and correlated random walk innovations and both normally and t-distributed residuals. We also apply a Markov switching ARX model, which allows for the VARX parameters and residual variances to differ across two regimes. The accuracy of the point forecasts as well as the density forecasts of all models are assessed on an out-of-sample basis using a set of three measures. OLS estimations of ARX and VARX models serve as benchmarks in these comparisons.

We document strong time-varying volatility in all model residuals, justifying our choice of methodology. In this case, the correlated stochastic volatility models seem to be especially suitable, as exceptionally high cross-correlations are observed in residual standard deviations across the three return series. RMSE and CRSP analyses show that there is small but insignificant outperformance of these particular models. Allowing for fat tails using a t-distribution slightly improves the accuracy of the point forecast, but not of the density forecasts. Nevertheless, it appears that there are nonlinearities present in the relation between our predictor variables and the cryptocurrency returns that cannot be captured by the straightforward (V)ARX specification of our benchmark and stochastic volatility models. Our MS-ARX model supports this statement by significantly improving upon the predictions of all other methods. We thus conclude that there are volatility regimes within the dynamics of the cryptocurrencies. Moreover, our research shows that our predictor variables affect the return series differently across regimes, suggesting that it is useful to further investigate these features in future research. We also note that machine learning methods can be especially useful to capture these apparent complex relations. For instance, Akyildirim et al. (2021) show that cryptocurrency return predictions of machine learning algorithms are significantly more accurate than those of statistical time series models.

There are multiple limitations to our research. First, we are aware of the fact that the cryptocurrency market remains somewhat influenced by manipulative players. Even though the inflow of institutional investors is growing rapidly at the time of writing, a lot of active retail investors have an unrealistic view on the market as a whole. We thus find that a lot of variation in the returns cannot be explained by either economic variables or on-chain metrics and are simply the result of behavioural biases. Furthermore, price data is limited in this particular market. Cryptocurrencies have only been around for approximately ten years and we deem most price data in this period unpredictable due to the dominance of previously described investor behaviour. This drastically reduces the number of valuable data points. There is also the case of on-chain data availability. We consider three of the most established tokens by market capitalisation and lifespan, but even then we encounter the problem that not every interesting metric is available for every token. We hope that, over time, these issues can be resolved, such that future research can build upon the idea of investigating predictive power of on-chain metrics. Lastly, concerning our methodology, simulation of our Bayesian models is rather time-expensive. We believe that the performance of these models could have been improved if we re-estimated our model parameters using the additional information at each point in time in our out-of-sample period. This way, we would end up with h-step-ahead forecasts, keeping h reasonably low, which would be more reliable for performance comparisons between models. For example, we would then be able to implement the performance measures as described in Giacomini and Rossi (2010), accounting for the unstable time-varying nature of our models.

References

Akyildirim, E., Goncu, A. and Sensoy, A. (2021), 'Prediction of cryptocurrency returns using machine learning', *Annals of Operations Research* **297**(1), 3–36.

Aoyagi, J. and Hattori, T. (2019), 'The empirical analysis of bitcoin market in the general equilibrium framework', *Available at SSRN 3433833* .

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

Bariviera, A. F., Basgall, M. J., Hasperué, W. and Naiouf, M. (2017), 'Some stylized facts of the bitcoin market', *Physica A: Statistical Mechanics and its Applications* **484**, 82–90.

Baur, D. G. and Dimpfl, T. (2018), 'Asymmetric volatility in cryptocurrencies', *Economics Letters* **173**, 148–151.

Bianchi, D. (2020), 'Cryptocurrencies as an asset class? an empirical assessment', *The Journal of Alternative Investments* **23**(2), 162–179.

Bohte, R. and Rossini, L. (2019), 'Comparing the forecasting of cryptocurrencies by bayesian time-varying volatility models', *Journal of Risk and Financial Management* **12**(3), 150.

Bouri, E., Lucey, B. and Roubaud, D. (2020), 'Cryptocurrencies and the downside risk in equity investments', *Finance Research Letters* **33**, 101211.

Caporale, G. M. and Zekokh, T. (2019), 'Modelling volatility of cryptocurrencies using markov-switching garch models', *Research in International Business and Finance* **48**, 143–155.

Catania, L., Grassi, S. and Ravazzolo, F. (2019), 'Forecasting cryptocurrencies under model and parameter instability', *International Journal of Forecasting* **35**(2), 485–501.

Cheah, E.-T. and Fry, J. (2015), 'Speculative bubbles in bitcoin markets? an empirical investigation into the fundamental value of bitcoin', *Economics letters* **130**, 32–36.

Chung, H., Laforte, J.-P., Reifschneider, D. and Williams, J. C. (2012), 'Have we underestimated the likelihood and severity of zero lower bound events?', *Journal of Money, Credit and Banking* **44**, 47–82.

Clark, T. E. and Ravazzolo, F. (2015), 'Macroeconomic forecasting performance under alternative specifications of time-varying volatility', *Journal of Applied Econometrics* **30**(4), 551–575.

Cogley, T. and Sargent, T. J. (2005), 'Drifts and volatilities: monetary policies and outcomes in the post wwii us', *Review of Economic Dynamics* **8**(2), 262–302.

Corbet, S., Lucey, B. and Yarovaya, L. (2018), 'Datestamping the bitcoin and ethereum bubbles', *Finance Research Letters* **26**, 81–88.

Corbet, S., Meegan, A., Larkin, C., Lucey, B. and Yarovaya, L. (2018), 'Exploring the dynamic relationships between cryptocurrencies and other financial assets', *Economics Letters* **165**, 28–34.

Dueker, M. and Neely, C. J. (2007), 'Can markov switching models predict excess foreign exchange returns?', *Journal of Banking & Finance* **31**(2), 279–296.

Dyhrberg, A. H. (2016), 'Hedging capabilities of bitcoin. is it the virtual gold?', *Finance Research Letters* **16**, 139–144.

Fantazzini, D. and Kolodin, N. (2020), 'Does the hashrate affect the bitcoin price?', *Journal of Risk and Financial Management* **13**(11), 263.

Fernández-Villaverde, J. and Rubio-Ramírez, J. (2010), Macroeconomics and volatility: Data, models, and estimation, Technical report, National Bureau of Economic Research.

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

Giacomini, R. and Rossi, B. (2010), 'Forecast comparisons in unstable environments', *Journal of Applied Econometrics* **25**(4), 595–620.

Greer, R. J. (1997), 'What is an an asset class, anyway?', *Journal of Portfolio Management* **23**(2), 86.

Harvey, D., Leybourne, S. and Newbold, P. (1997), 'Testing the equality of prediction mean squared errors', *International Journal of Forecasting* **13**(2), 281–291.

Hayes, A. S. (2017), 'Cryptocurrency value formation: An empirical study leading to a cost of production model for valuing bitcoin', *Telematics and Informatics* **34**(7), 1308–1321.

Jacquier, E., Polson, N. G. and Rossi, P. E. (2004), 'Bayesian analysis of stochastic volatility models with fat-tails and correlated errors', *Journal of Econometrics* **122**(1), 185–212.

Jacquier, E., Polson, N. G., Rossi, P. et al. (1999), *Stochastic volatility: Univariate and multivariate extensions*, CIRANO.

Kadiyala, K. R. and Karlsson, S. (1997), 'Numerical methods for estimation and inference in bayesian var-models', *Journal of Applied Econometrics* **12**(2), 99–132.

Krueckeberg, S. and Scholz, P. (2019), Cryptocurrencies as an asset class, *in* 'Cryptofinance and Mechanisms of Exchange', Springer, pp. 1–28.

Lahmiri, S., Bekiros, S. and Salvi, A. (2018), 'Long-range memory, distributional variation and randomness of bitcoin volatility', *Chaos, Solitons & Fractals* **107**, 43–48.

Li, J. and Yi, G. (2019), 'Toward a factor structure in crypto asset returns', *The Journal of Alternative Investments* **21**(4), 56–66.

Ma, F., Liang, C., Ma, Y. and Wahab, M. (2020), 'Cryptocurrency volatility forecasting: A markov regime-switching midas approach', *Journal of Forecasting* **39**(8), 1277–1290.

Matta, M., Lunesu, I. and Marchesi, M. (2015), Bitcoin spread prediction using social and web search media., *in* 'UMAP workshops', pp. 1–10.

Meynkhard, A. (2019), 'Fair market value of bitcoin: Halving effect', *Investment Management and Financial Innovations* **16**(4), 72–85.

Nakamoto, S. (2019), Bitcoin: A peer-to-peer electronic cash system, Technical report, Manubot.

Platanakis, E. and Urquhart, A. (2020), 'Should investors include bitcoin in their portfolios? a portfolio theory approach', *The British Accounting Review* **52**(4), 100837.

Poyser, O. (2018), 'Herding behavior in cryptocurrency markets'.

Schaller, H. and Norden, S. V. (1997), 'Regime switching in stock market returns', *Applied Financial Economics* **7**(2), 177–191.

Shanaev, S., Sharma, S., Shuraeva, A. and Ghimire, B. (2019), 'The marginal cost of mining, metcalfe's law and cryptocurrency value formation: Causal inferences from the instrumental variable approach', *Metcalfe's Law and Cryptocurrency Value Formation: Causal Inferences from the Instrumental Variable Approach (June 7, 2019)* .

Tan, C.-Y., Koh, Y.-B., Ng, K.-H. and Ng, K.-H. (2021), 'Dynamic volatility modelling of bitcoin using time-varying transition probability markov-switching garch model', *The North American Journal of Economics and Finance* **56**, 101377.

Zhang, W., Wang, P., Li, X. and Shen, D. (2018), 'Some stylized facts of the cryptocurrency market', *Applied Economics* **50**(55), 5950–5965.

A Appendix

A.1 Figures

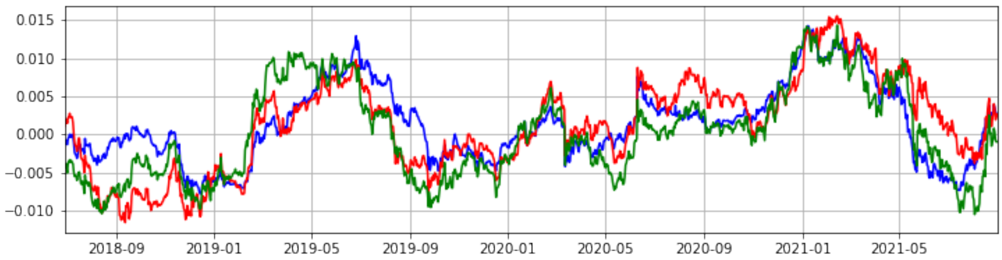

(a) Mean

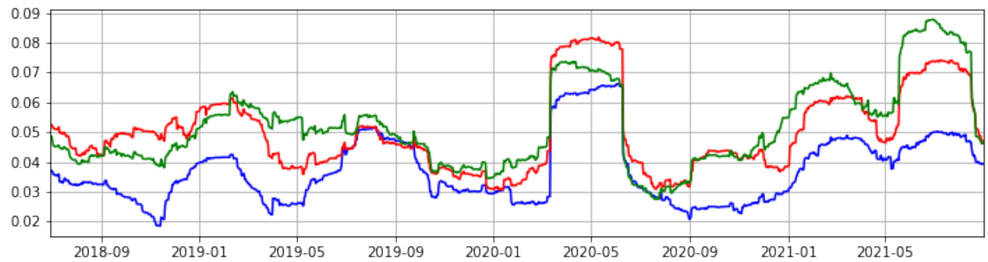

(b) Standard deviation

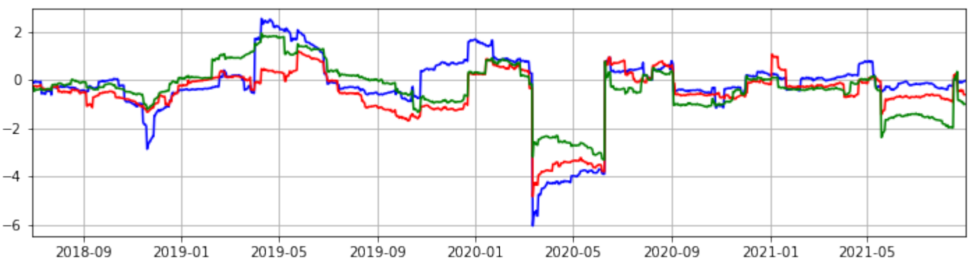

(c) Skewness

Figure A1: 90-day rolling window mean, standard deviation and skewness estimates of the three cryptocurrency return series in the sample period 1 July 2018 - 31 August 2021.

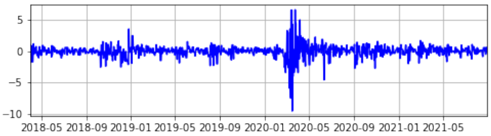

(a) S&P 500

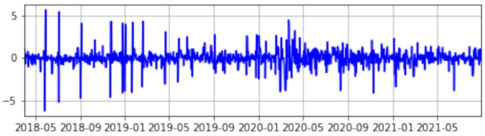

(b) Gold futures

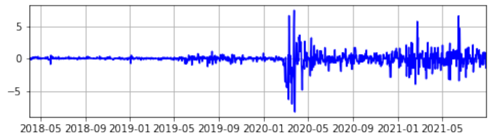

(c) 1-year US bond

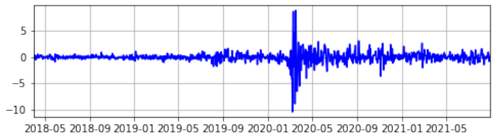

(d) 10-year US bond

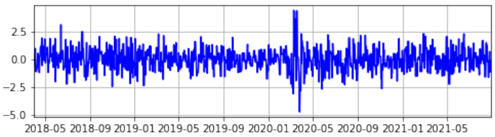

(e) DXY

Figure A2: Standardised log returns of the five economic variables.

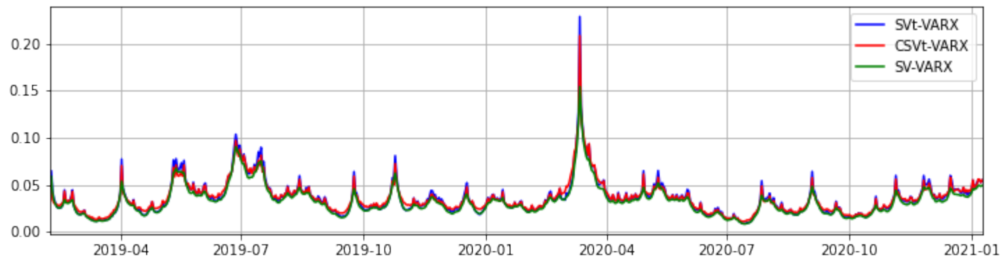

(a) BTC

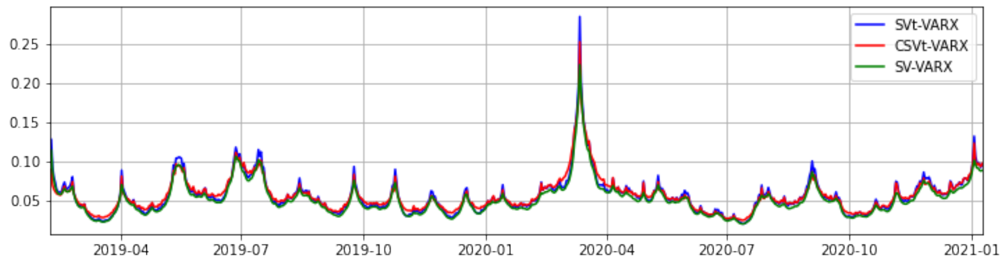

(b) ETH

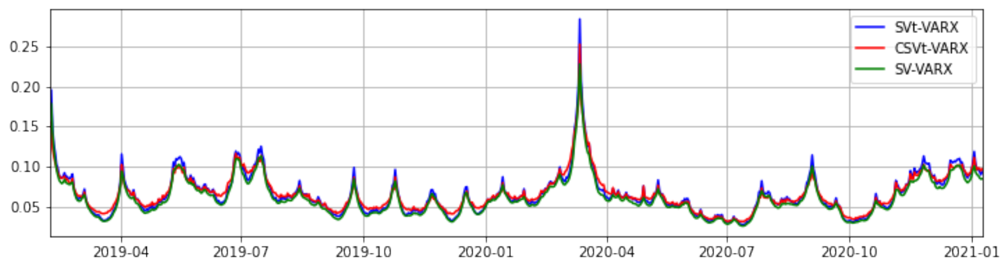

(c) LTC

Figure A3: Posterior medians of the square roots of the diagonal elements of Σ_t for the SVt-VARX and SV-VARX models in the sample period 8 February 2019 - 9 January 2021.

A.2 Tables

Table A1: Estimated parameters of the ARX and VARX models

	ARX			VARX		
	BTC	ETH	LTC	BTC	ETH	LTC
α	0.0025**	0.0014	0.0002	0.0027**	0.0018	0.0005
BTC-1	−0.0732**	-	-	−0.0634	−0.1680**	−0.1042
ETH-1	-	−0.0505	-	−0.0940*	0.0043	−0.1619**
LTC-1	-	-	−0.0449	0.0878*	0.0575	0.1557**
BTC-2	0.0526	-	-	−0.0261	−0.0198	0.0042
ETH-2	-	0.0956***	-	0.0896*	0.0900	0.0268
LTC-2	-	-	0.0571*	−0.0189	0.0191	0.0301
BTC-FRM	0.0028**	-	-	0.0027**	-	-
BTC-MVRV	0.0230**	-	-	0.0235**	-	-
BTC-NVT	−0.0018	-	-	−0.0017	-	-
ETH-FRM	-	0.0013	-	-	0.0021	-
ETH-MVRV	-	0.0401	-	-	0.0577	-
ETH-NVT	-	−0.0011	-	-	−0.0021	-
LTC-FRM	-	-	−0.0018	-	-	−0.0014
LTC-MVRV	-	-	−0.0143	-	-	−0.0123
LTC-NVT	-	-	0.0018	-	-	0.0019
S&P 500	−0.0027**	−0.0050	−0.0003	−0.0028**	−0.0064	−0.0001
Gold	−0.0019	−0.0034	−0.0004	−0.0018	−0.0041	0.0000
US 1yr	−0.0001	0.0000	−0.0026	−0.0003	0.0012	−0.0027
US 10yr	0.0005	0.0014	0.0029	0.0007	0.0011	0.0032*
DXY	0.0032**	0.0027	−0.0008	0.0033**	0.0035	−0.0007

This table contains the parameters of the ARX and VARX models that are estimated using OLS over the sample period 1 April 2018 - 31 August 2021. Statistical significance is denoted with *, ** and ***, corresponding to a significance level of 10%, 5% and 1% respectively.

Table A2: Posterior means of the VARX-SV, VARX-CSV, VARX-SVt and VARX-CSVt model parameters

	VARX-SV			VARX-CSV			VARX-SVt			VARX-CSVt		
	BTC	ETH	LTC	BTC	ETH	LTC	BTC	ETH	LTC	BTC	ETH	LTC
α	0.0040	0.0042	0.0027	0.0040	0.0041	0.0027	0.0040	0.0042	0.0027	0.0040	0.0041	0.0027
BTC-1	-0.0321	-0.1541	-0.0626	-0.0313	-0.1536	-0.0621	-0.0318	-0.1540	-0.0625	-0.0312	-0.1536	-0.0623
ETH-1	-0.1425	-0.0291	-0.2308	-0.1422	-0.0290	-0.2303	-0.1425	-0.0290	-0.2309	-0.1422	-0.0290	-0.2303
LTC-1	0.1063	0.0816	0.1731	0.1060	0.0814	0.1732	0.1063	0.0816	0.1733	0.1060	0.0814	0.1733
BTC-2	-0.0413	-0.0050	0.0084	-0.0403	-0.0051	0.0089	-0.0409	-0.0050	0.0085	-0.0402	-0.0052	0.0089
ETH-2	0.1122	0.0581	0.0046	0.1127	0.0590	0.0046	0.1124	0.0585	0.0047	0.1129	0.0591	0.0048
LTC-2	-0.0293	0.0359	0.0183	-0.0298	0.0356	0.0187	-0.0295	0.0358	0.0184	-0.0300	0.0355	0.0185
BTC-FRM	0.0007	0.0001	0.0000	0.0007	0.0001	0.0000	0.0007	0.0001	0.0000	0.0006	0.0001	0.0000
BTC-MVRV	0.0126	0.0000	0.0000	0.0117	0.0000	0.0000	0.0123	0.0000	0.0000	0.0115	0.0000	0.0000
BTC-NVT	0.0002	-0.0000	0.0000	0.0003	-0.0000	0.0000	0.0002	-0.0000	0.0000	0.0003	-0.0000	0.0000
ETH-FRM	0.0000	-0.0008	0.0000	0.0000	-0.0010	0.0000	0.0000	-0.0009	0.0000	0.0000	-0.0010	0.0000
ETH-MVRV	-0.0000	0.0652	0.0000	-0.0000	0.0605	0.0000	-0.0000	0.0636	0.0000	-0.0000	0.0601	0.0000
ETH-NVT	-0.0000	-0.0019	0.0000	-0.0000	-0.0017	0.0000	-0.0000	-0.0018	0.0000	-0.0000	-0.0017	0.0000
LTC-FRM	0.0000	0.0000	0.0041	0.0000	0.0000	0.0037	0.0000	0.0001	0.0041	0.0000	0.0000	0.0037
LTC-MVRV	0.0000	-0.0000	0.0506	0.0000	-0.0000	0.0466	0.0000	-0.0000	0.0500	0.0000	-0.0000	0.0465
LTC-NVT	-0.0000	0.0001	-0.0018	-0.0000	0.0000	-0.0017	-0.0000	0.0001	-0.0018	-0.0000	0.0000	-0.0017
S&P 500	-0.0025	-0.0072	-0.0044	-0.0025	-0.0068	-0.0041	-0.0025	-0.0071	-0.0043	-0.0025	-0.0068	-0.0041
Gold	0.0002	-0.0023	-0.0010	0.0002	-0.0021	-0.0008	0.0002	-0.0022	-0.0010	0.0002	-0.0021	-0.0008
US 1yr	-0.0010	0.0016	-0.0003	-0.0010	0.0013	-0.0005	-0.0010	0.0015	-0.0004	-0.0010	0.0013	-0.0005
US 10yr	0.0017	0.0012	0.0016	0.0018	0.0012	0.0017	0.0017	0.0012	0.0016	0.0018	0.0013	0.0017
DXY	0.0031	0.0049	0.0027	0.0031	0.0047	0.0025	0.0031	0.0048	0.0027	0.0031	0.0046	0.0025

This table contains the posterior means of the VARX-SV, VARX-CSV, VARX-SVt and VARX-CSVt model parameters that are obtained using our MCMC algorithm over the sample period 8 February 2019 - 31 August 2021 after initialization over 1 April 2018 - 8 February 2019.

Table A3: Estimated parameters of the MS-ARX model

	BTC		ETH		LTC	
Regime	0	1	0	1	0	1
σ	0.0002	0.0037	0.0007	0.0069	0.0006	0.0058
α	0.0028***	0.0028	0.0035**	0.0014	−0.0014	0.0055
BTC-1	−0.0714**	−0.1759***	-	-	-	-
ETH-1	-	-	−0.1457***	0.0352	-	-
LTC-1	-	-	-	-	−0.1864***	0.0442
BTC-2	0.0943***	−0.0880	-	-	-	-
ETH-2	-	-	0.0665**	0.0419	-	-
LTC-2	-	-	-	-	0.0591**	0.0571
BTC-FRM	0.0018*	0.0009	-	-	-	-
BTC-MVRV	0.0240***	−0.0026	-	-	-	-
BTC-NVT	−0.0020	−0.0031	-	-	-	-
ETH-FRM	-	-	0.0031	−0.0041	-	-
ETH-MVRV	-	-	0.1423	−0.1345	-	-
ETH-NVT	-	-	−0.0078	0.0127	-	-
LTC-FRM	-	-	-	-	0.0110***	−0.0091*
LTC-MVRV	-	-	-	-	0.0857***	−0.0271*
LTC-NVT	-	-	-	-	−0.0028	0.0074*
S&P 500	−0.0013	−0.0054***	−0.0168**	0.0298*	−0.0084***	0.0218***
Gold	−0.0028***	0.0033**	−0.0058	−0.0075	−0.0020	−0.0105**
US 1yr	−0.0002	0.0027	0.0101	−0.0260*	0.0035**	−0.0088**
US 10yr	−0.0013	0.0019	0.0025	−0.0031	0.0022	−0.0070**
DXY	0.0022**	−0.0032	0.0078	−0.0090	0.0046	−0.0028

This table contains the parameters of the MS-ARX model that are estimated using maximum likelihood estimation over the sample period 1 April 2018 - 31 August 2021. Statistical significance is denoted with *, ** and ***, corresponding to a significance level of 10%, 5% and 1% respectively.