

Long-Memory Property in Cryptocurrency Realized Volatility

Burak Korkusuz¹

Abstract

This study investigates the realized volatility of major cryptocurrencies with a focus on long memory properties. Using high-frequency data, the study evaluates whether long-memory properties improve volatility modelling and forecasting performance. To this end, the four well-known time series models are compared: the ARFIMA, which explicitly captures fractional integration and long memory; the HAR model, which approximates long-memory behaviour through multi-horizon components; and the benchmark AR model and ARMA model specifications. The results reveal strong persistence in cryptocurrency volatility consistent with long-memory behaviour. However, the HAR model consistently outperforms alternatives in both in-sample and out-of-sample forecasts, suggesting that multi-scale structures capture long-memory dynamics more effectively than the ARFIMA model. Overall, the findings highlight that while long memory is a key feature of cryptocurrency volatility, it is better approximated through heterogeneous temporal components, with important implications for forecasting and risk management.

1. INTRODUCTION

Cryptocurrencies have increasingly attracted attention in financial markets, largely due to their distinctive features and their appeal as alternative investment instruments. As trading activity expands and institutional participation grows, accurately understanding and forecasting cryptocurrency volatility has become crucial for market participants, including traders, risk managers, and policymakers. However, the high-frequency nature of trading and the unique volatility patterns observed in these assets make it challenging for traditional

1 Dr., Osmaniye Korkut Ata University, Faculty of Economics and Administrative Sciences, Department of Econometrics, 80000, Osmaniye, Türkiye, burakkorkusuz@osmaniye.edu.tr, ORCID ID: 0000-0001-9374-2350

volatility models to fully capture their complex dynamics (Katsiampa, 2017, 2019).

The volatility behaviour of cryptocurrencies is influenced by several factors, such as speculative trading, evolving regulatory frameworks, and the relatively early stage of market development (Corbet et al., 2019). Previous studies have emphasized the difficulty of modelling such volatility, given the presence of extreme price swings and potential long-memory characteristics (Phillip et al., 2018). In this regard, long-memory processes, captured by models like ARFIMA, are particularly relevant, as shocks in cryptocurrency markets can have persistent effects over extended periods (Bariviera, 2017). Supporting this, Andersen et al. (2003) highlight the effectiveness of ARFIMA models in modelling realized volatility, while McMillan and Speight (2010) note that fractional integration allows for a deeper understanding of persistence in volatile time series.

In addition, the Heterogeneous Autoregressive (HAR) model, which incorporates daily, weekly, and monthly components, is well-suited to capturing volatility dynamics across multiple time horizons. Originally developed for traditional financial markets (Corsi, 2009), this framework may also provide valuable insights when applied to cryptocurrency volatility forecasting.

Empirical evidence on the performance of volatility models in cryptocurrency markets remains mixed. Early research has largely concentrated on Bitcoin, often employing GARCH-type models. For example, Bouoiyour and Selmi (2015, 2016) evaluate different GARCH specifications over specific periods, while Katsiampa (2017) extends this analysis over a longer sample. Balcilar et al. (2017) explore the relationship between trading volume and volatility, finding no predictive power of volume for Bitcoin volatility. Meanwhile, Troster et al. (2019) focus on the heavy-tailed nature of Bitcoin returns and show that models accommodating such features outperform those assuming normality, with the GAS model yielding more accurate Value-at-Risk forecasts. Despite these developments, GARCH-based approaches using daily data continue to be widely applied in the literature.

Against this background, the present study examines the realized volatility dynamics of four major cryptocurrencies—Bitcoin (BTC), Ethereum (ETH), Litecoin (LTC), and Ripple (XRP)—using high-frequency data spanning May 2018 to January 2023. The analysis evaluates the performance of several time series models, including ARFIMA, HAR, AR, and ARMA, in capturing and forecasting volatility patterns. By comparing these models across multiple cryptocurrencies, this study contributes to the existing literature by highlighting differences in volatility behaviour across assets. The results indicate that each

cryptocurrency exhibits distinct dynamics, with HAR and ARMA models generally delivering superior performance due to their ability to capture multi-scale dependencies and persistence effects. Both in-sample and out-of-sample analyses are conducted to ensure robust evaluation. Overall, the findings underline the importance of selecting appropriate volatility models for different cryptocurrencies and demonstrate the usefulness of the HAR framework in capturing medium- and long-term volatility structures.

The remainder of this work is organized as follows. Section 2 presents the methods used in this work. Section 3 describes the characteristics of high frequency data. Section 4 gives the empirical findings. Finally, Section 5 provides the summary and conclusion of this study.

2. METHODS

2.1. Realized Volatility

Volatility is not directly observable, which means that an appropriate proxy must be used to approximate its true value. Traditionally, daily squared returns have been widely employed for this purpose. However, Andersen and Bollerslev (1998) show that this measure provides a relatively weak representation of actual volatility when compared to the information contained in high-frequency data. They demonstrate that aggregating squared intraday returns offers a more accurate and informative measure. Based on this approach, Andersen and Bollerslev (1998) define realized volatility (RV_t) as follows:

	$RV_t = \sum_{i=1}^m r_{t,i}^2$	(1)
--	---------------------------------	-----

where $r_{t,i}^2$ represents the logarithm of an asset's squared return on day t and i^{th} intraday interval. $r_{t,i} = p_{t,i} - p_{t,i-1}$, ($t = 1, \dots, T; i = 1, \dots, N$) is the formula of intraday return, where T the total number of trading days, N the number of intraday intervals and $p_{t,i}$ is the log closing price for a respective asset.

2.2. ARFIMA models

The ARFIMA (Autoregressive Fractionally Integrated Moving Average) model, originally proposed by Granger and Joyeux (1980), is widely recognized for its ability to capture long-memory behaviour in time series data. It is particularly useful in modelling volatility persistence, as it allows shocks to have effects that decay slowly over time. By incorporating both short-

term dynamics and long-term dependence, the ARFIMA framework offers improved forecasting performance in series exhibiting such characteristics. In the context of realized volatility, Andersen et al. (2003) highlight the suitability of ARFIMA models for effectively capturing these persistence patterns. The general specification of the ARFIMA (p, d, q) model is given as follows:

$$\varphi(L)(1-L)^d(RV_t - \mu) = \theta(L)\varepsilon_t \quad (2)$$

In Equation 2, $\varphi(L)$ and $\theta(L)$ represent the lag polynomials of the Autoregressive (AR) and Moving Average (MA) parts, respectively. The term ε_t symbolises to the error term, which is assumed to follow a Normal distribution with a mean of zero and constant variance [$N(0, \sigma_u^2)$]. The parameter d in this formula represents the fractional differencing, which captures the persistence behaviour of volatility. Typical AR and MA parts explain the short memory features of volatility.

2.3. AR and ARMA

From Equation 2, AR(p, 0, 0) and ARMA(p, 0, q) models can be simply derived. In the case of Moving Average (MA) parameter and Fractional Differencing Parameter (d) being zero in the equation, AR model is obtained. In a similar vein, ARMA could be obtained by omitting “ d ” parameter from Equation 2.

2.4. HAR model

The HAR-RV model is grounded in the Heterogeneous Market Hypothesis proposed by Muller et al. (1997), which argues that market participants differ in their investment horizons, risk attitudes, and responses to new information. According to this perspective, various groups of investors interpret and react to the same market signals in different ways, leading to heterogeneous dynamics in volatility. Building on this idea, and extending the HARARCH framework, Corsi (2009) introduces the Heterogeneous Autoregressive (HAR) model as a practical approach for modelling realized volatility (RV).

$$RV_{t+h}^d = \beta_0 + \beta_d RV_t^d + \beta_w RV_t^w + \beta_m RV_t^m + \varepsilon_{t+h} \quad (2)$$

where RV_t^d is daily RV part; RV_t^w refers to weekly RV component, and then RV_t^m indicates monthly RV parameter. RV_t^w and RV_t^m can be computed as follows:

$$RV_t^w = \frac{1}{5} (RV_{t-5}^d + RV_{t-4}^d + \dots + RV_{t-1}^d) \quad (3)$$

$$RV_t^m = \frac{1}{22} (RV_{t-22}^d + RV_{t-21}^d + \dots + RV_{t-1}^d) \quad (4)$$

This model functions as an additive cascade of different volatility components, designed to capture volatility across various time horizons, specifically daily, weekly, and monthly. Its specification combines these components to provide a more holistic approach to volatility forecasting.

2.5. Recursive Window Forecasting Technique

The recursive window forecasting approach expands the estimation sample progressively over time. In contrast to rolling window methods, which replace older observations with newer ones, the recursive method keeps all previously available data in the sample. As additional observations become available, they are simply appended to the existing dataset, allowing the model to rely on an increasingly larger information set. Typically, the procedure begins with an initial sample to produce one-step-ahead forecasts, and with each subsequent iteration, the sample is updated by including the new data without discarding earlier observations. This expanding framework can improve forecast accuracy by utilizing the full history of the series.

2.6. Loss Functions

In this analysis, several well-established loss functions are employed to evaluate the forecasting accuracy of volatility models. The primary loss functions include Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). Together, these loss functions provide a comprehensive assessment of model performance, helping to identify the most effective volatility forecasting methods.²

2 To save more space, this study refers to the methodology section of Sahiner et al. (2023) for the formulas of these criteria.

3. DATA

This study employs high-frequency cryptocurrency data for Bitcoin, Ethereum, Litecoin, and Ripple, obtained from the “CryptoDataDownload” website. The data spans nearly five years, from May 15, 2018, to January 24, 2023, and includes a total of 493,968 intra-day observations at 5-minute intervals. Hol and Koopman (2002) recommend a frequency interval of 5 to 30 minutes for accurate realized volatility measurements. Furthermore, Liu, Patton, and Sheppard (2015) emphasize the effectiveness of the 5-minute realized variance, reinforcing its acceptance as a standard in volatility measurement.

After collecting this extensive high-frequency dataset, the 5-minute interval data is aggregated to daily realized variance values, resulting in 1,716 daily observations. For forecasting purposes, approximately the first 716 data points (around 40%) are designated for in-sample estimation, while the remaining 1000 observations (around 60%) are reserved for out-of-sample forecasting. The study employs a one-step-ahead recursive window approach to produce volatility forecasts.

The Hurst exponent graphs for the realized variance (RV) series of BTC, ETH, LTC, and XRP are shown in Figure 2. Each Hurst exponent exceeds 0.5, indicating a persistent long-memory behaviour for these four cryptocurrencies. Specifically, BTC and LTC have an estimated Hurst exponent of 0.71, ETH has 0.69, and XRP shows 0.78. These values suggest that the realized volatility of these assets is highly persistent, where past volatility impacts future volatility. This persistence is critical for modelling, as it indicates that volatility does not follow a random walk but instead exhibits trends that can be leveraged for forecasting and risk management. This insight into volatility patterns underscores the importance of using models that incorporate long-memory features, such as the ARFIMA and HAR models, in the cryptocurrency market.

Figure 1: Graphs of Realized Volatility Series for Each Coin

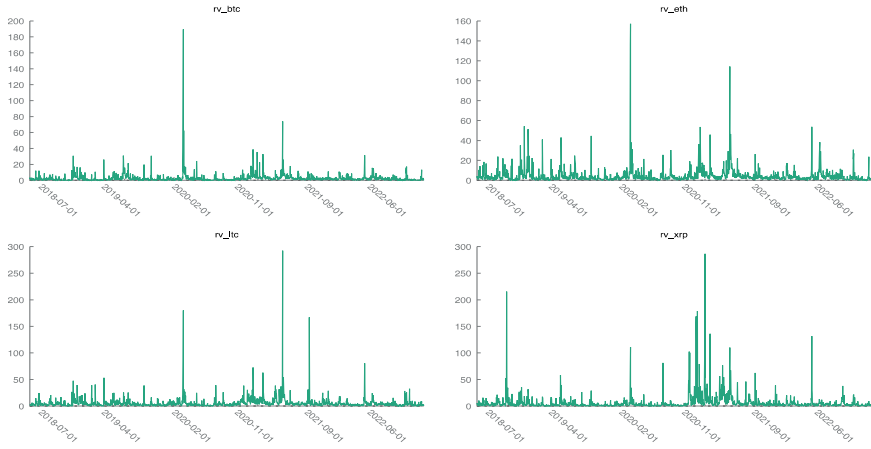
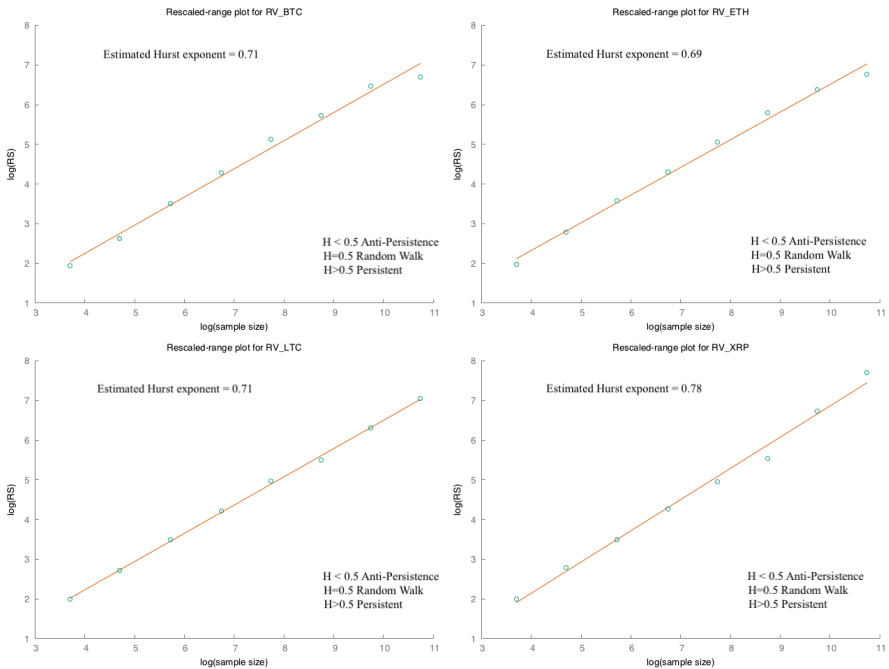


Figure 2: Hurst Exponent Estimation Results for Each Coin



4. EMPIRICAL FINDINGS

4.1. In-sample estimation results

The estimation results provide valuable insights into the volatility dynamics of Bitcoin (BTC), Ethereum (ETH), Litecoin (LTC), and Ripple (XRP). Statistical significance at the 5% level is indicated by an asterisk (*), while coefficients without this marker are not statistically significant. The findings suggest that each cryptocurrency displays distinct patterns in terms of volatility persistence and dependence on past values across the ARFIMA, HAR, AR, and ARMA models.

For Bitcoin, the ARFIMA model highlights the presence of long-memory behaviour, as reflected by the statistically significant fractional differencing parameter ($D = 0.292$). The HAR model also captures persistence effectively, particularly through its daily and weekly components, which are statistically significant. In addition, both AR and ARMA specifications exhibit significant autoregressive effects for BTC. However, the ARMA model provides a slightly better fit by combining a strong autoregressive term with a significant negative moving average component, allowing it to better represent short-term adjustments in volatility. Overall, while the ARFIMA model is effective in capturing long-run dependence, both the HAR model with its multi-horizon structure and the ARMA model with its combined dynamics also perform well in modelling Bitcoin's volatility behaviour.

Table 1: Full sample estimation results of competing models for BTC and ETH

BITCOIN (BTC)				
Parameter/Model	ARFIMA	HAR	AR	ARMA
Constant	2.640	1.065*	2.816*	2.814*
D (fractional parameter)	0.292*	-	-	-
AR(1)	-0.456*	-	0.456*	0.690*
MA(1)	-0.571*	-	-	-0.307*
HAR week component	-	0.174*	-	-
HAR month component	-	0.080	-	-
HAR day component	-	0.369*	-	-
R-squared	0.22	0.22	0.20	0.21
ETHEREUM (ETH)				
Parameter/Model	ARFIMA	HAR	AR	ARMA
Constant	4.412	1.481*	4.583*	4.581*
D (fractional parameter)	0.275*	-	-	-

AR(1)	-0.196*	-	0.498*	0.656*
MA(1)	0.386*	-	-	-0.219*
HAR week component	-	0.195*	-	-
HAR month component	-	0.086	-	-
HAR day component	-	0.394*	-	-
R-squared	0.26	0.26	0.25	0.25

For Ethereum, all the models provide a similar fit, with their R-squared values. The ARFIMA and HAR models show significance for some persistence in realized volatility. In comparison, the AR and ARMA models also perform well, particularly in capturing the autoregressive structure, but they lack the additional explanatory power provided by the HAR model's time components.

Litecoin displays the lowest R-squared values across all models, indicating limited explanatory power, although both the ARFIMA and ARMA models show relatively significant coefficients, especially for the autoregressive and moving average terms. Despite the low overall fit, the ARMA model's significant AR(1) and MA(1) terms suggest that it is such a promising model for LTC by effectively capturing both immediate past effects and moving average patterns in comparison to the long memory ARFIMA and HAR models.

For Ripple, the HAR model stands out with significant coefficients for both weekly and monthly components, indicating that XRP's volatility is influenced by these specific time frames. The ARFIMA model can also capture this long memory effect with its significant fractional differencing parameter ($D = 0.246$). In terms of the AR and ARMA models, they also seem to perform well with their R-squared values and both significant AR(1) and MA(1) terms.

Table 2: Full sample estimation results of competing models for LTC and XRP

LITECOIN (LTC)				
Parameter/Model	ARFIMA	HAR	AR	ARMA
Constant	5.532	2.401*	5.763*	5.750*
D (fractional parameter)	0.234*	-	-	-
AR(1)	-0.743*	-	0.270*	0.851*
MA(1)	0.714*	-	-	-0.670*
HAR week component	-	0.269*	-	-
HAR month component	-	0.164*	-	-
HAR day component	-	0.152*	-	-

R-squared	0.11	0.10	0.07	0.10
RIPPLE (XRP)				
Parameter/Model	ARFIMA	HAR	AR	ARMA
Constant	6.242	1.783*	6.591*	6.586*
D (fractional parameter)	0.246*	-	-	-
AR(1)	0.167	-	0.474*	0.713*
MA(1)	-0.017	-	-	-0.322*
HAR week component	-	0.172*	-	-
HAR month component	-	0.196*	-	-
HAR day component	-	0.361*	-	-
R-squared	0.25	0.25	0.22	0.23

In general, the HAR and ARFIMA models tend to perform better in modelling cryptocurrency realized volatility patterns with their slightly higher R-squared values, particularly for assets like Ethereum and Ripple. Overall, the HAR model's ability to incorporate weekly and monthly components provides a stronger explanatory power for most coins, especially when compared to models that rely solely on autoregressive structures. Thus, while the ARMA model captures specific lagged effects, the HAR model offers a more comprehensive view, particularly for cryptocurrencies with pronounced short- and medium-term volatility dynamics.

4.2. Out-of-sample recursive window forecasting results

The out-of-sample forecasting results for Bitcoin (BTC), Ethereum (ETH), Litecoin (LTC), and Ripple (XRP) are assessed based on Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). These metrics allow us to identify the most accurate model for each coin and compare how well the out-of-sample forecasting performance.

• Bitcoin (BTC)

For Bitcoin, the HAR model performs best in terms of MAPE (2.89) and MAE (2.37), while also achieving relatively low MSE (95.52). This indicates that the HAR model is the most accurate in forecasting Bitcoin realized volatility. Despite having a long memory model, ARFIMA does not perform well in the case of Bitcoin. In general, both the HAR and ARMA models demonstrate strong out-of-sample performance, closely aligning with their in-sample fit results. This suggests that capturing both lagged and aggregated lag components enhances the predictive accuracy for Bitcoin, with the HAR

model proving particularly useful for minimizing loss functions. Those findings are supported by the MCS procedure results in the fifth and sixth columns, which are significant with being over 0.20 critical value.

- **Ethereum (ETH)**

Ethereum exhibits similar results, which is in line with BTC results. For instance, the HAR model again shows superior forecasting performance with its accurate MAE (3.24) and MAPE (3.44) values, the lowest among the models, which indicates that it generates the out-of-sample forecasts of the ETH's realized volatility more accurately. On the other hand, the AR model achieves the lowest MSE (72.51) value against the HAR model's performance on MAE and MAPE. This suggests simple AR model as a competing counterpart for the HAR model in the case of ETH. This out-of-sample performance aligns with the in-sample results, where the HAR model captured significant weekly and monthly components for Ethereum. Thus, the HAR model's strength in predicting both short- and medium-term movements could be a better and robust choice over simple AR model for Ethereum.

Table 3: Out-of-sample recursive window forecasting results of competing models

BITCOIN	MSE	MAE	MAPE	MCS (MSE)	MCS (MAE)
AR(1)	79.466	2.514	3.623	1.000	0.000
ARFIMA(1,d,1)	102.062	2.852	5.122	0.000	0.000
ARMA(1,1)	90.914	2.463	3.058	0.000	0.000
HAR	95.516	2.373	2.890	0.000	0.690
ETHEREUM	MSE	MAE	MAPE	MCS (MSE)	MCS (MAE)
AR(1)	72.510	3.272	4.094	1.000	0.000
ARFIMA(1,d,1)	80.860	3.589	4.667	0.000	0.000
ARMA(1,1)	77.137	3.357	3.573	0.000	0.000
HAR	92.595	3.239	3.435	0.000	0.580
LITECOIN	MSE	MAE	MAPE	MCS (MSE)	MCS (MAE)
AR(1)	199.266	4.947	1.831	0.000	0.000
ARFIMA(1,d,1)	198.888	5.490	2.208	0.000	0.000
ARMA(1,1)	189.750	4.558	1.519	1.000	0.000
HAR	201.081	4.437	1.439	0.000	0.680

RIPPLE	MSE	MAE	MAPE	MCS (MSE)	MCS (MAE)
AR(1)	287.831	6.596	2.895	0.000	0.000
ARFIMA(1,d,1)	294.657	7.426	4.011	0.000	0.000
ARMA(1,1)	285.398	6.428	2.499	0.000	0.000
HAR	281.102	6.258	2.235	1.000	0.740

- **Litecoin (LTC)**

In Litecoin's case, the HAR model again achieves the lowest MAE (4.43) and MAPE (1.44), while the ARMA model has the lowest MSE (189.75) value, which are underpinned by the MCS test results in the fifth and sixth columns of corresponding coin. The HAR model's strength in MAPE highlights its accuracy in percentage-based predictions, while the ARMA model excels in minimizing squared errors, making both models viable choices depending on the focus of forecasting precision. This is consistent with the in-sample findings, where the ARMA model captures short-term dynamics effectively for Litecoin, but the HAR model is also seen as a reasonable fit.

- **Ripple (XRP)**

For Ripple, the HAR model emerges as the top performer with the lowest MSE (281.10), MAE (6.26), and MAPE (2.24) values, making it the best model across all error metrics for XRP. This strong out-of-sample performance mirrors the in-sample results, where the HAR model effectively captures both weekly and monthly components for Ripple. The consistency between in-sample and out-of-sample results further supports the HAR model as the most reliable choice for forecasting Ripple's realized volatility. While ARFIMA model fits well in the in-sample analysis, it does not perform in the out-of-sample period.

5. SUMMARY AND CONCLUSION

This study provides valuable insights into the realized volatility dynamics of Bitcoin, Ethereum, Litecoin, and Ripple, emphasizing the importance of model selection in capturing unique volatility characteristics. The HAR model consistently outperforms other models, particularly in out-of-sample forecasts for each cryptocurrency, as reflected by its superior MSE, MAE, and MAPE values. Following the HAR model, ARMA model also performs well, which demonstrates strong out-of-sample performance over ARFIMA and AR models. These results suggest that the HAR model's ability to incorporate multiple time scales—daily, weekly, and monthly—enables it to more accurately

forecast realized volatility, especially for cryptocurrencies with pronounced medium-term volatility trends.

Our findings align with prior literature that advocates for the inclusion of long-memory and multi-scale components in volatility modelling for financial assets, particularly in high-frequency and volatile markets such as cryptocurrency. While the AR and ARMA models demonstrate some success in capturing short-term autoregressive structures, their forecasting accuracy was generally surpassed by the HAR model's comprehensive approach. Additionally, the ARFIMA model, which accounts for long-memory characteristics, performs poorly in the period of out-of-sample forecasting.

Overall, this study highlights the value of using HAR model for capturing complex realized volatility patterns in cryptocurrencies. By proving robust forecasting accuracy, the results of this work includes relevant information in risk management and trading strategies, supporting more informed decision-making in the rapidly evolving cryptocurrency market. Future research could further explore hybrid modelling approaches or apply these findings to other digital assets, contributing to an increasingly nuanced understanding of cryptocurrency market dynamics.

References:

- Andersen, T. G. & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4), 885-905.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modelling and Forecasting Realized Volatility. *Econometrica*, 71(2), 579-625.
- Balcilar, Mehmet, Elie Bouri, Rangan Gupta, and David Roubaud. (2017). Can Volume Predict Bitcoin Returns and Volatility? A Quantiles-Based Approach. *Economic Modelling* 64: 74–81.
- Bariviera, A. F. (2017). The inefficiency of Bitcoin revisited: A dynamic approach. *Economics Letters*, 161, 1-4.
- Bouoiyour, Jamal, and Refk Selmi. (2015). Bitcoin Price: Is It Really That New Round of Volatility Can Be on Way? *MPRA Paper*.
- Bouoiyour, Jamal, and Refk Selmi. (2016). Bitcoin: A Beginning of a New Phase? *Economics Bulletin* 36: 1430–40.
- Corbet, S., Lucey, B., Urquhart, A., & Yarovaya, L. (2019). Cryptocurrencies as a financial asset: A systematic analysis. *International Review of Financial Analysis*, 62, 182-199.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174-196.
- Granger, C.W.J. and Joyeux, R. (1980). An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis*, 1(1), 15–29.
- Hol, E., & Koopman, S. J. (2002). Stock index volatility forecasting with high frequency data (No. 02-068/4). *Tinbergen Institute Discussion Paper*.
- Katsiampa, P. (2017). Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters*, 158, 3-6.
- Katsiampa, P. (2019). An empirical investigation of volatility dynamics in the cryptocurrency market. *Journal of Risk and Financial Management*, 12(1), 41.
- Liu, L. Y., Patton, A. J., & Sheppard, K. (2015). Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes. *Journal of Econometrics*, 187(1), 293-311.
- McMillan, D. G., & Speight, A. E. H. (2010). Daily volatility forecasts: New evidence on the accuracy of GARCH models. *International Journal of Forecasting*, 26(2), 276-288.
- Müller, U. A., Dacorogna, M. M., Davé, R. D., Olsen, R. B., Pictet, O. V., Von Weizsäcker, J. E. (1997). Volatilities of different time resolutions, Analysing the dynamics of market components, *Journal of Empirical Finance*, 4 (2-3) (1997), pp. 213-239.

- Phillip, A., Chan, J., & Peiris, S. (2018). A new look at cryptocurrencies. *Economics Letters*, 163, 6-9.
- Sahiner, M., McMillan, D. G., & Kambouroudis, D. (2023). Do artificial neural networks provide improved volatility forecasts: Evidence from Asian markets. *Journal of Economics and Finance*, 47(3), 723-762.
- Troster, Victor, Aviral Kumar Tiwari, Muhammad Shahbaz, and Demian Nicolás Macedo. (2019). Bitcoin Returns and Risk: A General GARCH and GAS Analysis. *Finance Research Letters* 30: 187–93.

