NOVA IMS Information Management School

MDSAA

Mestrado em

Data Science and Advanced Analytics

When Bitcoin Sneezes, Does Cardano Catch a Cold? Spillovers Between Brown and Green Cryptos:

A Graph Neural Networks approach using intraday data

Francisco Catarino Pires Lacerda

Master Thesis

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação
Universidade Nova de Lisboa

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação
Universidade Nova de Lisboa

When Bitcoin Sneezes, Does Cardano Catch a Cold? Spillovers Between Brown and Green Cryptos:

A Graph Neural Networks approach using intraday data

by

Francisco Catarino Pires Lacerda

Master Thesis presented as partial requirement for obtaining the Master's degree in Data Science and Advanced Analytics, with a specialization in Business Analytics

Supervised by

Professor Ian Scott, PhD, NOVA Information Management School

Professor Bruno Damásio, PhD, NOVA Information Management School

November, 2025

STATEMENT OF INTEGRITY

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism, any form of undue use of information or falsification of results along the process leading to its elaboration. I further declare that I have fully acknowledged the Rules of Conduct and Code of Honor from the NOVA Information Management School.

Lisboa, 2025

Francisco Lacerda

DEDICATION

Dedico este trabalho à minha família e em especial à minha tia Fernanda Lacerda que sempre puxou por mim para acabar este projeto, especialmente na fase final, quando conciliar a mudança para Barcelona e o início do meu primeiro emprego se tornou particularmente desafiante.

Dedico também este trabalho às amizades para a vida que construí na Nova IMS, recordarei para sempre com carinho os 2 anos que passámos juntos.

No entanto, o maior agradecimento tem de ir para o Pedro Fidalgo, treinador e amigo, pela constante disponibilidade e auxílio ao longo de todo este projeto.

Recordarei sempre as nossas brainstorms pós treinos, nas quais, graças à sua criatividade inesgotável, me ajudou a pensar formas engenhosas e interessantes de ultrapassar cada obstáculo que foi surgindo ao longo deste trabalho.

ACKNOWLEDGEMENTS

I would like to express my gratitude to my supervisors, Professor Ian Scott and Professor Bruno Damásio, for their guidance throughout the development of this thesis.

Beyond the academic support, I am especially grateful for the personal insights they shared on how to approach research not only with rigor but also curiosity.

Professor Ian´s reflections on how to structure our time, protect our focus, and maintain a healthy relationship between work and personal life have had a lasting impact on me.

I am sincerely thankful for their mentorship, which extended well beyond the technical aspects of this project.

It has been a privilege to learn from both.

ABSTRACT

This thesis investigates volatility spillover effects in the cryptocurrency market by applying graph-based volatility forecasting models to a portfolio of Proof-of-Work (PoW) and Proof-of-Stake (PoS) cryptocurrencies. Four models are developed and compared, HAR, GHAR, GNNHAR-1L, and GNNHAR-2L to address three core research questions: (1) Do volatility spillovers exist among cryptocurrencies? (2) Are these spillovers nonlinear? and (3) Do indirect (multi-hop) neighbors contribute meaningfully to volatility transmission? High-frequency price data sampled at 5-minute intervals over the 2023–2024 period is used to construct realized volatility measures, HAR-style lagged features, and dynamic graph structures estimated via rolling-window Graphical LASSO. Model performance is evaluated using the Model Confidence Set (MCS) procedure to ensure statistically robust comparisons. The empirical results show that GNNHAR-1L consistently outperforms both GHAR and GNNHAR-2L across all assets, providing strong evidence for the presence of nonlinear spillover effects. By contrast, incorporating second-hop neighbors in GNNHAR-2L does not yield significant forecasting improvements, indicating limited indirect spillover dynamics in this dense volatility network. Furthermore, no meaningful differences are observed between PoW ("brown") and PoS ("green") cryptocurrencies in terms of volatility transmission behavior. Instead, the analysis uncovers clear volatility-based clustering: assets with similar volatility levels exhibit stronger connections and more pronounced spillover linkages. These findings contribute to the growing literature on graph-based financial modeling and provide new insights into the structural and nonlinear nature of volatility interdependence in cryptocurrency markets. The study concludes with several directions for future research, including alternative graph-construction methods, multi-horizon volatility modeling, cross-market spillover analysis, and the integration of investor heterogeneity into network-based volatility frameworks.

KEYWORDS

Cryptocurrency; Bitcoin; Realized Volatility; Graph Neural Networks; Intraday data

Sustainable Development Goals (SDG):

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS AND ACRONYMS

ADA – Cardano

ALGO – Algorand

Altcoins – All cryptocurrencies other than Bitcoin

AVAX – Avalanche

BCH – Bitcoin Cash

BTC – Bitcoin

CELO – Celo

CFX – Conflux

CKB – Nervos Network

DeFi – Decentralized Finance is an emerging peer-to-peer system attempting to remove third parties and centralized institutions from financial transactions

DOGE – Dogecoin

DOT – Polkadot

EGLD – MultiversX

ETH – Ethereum

ETC – Ethereum Classic

FOMO – Fear of Missing Out

GHAR – Graph Heterogeneous Autoregressive

GLASSO - Graphical Least Absolute Shrinkage and Selection Operator

GNN – Graph Neural Networks

GNNHAR – Graph Neural Networks Heterogeneous Autoregressive

HAR – Heterogeneous Autoregressive

INJ – Injective

LTC – Litecoin

MCS – Model Confidence Set

MINA – Mina

ML – Machine Learning

MLP – Multilayer Perceptron

PoS – Proof of Stake

PoW – Proof of Work

RV – Realized Volatility

RVN – Ravencoin

SC – Siacoin

SOL – Solana

TRX – Tron

XTZ – Tezos

ZEC – Zcash

1. INTRODUCTION

The beginning of cryptocurrencies and DeFi

The first cryptocurrency (Bitcoin) emerged in 2009 as a direct response to the vulnerabilities exposed during the 2008 global financial crisis (Nakamoto, 2008). The crisis, marked by the collapse of major financial institutions, constituted a breakdown of trust in centralized monetary systems (Scalera & and Dixon, 2016)(Zingales, 2011), simultaneously showing the need and creating a path for alternative financial mechanisms that could operate independently of traditional intermediaries to be born. Bitcoin (BTC), the first cryptocurrency, was introduced as a decentralized peer-to-peer digital currency, offering an innovative solution to the flaws inherent to centralized financial systems. Built on blockchain technology, BTC enabled secure transactions without the need for intermediaries, creating an environment of transparency, and decentralization. These key features are the foundation of a new financial concept independent from traditional institutions, DeFi.

After Bitcoin's success, a wave of alternative cryptocurrencies (altcoins) like Ethereum, Litecoin among others emerged. These assets introduced new features, such as smart contracts and decentralized finance (DeFi) platforms, further expanding the market and attracting speculative interest (Krupa et al., 2021). This interest is in part due to the demonstrated extreme price volatility cryptocurrencies are prone to. (Chaim & Laurini, 2018) evidence that BTC is much more volatile than traditional financial assets stating that cryptocurrencies have a high unconditional volatility and are subject to large price swings.

(Bouri et al., 2019) discuss the "price explosivity" in BTC, reinforcing the idea of extreme fluctuations in cryptocurrency prices and co-explosivity while (Nekhili, 2020) emphasized BTC´s speculative trading nature and the role of speculators in the BTC's futures market, concluding BTC market was (at least at the time) price inefficient.

The relatively unregulated nature of cryptocurrency markets (Krupa et al., 2021) (Keidar & Blemus, 2018), coupled with hype cycles and media attention, can fuel speculative bubbles evidence by (Haykir & Yagli, 2022) (Ben Osman et al., 2024) (Chaim & Laurini, 2018). High-profile events, such as regulatory decisions, hacking incidents, or endorsements by influential figures, have often triggered significant price fluctuations. China's Crypto Ban (May-June

2021) contributed to a significant decline in BTC while in February 2021 Tesla's Bitcoin purchase announcement pushed prices to over $50,000.

Over time, institutional investors began participating in cryptocurrency markets, legitimizing them as an asset class. Products like cryptocurrency futures (first introduced for Bitcoin in late 2017 by Chicago Board Options Exchange (CBOE)), made it easier for institutions and investors to speculate on price movements, (Bouoiyour & Selmi, 2019) found that futures trading significantly drove up the price of Bitcoin immediately after the announcement day of the futures contracts while (Sebastião & Godinho, 2020) deemed CBOE BTC futures an effective hedging tool for hedging BTC at a daily horizon.

The genesis and brief history of the crypto market presented above along with evidence of this market assets interconnections, (Gkillas et al., 2024) shows evidence of co-jumping behavior among cryptocurrencies while (Vardar et al., 2022) provides evidence for volatility transmission between 8 major cryptocurrencies and (Balcilar & Ozdemir, 2023) found evidence for the existence of spillover between 11 major cryptocurrencies both in bullish, bearish and mixed market conditions, are important to understand the motivations behind the objectives of this thesis, which are to further investigate spillover effects nature on the crypto market regarding distinct types of crypto assets.

Spillover Effects

One of the most intriguing phenomena in financial markets is the spillover effect, the idea that volatility in one asset or market can be transmitted to others. In traditional markets, volatility spillovers have been observed between commodities and equities markets (Naeem et al., 2023) and equity and foreign exchange markets (Warshaw, 2020), especially during periods of market stress.

There is also evidence of the existence of spillover effects between the crypto market and other traditional financial markets, (Pacelli et al., 2025) highlights existence of spillover effects between crypto market along with global equity indexes from Europe, the United States, and China, (Mensi et al., 2023) evidences spillover between cryptocurrencies and CBOE uncertainty indices while (Hung, 2022) and (Akin et al., 2024) find volatility spillover between BTC and S&P500, gold and crude oil.

In the context of cryptocurrencies, which have been observed to operate within an interconnected ecosystem by (J. Liu et al., 2024), the authors show evidence of strong contemporaneous dependencies between cryptocurrencies and the importance of spillover effects modelling to understand the relationships between assets. Understanding these effects is essential for accurate risk modeling and for developing robust portfolio strategies.

As for this thesis, it will focus on the underlying nature of the relationships between two distinct groups of cryptocurrencies which are going to be labelled “Brown” and “Green” cryptocurrencies.

Distinction between Green and Brown Cryptocurrencies

It is important to note that the cryptocurrencies used for this study were categorized into two distinct groups regarding energy consumption. The idea stems from (Wang et al., 2024) which does a similar distinction between cryptocurrencies with energy intensive validation mechanisms, labeled as “brown”, and cryptocurrencies less energy reliant, labeled as “green”. Prior studies suggest that these structural differences may influence market behavior, (Wang et al., 2024) found differences in jump behavior regarding the Energy Uncertainty Index (EUI) developed by (Dang et al., 2023), between the two groups of cryptocurrencies.

In the case of the present study, the first group is categorized for its reliance on energy intensive mining processes to validate transactions and maintain network security (Proof-of-Work).

Proof-of-Work (PoW) has miners competing to solve complex mathematical puzzles by using computational power, the miner to solve the problem first gets to add a new block of transactions to the blockchain. After the network verifies the solution is in fact correct, the miner is rewarded with new coins.

The second group is composed of cryptocurrencies who are less energy-reliant in comparison to cryptocurrencies that use PoW as their validation method because they use Proof-of-Stake (PoS) or staking-based consensus blockchains.

PoS does not require miners to solve complex puzzles using massive computational power. Instead, it selects validators to confirm transactions based on the amount of cryptocurrency they stake (lock up as collateral), rather than using energy-intensive mining.

Because of this, PoS or staking-based consensus blockchain cryptocurrencies are often perceived as “green” alternatives to PoW assets (“brown”). These differences in energy consumption can have significant implications in the way these two different groups of cryptocurrencies behave to changes in the market as evidenced by (Wang et al., 2024) and (Gualpa et al., 2023). Moreover, although (Milunovich, 2022) does not explicitly mention energy consumption as a factor influencing the market behavior or connectedness between these two groups of cryptocurrencies, it does highlight PoW cryptocurrencies appear to be more strongly connected within the network compared to PoS and, on average, PoW coins export more uncertainty to other cryptocurrencies (stronger volatility spillover effects).

Although prior research has identified differences in the sensitivity of PoW (“brown”) and PoS (“green”) cryptocurrencies to macroeconomic uncertainty, no existing study examines whether these technological differences translate into distinct patterns of volatility spillover. The Green vs Brown distinction has therefore not yet been analyzed through a modern network-based volatility model such as GHAR or GNNHAR. This constitutes a clear gap: despite growing policy and investor interest in the environmental footprint of cryptocurrencies, it remains unknown whether cleaner consensus mechanisms exhibit different contagion dynamics and whether they amplify or mitigate systemic volatility transmission.

Motivated by this lack of evidence, the present study investigates whether these technological and energy-related characteristics influence how volatility spillovers propagate across the cryptocurrency ecosystem.

Main Research Question

The overarching research question guiding this thesis is:

Do Brown (PoW) and Green (PoS) cryptocurrencies differ in how they transmit and receive volatility spillovers?

Addressing this question first requires establishing the fundamental properties of spillover dynamics within the crypto market. Accordingly, the analysis focuses on three supporting research questions.

Research Questions

RQ1 - Do volatility spillovers exist between cryptocurrencies?

Before assessing differences between Brown and Green assets, it is necessary to determine whether volatility transmission occurs at all, and whether incorporating information from neighboring assets improves forecasting performance relative to a univariate benchmark.

RQ2 - Are volatility spillovers nonlinear?

If spillovers are present, the next step is to examine whether they follow a linear pattern or whether their intensity varies nonlinearly with market conditions, shock magnitude, or other contextual factors.

RQ3 - Do indirect (multi hop) neighbors contribute to volatility transmission? Volatility shocks may propagate only through immediate connections or may spread more broadly across multiple steps in the network. Understanding the depth of propagation is essential for characterizing systemic interconnectedness.

Modeling Strategy

To address these research questions, four forecasting models are employed:

(1) the heterogeneous autoregressive model (HAR),

(2) the Graph HAR model (GHAR),

(3) a one-layer Graph Neural Network HAR (GNNHAR-1L), and

(4) a two-layer Graph Neural Network HAR (GNNHAR-2L).

These models are compared in a structured manner that aligns directly with the research questions:

• The HAR vs GHAR comparison tests **RQ1** by evaluating whether incorporating linear neighbor information yields improved forecasts.

• The GHAR vs GNNHAR-1L comparison addresses **RQ2** by assessing whether nonlinear aggregation enhances predictive accuracy.

• The GNNHAR-1L vs GNNHAR-2L comparison informs **RQ3** by testing whether expanding the receptive field to include indirect neighbors provides additional forecasting benefits.

Evaluating model performance at the individual asset level further enables an investigation of whether spillover characteristics differ between Brown and Green cryptocurrencies.

Purpose of the Study

By combining high-frequency realized volatility measures, time-varying network structures, and graph neural network-based forecasting models, this thesis aims to: (1) characterize the structure and nature of volatility spillovers in the cryptocurrency market; (2) determine the role of nonlinear and multi-hop transmission mechanisms; and (3) provide empirical evidence on whether Brown and Green cryptocurrencies exhibit different volatility spillover behaviors.

2. LITERATURE REVIEW

WHAT HAS BEEN DONE

In the last years, the study of cryptocurrency volatility has attracted considerable attention as the market continues to evolve with rapid technological and financial innovations. For this reason, (Ahmed et al., 2024) article, which analyses 164 other articles from some of the most well-regarded scientific journals, provides the understanding of what has lately been done in this field of study and what research gaps and improvements to current studies are still there to be made.

(Ahmed et al., 2024) focused on 9 stylized facts, which were Volatility Clustering, Volatility Persistence, Asymmetric Volatility, Leverage effect, Volatility Spillover, Mean Reversion, Structural Breaks, Extreme Volatility and Tail Dependence.

From these stylized facts and considering the research that has been done posterior to this article publication the focus of this thesis will be centered on the stylized fact known as Volatility Spillovers and their nature in the crypto market, which refers to the phenomenon where fluctuations in the volatility of one asset or market influence the volatility of another, often indicating the transmission of risk or information across financial entities, very little is known about its determinants in the crypto market such as news and investors sentiment, herding behavior and loss aversion.

Despite the extensive documentation of stylized facts, (Ahmed et al., 2024) also identifies other substantial research gaps, particularly concerning the use of high frequency data and machine learning (ML) techniques. The integration of high frequency data and machine learning represents a promising frontier for cryptocurrency volatility research. High frequency data can reveal intraday patterns which could otherwise be smoothed out. When paired with ML models, this data can provide more accurate volatility forecasts, enhance risk management strategies, and uncover hidden drivers of market dynamics. Numerous studies such as (Peng et al., 2024)(Miura et al., 2019)(Mücher, 2022) have leveraged high frequency data to train Neural Networks based models and shown promising results.

Moreover, the use of GNN to model the relationships between different cryptocurrencies and assets will also help us expand our current knowledge of volatility spillover between these assets. We will expand more about this part of our research gap in the upcoming subchapters.

The Use of High-Frequency Data in Realized Volatility Modelling

(Ahmed et al., 2024) alerts to a relative shortage of studies using high frequency data. The smoothness of the data (the use of daily or weekly data) can overshadow certain intricacies in the intraday dynamics of the assets (Arora, 2017) (Petukhina et al., 2021) (Bariviera et al., 2018).

For this reason, the use of high frequency financial data has become foundational in the modelling of asset price volatility, especially through the concept of realized volatility (RV) (Yan & Yan, 2019) (Sheppard & Xu, 2019).

Introduced by (Andersen et al., 2003) and (Barndorff-Nielsen & Shephard, 2002), realized volatility provides a non-parametric, consistent estimator of integrated volatility (which is unobservable) by summing high frequency squared returns.

In contrast, while daily returns aggregate price changes over a 24-hour window and obscure the path of volatility within the day, high frequency returns capture intraday volatility bursts, market reactions to news, liquidity shifts, and short-lived contagion effects (Karam, 2018) (Ammar & Hellara, 2022).

Consequently, RV provides a more responsive and granular representation of actual market activity and higher frequency temporal information, which in turn can materialize into improved forecasting ability (X. Zhang, 2022) (G. Liu et al., 2024).

In the context of this study, the use of RV constructed from high frequency 5-min intervals data when working in the interconnected ecosystem that constitutes the crypto market (Balcilar & Ozdemir, 2023; Gkillas et al., 2024; J. Liu et al., 2024; Vardar et al., 2022) is especially important in the context of graph neural networks (GNNs) for capturing the temporal dynamics of each asset's volatility and leveraging information from cross-asset spillovers across different horizons (daily, weekly, monthly) as shown by (C. Zhang et al., 2025a). And while not using graph neural networks, both (Zhao et al., 2024) and (Z. Xu et al., 2024) recorded improved results by combining GARCH models with a Neural Networks component in comparison to the models in isolation.

The choice to use 5-min sampling intervals for computing RV lies in the balance between information quality and noise reduction. While valuable, high frequency data is also susceptible to microstructure noise (Z. Xu et al., 2024) (Z. M. Li et al., 2020).

Microstructure noise refers to distortions in asset price data that arise not from fundamental market forces, but from the mechanics of how trades are executed in real world financial markets such as bid-ask bounces or latency and reporting delays.

These distortions are more heavily felt when working with very high frequencies (e.g. 1-minute intervals) and can contaminate RV estimates in the form of spurious short-term fluctuations (random noise) (Z. M. Li et al., 2020).

In order to strike the said balance, (L. Y. Liu et al., 2015) shows that a 5-minute interval provides a near optimal tradeoff where it retains sufficient intraday variability to approximate integrated volatility (IV) accurately, while mitigating the adverse effects of market microstructure distortions.

To further add, the study of (C. Zhang et al., 2025a), which provides the foundation for this study’s methodology, also uses 5-min interval data to calculate daily RV.

As a result, both the network topologies used in the HAR, GHAR and GNNHAR models, as well as the asset level input features, are derived from high frequency data.

SPILLOVER EFFECTS

The study of volatility spillovers has gained increasing attention in cryptocurrency research. Unlike traditional financial systems, which are mediated and regulated by centralized institutions, the cryptocurrency ecosystem is decentralized, has no close trading hours (24/7 trading) and global investor participation (Makarov & Schoar, 2022), creating a unique environment where it becomes crucial to investigate, understand and model the relationships between the assets. In the case of this study, the information propagation in the form of volatility spillovers.

In such a decentralized context, cryptocurrencies are often interlinked via trading platforms and shared protocols (e.g. DeFi) (Mensi, Gubareva, et al., 2024)(Mazor & Rottenstreich, 2024), which facilitate the transmission of shocks from one asset to another. For example, a sudden drop in the price of Bitcoin, typically regarded as a bellwether for the crypto market, can trigger correlated volatility responses across altcoins. (Alamaren et al., 2024) investigated the existence of volatility spillovers and pairwise connectedness between cryptocurrencies, energy and tech companies in the US market and found that Bitcoin acted as a net transmitter of volatility to other cryptocurrencies.

(Moratis, 2021) concluded in its study that, out of the largest 30 cryptocurrencies by market cap, Bitcoin was the dominant volatility transmitter in the crypto market. Additionally, (Moratis, 2021) results showed market size (market cap) is not the sole driver of spillover importance as Ripple, the 3rd largest cryptocurrency by market cap at the time, acted as a net receiver.

(Aslanidis et al., 2021) showed evidence of intensifying spillover effects among cryptocurrencies using PCA (Principal Component Analysis).

These studies have all documented strong interdependencies in the volatility behavior of major and minor cryptocurrencies with Bitcoin and Ethereum often functioning in the role of volatility transmitters, while smaller tokens act as recipients. This phenomenon further motivates the development of models capable of capturing cross asset interactions and understanding the network nature of volatility in the crypto market.

As a result, traditional univariate models such as ARCH (R. F. Engle, 1982), GARCH (Bollerslev, 1986) and HAR (Corsi, 2009) or pairwise modeling approaches like Diebold-Yilmaz connectedness indices (Diebold & Yilmaz, 2009) may fail to detect the pathways through which information propagate.

Another widely used framework is the multivariate GARCH family, particularly the BEKK-GARCH (R. F. Engle & Kroner, 1995) and DCC-GARCH (R. Engle, 2002) models. These approaches allow for the modeling of time-varying conditional covariances and volatilities across multiple assets. By examining the dynamic evolution of the conditional covariance matrix, researchers can infer how volatility shocks in one market spillover into others.

Although these models are statistically rigorous, they cannot capture indirect (multi hop) spillovers. Furthermore, BEKK and DCC-GARCH models capture spillover effects through linear combinations of past returns and volatilities. Their functional forms assume proportional and symmetric transmission, restricting the dynamics to linear relationships between assets.

This creates a compelling argument for adopting network-based frameworks, such as the one used in this thesis, to better understand and forecast volatility in a highly connected digital asset landscape.

THE USE OF GRAPH NEURAL NETWORKS

As the complexity and interconnectedness of financial systems have increased (Raddant & Kenett, 2021), and as a consequence, traditional econometric models have struggled to scale and capture the nuanced nonlinear and multi hop dependencies that characterize modern asset networks (de la Concha et al., 2018). In response, recent literature has introduced graph based frameworks to model volatility spillovers, treating financial markets as structured systems where assets are nodes and spillover relationships are edges in a graph.

One notable development in this direction is the Graph HAR (GHAR) model. In its study, (C. Zhang et al., 2025b) extended the traditional Heterogeneous Autoregressive (HAR) (Corsi, 2009) model by incorporating graph based information into the forecasting equation. The GHAR model introduces a static adjacency matrix, derived from a precision (inverse covariance) matrix using Graphical Lasso (GLASSO), to represent the spillover network structure. Each asset's realized volatility is forecasted not only based on its own past values, but also through a linear aggregation of its direct neighbors' volatility features, allowing GHAR to model local, 1-hop spillover effects.

However, while GHAR marks a significant advancement by introducing cross-sectional dependence into HAR, it remains limited to linear interactions and immediate (1-hop) neighbors. This restricts its ability to capture more complex dynamics such as nonlinear volatility transmission and indirect spillovers that propagate through multi-step network pathways, which might incur in the loss of valuable information by the model.

To overcome these limitations, recent studies have applied Graph Neural Networks (GNNs) to financial volatility forecasting. Building on the GHAR foundation, (C. Zhang et al., 2025a) proposed the GNNHAR model, which integrates graph convolutional layers with HAR features to enable nonlinear transformation (GNNHAR-1L) and multi hop message passing across the asset graph (GNNHAR-2L). In this framework, the depth of the GNN determines the receptive field, a 1-layer GNNHAR captures direct neighbors, while a 2-layer GNNHAR incorporates 2-hop (indirect) neighbors, allowing for the modeling of volatility propagation through the network.

Several studies employing graph based models show promising forecasting improvements compared to traditional such as the above mentioned GARCH. (Kumar et al., 2024) outperforms traditional methods like GARCH and MLP using Temporal GAT model for volatility forecasting, effectively capturing temporal dependencies and interdependencies between global market indices.

(Son et al., 2023) outperformed multivariate GARCH using spatial-temporal GNN, attributed to the model’s ability to incorporate volatility spillover effects among eight of the biggest global market indices.

(Zhou et al., 2025) achieved improvements in volatility forecasting by leveraging EMGNN (Explainable Multilayer GNN) model ability to capture multi-scale interactions and dynamic spillovers on cryptocurrency markets using a graph structure, compared to HAR, Gaussian process regression (GPR), Random Forest (RF), LightGBM and LSTM.

(Brogaard, 2024) demonstrated GNN and GNN-attention outperformed other ML algorithms, including RF, Support Vector Machines (SVM), Gradient Boosting Decision Tress and even Convolutional Neural Networks (CNN) under the same training configurations, highlighting the importance of capturing inter-asset dependencies.

By introducing SpotV2Net, a multivariate intraday spot volatility forecasting model based on a graph attention network architecture, (Brini & Toscano, 2024) results show that SpotV2Net yields more improvements in multi-step intraday volatility forecasting in comparison to HAR, XGB and LSTM models. Also, SpotV2Net outperformed SpotV2Net-NE (NE=no edge), demonstrating that multi step intraday volatility forecasting add value in comparison to just direct spillover via neighbors.

(Heß & Damasio, 2025) use ANN and CNN models to overcome the limitations of traditional methods about the evolution of the banking system market and technology driven changes.

Taking into consideration the studies above mentioned and the fact (C. Zhang et al., 2025a) evidenced an improvement in the forecasting when employing GNNHAR models compared to GHAR, this thesis builds directly on the GNNHAR framework, applying it to the cryptocurrency market to investigate whether “green” and “brown” cryptoassets exhibit different spillover behaviors, and whether these behaviors are better captured through nonlinear modelling and multi hop network structures.

3. METHODOLOGY

3.1 Data Overview

The dataset spans from January 1, 2023 to December 31, 2024, covering two full calendar years of high frequency cryptocurrency price data.

The study period (2023-2024) was chosen to reflect a phase of market normalization following the extreme uncertainty and monetary interventions associated with the COVID-19 pandemic (Lahmiri & Bekiros, 2020) (Asiri et al., 2023) (Ünlü & Bayram, 2024). This window captures the behavior of cryptocurrency markets under more stable macroeconomic conditions and after key structural events (e.g. Ethereum’s shift to PoS in 2022), providing a more consistent environment for analyzing cross asset volatility dynamics. Despite this, it is important to note an increase in uncertainty in the market as we got closer to the US presidential election in early November 2024, as it is shown in the RV graphs (Appendix J : Daily Realized volatility plots).

The data was retrieved from the official Binance API, providing cryptocurrencies prices in Tether (USDT, a USD-pegged stablecoin) with a granularity of 5-min intervals.

3.2 Realized Volatility Construction

To model and forecast the volatility of cryptocurrencies, this study employs daily realized volatility (RV) as the target variable. RV is a non-parametric estimate of the latent integrated volatility (IV) of an asset, constructed by aggregating squared high frequency intraday returns over a fixed interval. This approach was first formalized by (Andersen & Bollerslev, 1998)*)* and later expanded in econometric detail by (Barndorff-Nielsen & Shephard, 2002) and (Andersen et al., 2003).

Let $P_{i,t}^{(m)}$ denote the price of asset *i* at the *m*-th 5-min interval of day *t*. The realized volatility of asset *i* on day *t* is calculated as:

$$RV_{i,t} = \sum_{m=1}^{M} (logP_{i,t}^{(m)} - logP_{i,t}^{(m-1)})^2 \qquad (1)$$

where *M* is the total number of 5-min intervals in a trading day (for 24-hour crypto markets, *M* = 288).

The integrated variance over a day t is defined as:

$$IV_t = \int_t^{t+1} \sigma_s^2 ds \qquad (2)$$

where σ_s^2 is the instantaneous variance at time s. This is unobservable in practice.

However, the quadratic variation of the log price process over a day can be consistently estimated using observed high frequency data, and under ideal conditions (e.g. no microstructure noise), this quadratic variation converges to the integrated variance as the sampling frequency increases ($M \to \infty$) (Andersen & Bollerslev, 1998) (Barndorff-Nielsen & Shephard, 2001).

This convergence property provides a robust theoretical foundation for using RV as a proxy for the unobservable integrated variance (Barndorff-Nielsen & Shephard, 2001). Furthermore, (Andersen et al., 2003) validate the forecasting power of RV and introduce the HAR-RV framework.

The use of log returns is preferred over simple returns in financial modeling due to several theoretical and practical advantages. They are additive over time, allowing for straightforward aggregation across periods. Log returns tend to exhibit a distribution that is closer to normality with the log transformation helping to reduce skewness and kurtosis, especially at high frequencies. Log returns are scale invariant, making them comparable across assets with different price levels and reflect continuous compounding, a foundational concept in financial mathematics. For these reasons and the fact (C. Zhang et al., 2025a), the paper which inspired this thesis methodology, uses log returns in the computation of realized volatility, this thesis uses log returns to ensure both statistical robustness and theoretical consistency.

In the context of this study, the use of RV lagged values serves as the input features for the HAR, GHAR and GNNHAR models as they reflect the heterogeneous memory structure proposed by the HAR model (Corsi, 2009). These lagged features help to avoid look-ahead bias, all rolling features are constructed as backward-looking averages, using only information available prior to the prediction day. This ensures that the model learns from historical patterns without incorporating any future data, maintaining causal validity and simulating realistic out-of-sample forecasting conditions.

These lagged features are computed separately for each cryptocurrency over three-time horizons (daily, weekly and monthly) and used to construct the node feature matrix $X_t \in R^{N \times 3}$

in the forecasting models. The lagged RV is aggregated into daily (d), weekly (w), and monthly (m) components:

$$RV_{t-1}^{(d)} = \frac{1}{1} RV_{t-1} \qquad (3)$$

$$RV_{t-1}^{(w)} = \frac{1}{7} \sum_{i=1}^{7} RV_{t-i} \qquad (4)$$

$$RV_{t-1}^{(m)} = \frac{1}{30} \sum_{i=1}^{30} RV_{t-i} \qquad (5)$$

These three components form the node feature matrix $X_t \in R^{N\times 3}$ where each row $x_{i,t} = [RV_{i,t}^{(d)}, RV_{i,t}^{(w)}, RV_{i,t}^{(m)}]$ on day *t* is the feature vector for cryptocurrency *i* and N is the number of cryptocurrencies (nodes).

They represent volatility at different horizons and are designed to capture both short and long term memory effects in volatility dynamics (Corsi, 2009). This matrix is the input for the HAR, GHAR and GNNHAR models.

3.3 Graph Construction using GLASSO

3.3.1 From Return Covariance to Graph Structure

To capture the evolving nature of volatility spillovers between cryptocurrencies, this study employs a dynamic adjacency matrix constructed via a rolling Graphical Lasso (GLASSO) procedure.

GLASSO was proposed by (Friedman et al., 2008), it estimates a sparse precision matrix (the inverse of the covariance matrix (Σ)) of asset returns. The key idea is that in a multivariate Gaussian framework, the off-diagonal elements of the precision matrix $\Theta = \Sigma^{-1}$ correspond to conditional dependencies: if $\theta_{i,j} = 0$, then assets *i* and *j* are conditionally independent given all other assets (they are not connected).

The GLASSO estimator solves the following convex optimization problem:

$$\widehat{\Theta} = \arg\min_{\Theta \succ 0} \{tr(S\Theta) - \log\det(\Theta) + \lambda \|\Theta\|_1\} \qquad (6)$$

where:

- S is the empirical covariance matrix of the asset return series.
- λ is a regularization parameter that controls sparsity via cross-validation.

- $\|\Theta\|_1$ denotes the elementwise L^1 norm.
- $\widehat{\Theta}$ is the estimated sparse precision matrix.

The L^1 penalty encourages many off-diagonal elements of Θ to be exactly zero, resulting in a sparse and interpretable graph.

Nodes represent individual cryptocurrencies and edges represent conditional dependency relationships between their return series. These graphs form the backbone of the GHAR and GNNHAR models by encoding which assets may influence each other through volatility transmission.

In the methodology proposed by (C. Zhang et al., 2025a) paper, these adjacency matrices are not constructed on a daily basis but rather updated at fixed intervals corresponding to the model's forecast windows. A new adjacency matrix is computed once per forecast window, which spans 22 trading days (approximately one calendar month of trading data as the markets close for the weekends). At the beginning of each forecast window referred to as the forecast anchor date (e.g. July 1, 2011), a single adjacency matrix is constructed and then reused across all samples (training, validation, and testing) within that forecast period. This design ensures temporal consistency in the graph structure. To estimate each adjacency matrix, the model uses a rolling historical window of daily log returns for all assets included in the universe.

It is important to highlight some key differences between (C. Zhang et al., 2025a) and this thesis. For starters, in this thesis the adjacency matrices are estimated at a finer temporal resolution, using a sliding window approach. A new matrix is generated every day using a trailing window of fixed length (100 days). This leads to a larger set of time-varying adjacency matrices, capturing daily changes in the asset dependence structure rather than monthly, enabling better alignment with real market dynamics. (Jiang & Pu, 2023) highlights the advantages of using more regularly updated (finer-grain) temporal graphs although also emphasizing the importance of task-specific tuning, while (Javaheri et al., 2025) empirical results demonstrate that the use of more frequently updated graphs enhances the learning of dynamic relationships between financial assets, particularly in environments characterized by heavy-tailed data.

These studies support the approach taken in this thesis, suggesting that updating graph structures on a daily basis provides finer temporal resolution, which is particularly valuable in highly volatile environments such as cryptocurrency markets.

The paper explicitly uses Graphical LASSO with Cross-Validation (GLASSO-CV) to automatically select the optimal regularization parameter λ. By default, our framework also uses Glasso-CV, which performs internal cross-validation over a grid of regularization parameters.

Another key difference has to do with the use of backfilling, when GLASSO fails or returns unstable results (NaNs), (C. Zhang et al., 2025a) drops those estimations while in this thesis, a backfilling strategy is incorporated where if the estimation fails for a given day, the last valid adjacency matrix is reused. Although not quite the same, our idea stems from (Z. L. Li et al., 2023), where the static graph is incorporated as an inductive bias during dynamic graph construction. This mechanism serves to stabilize the dynamic adjacency matrix by grounding it in long-term structural patterns. However, while their model falls back implicitly on a globally static structure, this thesis introduces a more adaptive solution: a backfilling strategy that reuses the most recently valid dynamically estimated adjacency matrix. Rather than reverting to an initial static baseline, our approach ensures that the fallback reflects more recent market conditions, thereby maintaining temporal continuity and preserving short-term dependencies.

This ensures continuity and avoids disruptions in downstream model training, particularly important when estimating hundreds of matrices over extended periods and adds practical robustness to graph based modelling in highly volatile environments like cryptocurrency markets.

In regard to temporal alignment and reusability, (C. Zhang et al., 2025a) reuses the same matrix across all input samples in a given forecast window, regardless of whether they belong to training or testing sets. In our implementation, since a new matrix is computed for each day, the adjacency matrices are dynamically aligned to each input sample by matching the date of the input with the available adjacency matrix of that day based on estimation timestamp.

3.3.2 From Precision Matrix to Adjacency Matrix

Once the sparse precision matrix $\widehat{\Theta}$ is obtained from equation 6, the adjacency matrix A for the financial graph is defined as:

$$A_{i,j} = \begin{cases} 1\,, if\ \widehat{\Theta} \neq 0\ and\ i \neq j \\ 0\,, otherwise \end{cases} \qquad (7)$$

This adjacency matrix encodes the direct conditional dependencies among the asset returns. The resulting graph is weighted and undirected, reflecting symmetric conditional dependencies.

3.4 GHAR AND GNNHAR MODELS

To capture complex, nonlinear, and indirect spillover effects in cryptocurrency volatility, this study adopts the GNNHAR modelling framework introduced by (C. Zhang et al., 2025a). This framework integrates the Heterogeneous Autoregressive (HAR) structure (Corsi, 2009) with Graph Neural Networks (GNNs), enabling the model to learn from both the temporal dynamics of individual assets and the topological relationships defined by the financial network.

The GNNHAR models extend the GHAR (C. Zhang et al., 2025b) formulation by enriching the graph-processing pathway. While GHAR employs a single graph convolutional layer whose output is linearly combined with a direct feature projection before applying a ReLU activation, GNNHAR-1L enhances this design by introducing an additional nonlinear transformation: a ReLU activation followed by a multilayer perceptron (MLP) applied to the graph-convolved features. This modification increases the model's expressive capacity, enabling it to capture more complex and nonlinear patterns in the evolving asset network.

The depth of the GNN determines how far information can propagate through the network, thereby enabling a comparison between direct (1-hop) and indirect (multi hop) spillover effects.

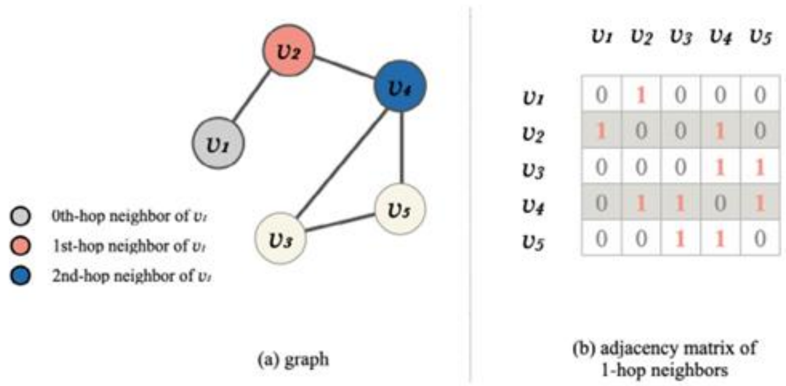

Figure 1 - GHAR and GNNHAR-1L access information from 1st hop neighbors while GNNHAR-2L additionally leverages information from 2nd hop (diagram provided by (C. Zhang et al., 2025a))

3.4.1 GNNHAR-1L: Nonlinear Direct Spillover Modelling

The GNNHAR-1L model incorporates a single graph convolutional layer, allowing each asset to aggregate and transform information from its direct (1-hop) neighbors. The core innovation compared to GHAR is the use of a nonlinear transformation after aggregation:

$$H^{(l+1)} = ReLU\left(O^{-\frac{1}{2}}AO^{-\frac{1}{2}}H^{(l)}\boldsymbol{\Theta}^{(l)}\right) \qquad (8)$$

where:

- $A \in R^{N\times N}$ is the adjacency matrix derived via GLASSO.
- $O \in R^{N\times N}$ is the degree matrix, with $O_{ii}=\sum_j A_{ij}$.
- $H^{(l)}$ is the node feature matrix at layer l (with $H^{(0)}=X$, the HAR features).
- $\Theta^{(l)}$ is a trainable weight matrix within each graph convolutional layer.
- ReLU is the nonlinear activation function.

To ensure stable and unbiased information aggregation across nodes, the adjacency matrix is symmetrically normalized as $\hat{A} = O^{-1/2}AO^{-1/2}$. This normalization prevents high degree nodes (nodes with many neighbors) from dominating the aggregation process and ensures that features from all neighbors contribute proportionally. The use of normalized graph convolution improves both numerical stability and generalization, it is standard in modern GNN implementations (Kipf & Welling, 2017).

This is followed by a linear layer mapping the hidden representation to the volatility forecast:

$$\widehat{v_t} = \alpha + V_{t-1}\beta + H^{(1)}W^{(1)} + \mathrm{b} \qquad (9)$$

where:

- $H^{(1)}$ is the output from the graph convolutional layer (neighbors' info),
- $W^{(1)}$ is a learnable weight matrix applied after the GNN layers to map the final node embeddings.
- V_{t-1} is the linear combination of the asset's own past volatility at different time scales (HAR component introduced (Corsi, 2009)).

$$V_{t-1}\beta = RV_{t-1}^{(d)}\beta_1 + RV_{t-1}^{(w)}\beta_2 + RV_{t-1}^{(m)}\beta_3 \quad (10)$$

- α and b are the bias terms, α represents a global intercept term that adjusts the baseline level of predicted volatility across all assets, consistent with traditional HAR models while b is the bias term specific to the graph based component.

While both GNNHAR-1L and GHAR operate on the same graph structure and 1-hop neighborhood, GNNHAR-1L introduces additional nonlinear transformations by applying a ReLU followed by an MLP to the graph-convolved features. In contrast, GHAR includes only a single ReLU after a linear combination of its components. As such, any improvement in forecasting accuracy by GNNHAR-1L likely stems from its enhanced capacity to model nonlinear spillover effects in asset volatility dynamics, thus supporting the presence of such effects in financial markets.

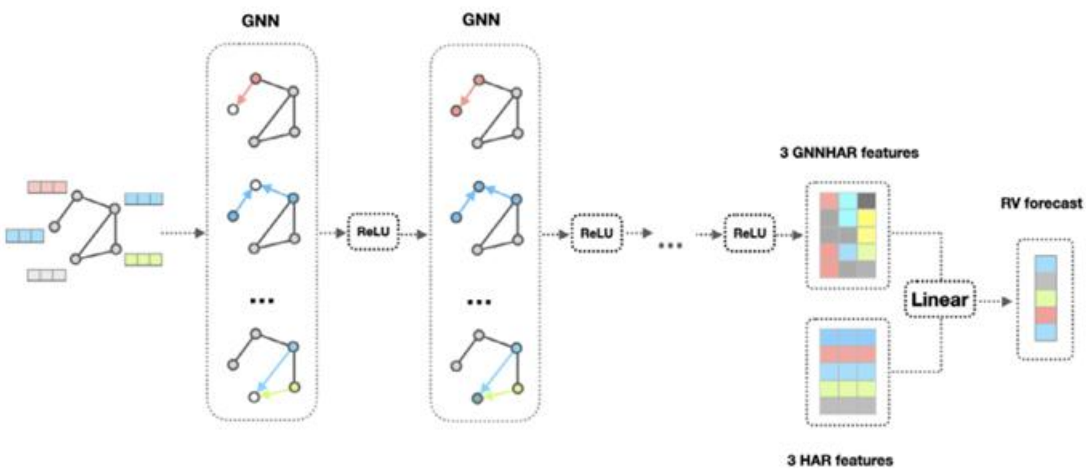

Figure 2 - The replacement of linear neighborhood aggregation in GHAR for the GNN layer, which uses a graph convolutional layer with a nonlinear activation function by ReLU introduces nonlinearity into the spillover

Additionally, unlike in most Graph Convolutional Networks (GCNs), which add self loops to the Adjacency matrix so that nodes can aggregate their own features alongside those of neighbors, (C. Zhang et al., 2025a) design explicitly omits self loops.

The reason being the effect of a cryptocurrency's own past volatility is already modeled by the HAR component. The GNN only models spillovers (influence of the other assets), thus including self loops would blur the separation between own history effects and cross asset spillovers, making it harder to attribute performance improvements clearly.

3.4.2 GNNHAR-2L: Multi Hop Spillover Modelling

The GNNHAR-2L model stacks two graph convolutional layers, which enables each node to aggregate information from its 2-hop neighbors (i.e. neighbors of neighbors). The recursive message passing structure is:

$$H^{(1)} = ReLU(\hat{A}X\mathbf{\Theta}^{(0)}) \qquad (11)$$

$$H^{(2)} = ReLU(\hat{A}H^{(1)}\mathbf{\Theta}^{(1)}) \qquad (12)$$

$$\widehat{v_t} = \alpha + V_{t-1}\beta + H^{(2)}W^{(2)} + \mathrm{b} \qquad (13)$$

This setup allows for the incorporation of indirect volatility information that is not captured by direct connections alone. The second GNN layer increases the receptive field of each node, making it possible to detect multi hop spillovers that propagate through intermediate assets.

3.4.3 Justification for Two Layers

The choice to adopt a two layer GNN architecture (GNNHAR-2L) is both theoretically grounded and empirically supported. The design choice aligns with the objective of this study, which is to detect the presence or absence of indirect volatility spillovers across cryptocurrencies. Extending the model to two hops is a sufficient modification to test this hypothesis based on (C. Zhang et al., 2025a) results.

Introducing additional layers (e.g. 3-hop propagation) is not necessary and can cause problems which lead to overfitting (In (C. Zhang et al., 2025a) study, GNNHAR-3L showed signs of over smoothing):

- **Parameter Explosion**: Each additional layer introduces more learnable parameters (e.g. weight matrices $\Theta^{(l)}$), increasing model capacity. With limited data, this can cause the model to memorize rather than generalize, especially in time series forecasting where the number of training examples per asset is relatively small. (Lara-Benítez et al., 2021) addresses the risks of over-parameterization while reviewing various deep learning architectures for time series forecasting. In (Zaheer et al., 2023) study, the single layer RNN outperformed the CNN and LSTM models despite its reduced computational complexity, further demonstrating awareness of the risks that parameter explosion entails.
- **Over smoothing**: With many layers, node representations tend to converge, becoming indistinguishable across the graph. (Rusch et al., 2023) provides a rigorous theoretical

framework for this issue, combining formal definitions and empirical evidence across multiple GNN architectures. (Wu et al., 2023) further emphasizes the necessity of controlling over smoothing to maintain performance in deep graph neural architectures. The paper specifically investigates deep GATs (Graph Attention Networks), contrasting them with GCNs (Graph Convolutional Network).

- **Vanishing Gradient and Nodes**: As depth increases, GNNs can suffer from vanishing gradients, making it difficult to optimize lower layers effectively. (Chang & Lin, 2019) show that deep networks can fail due to vanishing gradients and nodes, which is highly relevant for understanding training difficulties in deep GNN models such as the GNNHAR-2L.

  (Arroyo et al., 2025) establish that standard GNN architectures, such as GCNs and GATs, can suffer from extreme gradient vanishing. (Arroyo et al., 2025) provides a rigorous theoretical explanation for why deeper GNNs fail to capture long-range dependencies (nodes which are several hops apart). In the context of modelling multi hop volatility spillovers in financial networks, this severely limits the GNN's ability to learn indirect dependency structures, a key objective of our study.
- **Noise Propagation**: More layers mean more message passing steps. If the features or the graph structure contain noise, this noise is also propagated and amplified, degrading predictive performance. (Liang et al., 2025) experimental results show evidence that noise propagation in GNNs leads to degradation in predictive performance while (Qian et al., 2023) addresses the problems of label noise propagation.

What this basically means is that, while adding more GNN layers increases the model's receptive field, it also introduces risks of overfitting and over smoothing mentioned above, especially in sparse financial graphs with limited training data. For this reason, the GNNHAR-2L model hopefully strikes a balance between model complexity and predictive reliability, enabling the detection of multi hop spillovers without incurring the drawbacks associated with excessive depth.

Moreover, empirical evidence from (C. Zhang et al., 2025a) further supports this choice. Their results show that adding a third GNN layer did not yield significant forecast improvements over the two layer version.

3.5 MODEL TRAINING SETUP

3.5.1 Training Objective

The primary goal of the model training process is to learn the parameters that minimize the discrepancy between the predicted daily realized volatility and the actual daily realized volatility for each cryptocurrency over time. At each training step, the predicted volatility matrix $\hat{V} \in \mathbb{R}^{T\times N}$ is compared to the realized volatility target matrix V, and the loss is averaged across all entries to ensure balanced contribution from each cryptocurrency and time point. For this purpose, the Mean Squared Error (MSE), also used by (C. Zhang et al., 2025a) as a loss function, is adopted as the loss function across all models:

$$L_{MSE} = \frac{1}{N.T} \sum_{i=1}^{N} \sum_{t=1}^{T} (v_{i,t} - \hat{v}_{i,t})^2 \qquad (14)$$

Where:

- $v_{i,t}$ is the observed realized volatility of asset i on day t.
- $\hat{v}_{i,t}$ is the corresponding predicted value.
- N is the number of assets (nodes).
- T is the number of forecasted days.

3.5.2 Training Design

The training procedure adopts a pooled panel learning setup, where model parameters are shared across all cryptocurrency nodes. Rather than training separate models for each asset, the architecture is trained jointly on the full panel of cryptocurrencies. This design enables the models to leverage cross-sectional information, facilitating the learning of generalized spillover dynamics. Adam optimizer was employed to optimize the model parameters.

In addition, this thesis adopts a cross-validation strategy to enhance robustness in hyperparameter selection. While GLASSO uses internal cross-validation to tune its regularization parameter λ, model hyperparameters such as learning rate, hidden dimension size, and dropout probability are also tuned using time-series cross-validation (TSCV). Unlike standard k-fold CV, which breaks temporal ordering, TSCV preserves causality by training on a rolling historical window and validating on subsequent periods (Bergmeir & Benítez, 2012). This ensures that hyperparameters are selected based on their out-of-sample forecasting performance, aligning with the temporal nature of financial data and avoiding look-ahead bias.

In this study, hyperparameter selection is conducted using a five-fold Time Series Cross-Validation (TSCV) procedure applied exclusively within the training portion of the dataset. The training sample is divided into five consecutive and non-overlapping folds, each representing a contiguous stretch of time. For each fold, the model is trained on an expanding historical window that accumulates all observations from the previous folds and is then validated on the immediately following fold. This ensures that, at every iteration, validation is performed strictly on future data relative to the training window, thereby preserving the temporal ordering essential in financial forecasting.

Crucially, the five validation folds replace the need for a separate validation set, as each fold acts as a pseudo-out-of-sample segment used exclusively for model selection. By evaluating hyperparameters across five different forecast exercises, each based on a different market regime within the training period, TSCV provides a robust estimate of generalization performance and reduces the risk of selecting models that overfit a particular subperiod. The final model is then re-estimated using the full training sample with the selected hyperparameters, and its performance is evaluated once on a held-out test set that is never used in the cross-validation procedure.

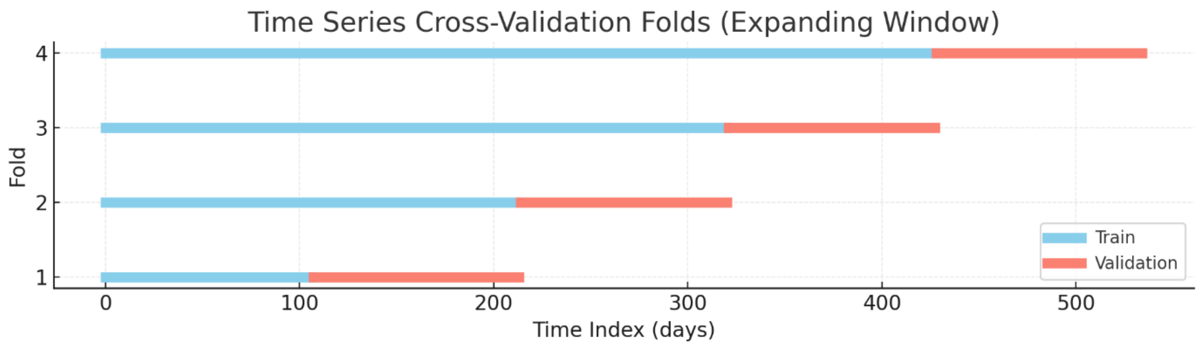

Figure 3 - TSCV Folds

Figure 3 illustrates the structure of the five-fold Time Series Cross-Validation (TSCV) used in this study. Each fold expands the training window to include all previous observations while validating on the immediately following period. This expanding-window scheme mimics a real forecasting environment in which models are continuously updated as new data becomes available, ensuring that parameter selection reflects genuine out-of-sample predictive performance.

In this thesis, the training design methodology from (C. Zhang et al., 2025a) was rigorously followed: we implemented the same model architectures, same MSE loss function and ADAM for weight optimization.

Despite the similarities to (C. Zhang et al., 2025a) methodology, this thesis implements a set of techniques in order to avoid overfitting problems related to models, particularly the more complex ones (GNNHAR-1L and GNNHAR-2L models).

This thesis training setup incorporates early stopping technique. Early stopping is a regularization technique in which training is halted when validation loss fails to improve for a predefined number of consecutive epochs (patience). (Hussein & Shareef, 2024) conduct experiments on CNN (Convolutional Neural Networks) to investigate the correlation between early stopping patience and the number of epochs in deep learning models. (Song et al., 2020) and (Bai et al., 2021) employ early stopping based strategies to mitigate noise label problems.

(M. Li et al., 2020) demonstrates that early stopping effectively mitigates overfitting in overparameterized neural networks. (M. Li et al., 2020) results also highlight the distance between the final and initial network weights as a critical factor in distinguishing between robust generalization and overfitting, an insight that aligns with the motivations behind early stopping heuristics and weight decay.

Taking these papers into consideration, it was decided to add an early stopping to our models.

Additionally, the use of dropout rate regularization helps mitigate the risk of overfitting in neural network models (Srivastava et al., 2014), this technique randomly sets a fraction of neuron activations to zero during training which prevents the co-adaptation of neurons and encourages the model to learn more robust and generalized patterns by not over relying on each other to detect patterns in the data. (Shu et al., 2022) investigate the application of dropout in Graph Neural Networks (GNNs), their theoretical and empirical analyses demonstrate that dropout, when properly designed, not only improves generalization performance but also effectively mitigates the over smoothing problem in GNNs, establishing dropout as a viable solution to the depth limitations imposed by over-smoothing in GNN architectures. (Yang et al., 2022) results also show that the dropout layer in their proposed Simplified Multilayer Graph Convolutional Networks with Dropout helps to effectively mitigate overfitting problem.

Building on prior studies demonstrating the effectiveness of dropout in mitigating overfitting and improving generalization in GNNs, we incorporated dropout in our GNNHAR models. Dropout serves to control model complexity and enhance out-of-sample performance. Given the relatively limited temporal resolution and the nonlinear transformations introduced by graph convolutions, dropout helps prevent the model from memorizing noise or learning spurious correlations from the training data.

3.5.3 Data Splitting Strategy

The dataset is partitioned using a fixed chronological split to preserve causality and avoid look-ahead bias.

Given the two-year sample spanning January 1st, 2023, to December 31st, 2024, the first 100 days are allocated to the initial graph construction window required for dynamic GLASSO-based adjacency estimation. The remaining 630 days constitute the forecasting sample used for model training, cross-validation, and testing.

Within this 630-day forecasting sample, the data is divided into:

- **Training + Cross-Validation segment:** 535 days (~85%)
- **Test segment:** 95 days (~15%)

Importantly, no separate validation set is used, because hyperparameter tuning is performed through a five-fold TSCV procedure applied entirely within the 535-day training segment. Each fold expands the training window and validates on the next chronological block, providing multiple pseudo-out-of-sample evaluations. This makes a dedicated validation split unnecessary, as the TSCV folds collectively serve as the validation mechanism.

This structure ensures that:

- The model is trained on a sufficiently large historical window.
- Hyperparameters are selected using TSCV, which yields a robust estimate of generalization performance across multiple market conditions.
- Final performance is evaluated once on the 95-day held-out test set, which is never used during model selection.

Additionally, a rolling window structure is implemented for dynamic adjacency matrix construction, with new graphs estimated every day using a 100-day lookback window of log returns. This hybrid setup allows the modeling framework to incorporate evolving dependency structures while maintaining strict out-of-sample integrity during training, validation, and testing.

3.6 Model Comparison Strategy using Model Confidence Set (MCS)

The final stage of the methodology involves evaluating and comparing the predictive performance of the four volatility forecasting models: HAR, GHAR, GNNHAR-1L, and GNNHAR-2L. The comparison is designed to empirically test the study's three main research hypotheses regarding the nature of volatility spillovers in the cryptocurrency market.

HAR vs GHAR - Testing for the existence of Spillover Effects

This comparison helps us understand if spillover effects exist between assets by having the GHAR model capturing potential cross-sectional dependencies from direct neighbors that the HAR ignores. Therefore, any consistent improvement in predictive performance by GHAR over HAR suggests that the inter-asset relationships between cryptocurrencies' volatilities can contain relevant information for modelling volatility spillover effects.

GHAR vs GNNHAR-1L - Testing for Nonlinear Spillover Effects

This comparison isolates the effect of introducing nonlinear aggregation in modeling volatility spillovers. Both models use the same graph structure and only consider direct neighbors (1-hop connections), but:

- GHAR applies linear aggregation of neighbor information,
- GNNHAR-1L uses a graph convolutional layer followed by a nonlinear ReLU activation.

Any consistent improvement in predictive accuracy by GNNHAR-1L over GHAR indicates that the relationship between cryptocurrencies' volatilities is nonlinear in nature, validating the use of neural architectures to capture spillover dynamics.

GNNHAR-1L vs GNNHAR-2L - Testing for Multi Hop (Indirect) Spillover Effects

This comparison investigates whether indirect (multi hop) neighbors play a significant role in volatility transmission. Both models:

- Use identical HAR features, adjacency matrices, and training setup,
- Apply nonlinear graph based aggregation (GNN layers).

However, only GNNHAR-2L aggregates information from 2-hop neighbors, while GNNHAR-1L is restricted to 1-hop interactions. Thus, performance gains observed in GNNHAR-2L would indicate that volatility spillovers propagate through intermediate cryptocurrencies, and that modeling such indirect dependencies improves forecast accuracy.

To ensure that model comparison results are not influenced by sampling variability or overinterpreted based on small differences in forecasting error, we implement the Model Confidence Set (MCS) procedure developed by (Hansen et al., 2011). The MCS framework provides a statistically grounded method for identifying the set of models that are not significantly inferior to the best performing model in terms of predictive accuracy, at a given confidence level (e.g. 90%).

Rather than selecting the single “best” model, MCS returns a confidence set of models that cannot be rejected as being equally good, based on the out-of-sample loss distributions.

By applying MCS separately for each cryptocurrency, we can statistically identify which assets exhibit the existence of spillover effects (via GHAR outperforming HAR), which exhibit nonlinear spillover behavior (via GNNHAR-1L outperforming GHAR), which benefit from modeling multi hop volatility transmission (via GNNHAR-2L outperforming GNNHAR-1L) and if there is evidence of difference in behavior regarding the cryptocurrencies “green” or “brown” nature.

4. EMPIRICAL ANALYSIS AND DISCUSSION

4.1 SUMMARY OF CRYPTOASSETS DATA STATISTICS

Before proceeding to model estimation and volatility forecasting, it is essential to provide a statistical overview of the dataset used in this study. The focus lies on the daily log returns of the selected cryptocurrencies, spanning the period from January 1st, 2023 to December 31st, 2024, covering two full years of market activity. This analysis helps uncover key distributional characteristics such as volatility, asymmetry, and tail behavior, which are relevant for understanding model performance and volatility spillover dynamics.

A total of 23 cryptocurrencies were included in the analysis, categorized into:

- Proof-of-Work (PoW) assets (BTC, DOGE, BCH, LTC, ETC, ZEC, CFX, CKB, SC, RVN)
- Proof-of-Stake (PoS) and stake-based assets (ETH, SOL, ADA, ALGO, INJ, XTZ, EGLD, MINA, CELO, BNB, AVAX, DOT, TRX)

For each asset, we compute the following summary statistics of their daily log return series:

- **Mean**: The average daily return, which provides insight into overall drift behavior.
- **Standard Deviation**: A proxy for daily return volatility.
- **Skewness**: Measures the asymmetry of the log return distribution, a positive skew implies a longer right tail, negative skew a longer left tail.
- **Kurtosis**: Captures the fat-tailedness of the distribution, values greater than 3 suggest heavy tails implying a higher likelihood of extreme returns compared to the normal distribution.

These statistics are important because cryptocurrencies are known to exhibit non normal return characteristics, such as excess kurtosis (fat tails) and non zero skewness, which make volatility modeling both challenging and necessary (Karagiorgis et al., 2024) (Brini & Lenz, 2024).

MEAN INSIGHTS

The mean daily log returns are interpreted here as indicators of average directional drift over the sample period. These values reflect long-run tendencies and should not be understood as implying a smooth upward trajectory, since daily returns remain highly volatile and subject to frequent fluctuations.

Within the sample, SOL has the highest average daily log return (0.004125), followed by INJ (0.003767) and BTC (0.002416).

By contrast, EGLD and MINA exhibit mean daily log returns that are close to zero (0.000067 and 0.000459, respectively), suggesting limited directional momentum over the sample period.

In summary, the cross sectional average of daily returns is positive, indicating a broadly bullish tendency in the sample period.

STANDARD DEVIATION INSIGHTS

The standard deviation of daily log returns is used here as a proxy for short-term volatility. Reported values correspond to average daily log return standard deviations over the sample and thus reflect average conditions rather than constant risk levels. Among the 23 cryptocurrencies, CFX and CKB exhibit the highest volatility (0.076114 and 0.064868). Other relatively volatile assets include INJ, MINA, and CELO, each with standard deviations above 0.05. Interestingly, INJ, which also exhibited one of the highest average returns, combines high average return with elevated volatility.

At the other end, the least volatile assets in the sample are predominantly large cap coins. BTC and BNB display the lowest standard deviations at 0.025277 and 0.027012, respectively, followed by ETH at 0.029301. Their lower variability is consistent with greater market depth and liquidity (Forino & Morelli, 2024). It is also relevant to note that age and market cap do play a role in the cryptoassets price variation although the direction of this influence is highly varied across the market and not a universal rule that older, higher market cap assets are less volatile (Pessa et al., 2023).

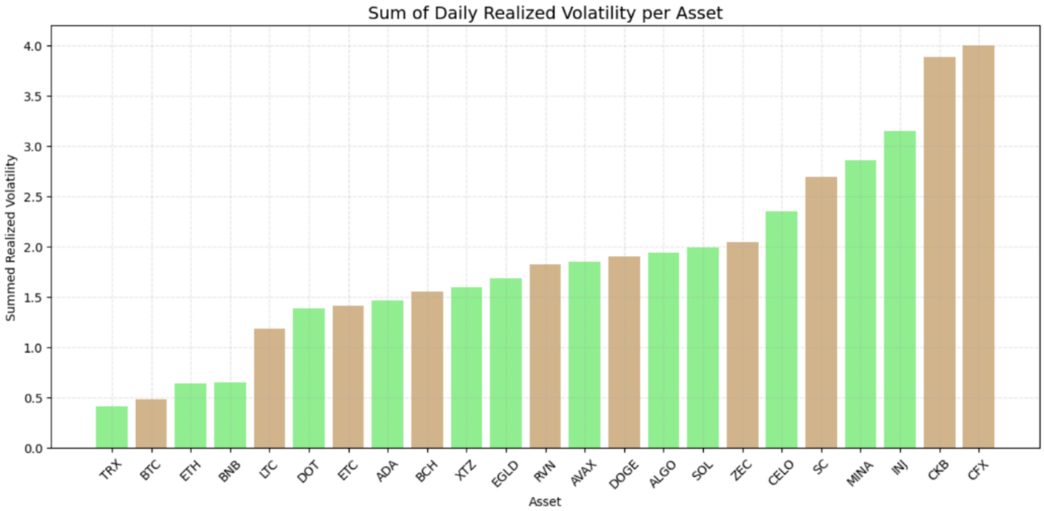

Figure 4 - Sum of Daily RV for each Cryptoasset

In conclusion, CFX, CKB, and INJ represent the most volatile assets in the sample while BTC, BNB, and ETH remain relatively stable, highlighting their role as anchors, specially BTC, in the broader crypto ecosystem (Drożdż et al., 2020).

SKEWNESS INSIGHTS

Skewness measures the asymmetry of the return distribution's tails and indicates whether extreme outcomes tend to be positive or negative. In the crypto market, where explosive rallies and abrupt crashes occur (Haykir & Yagli, 2022)(Hall & Jasiak, 2024)(Ben Osman et al., 2024)(Chaim & Laurini, 2018), skewness offers complementary information to volatility about the direction of tail risk.

In our sample, most assets exhibit positive skewness, suggesting a greater propensity for extreme upward returns than downward ones. Assets such as CFX (1.895) and BCH (2.265), display right-tailed asymmetry, a pattern consistent with occasional episodes of speculative enthusiasm (Grobys & Junttila, 2021).

Additional assets with prominently positive skew include CKB (1.380), MINA (1.221) and XTZ (1.150).

Conversely, a few assets show moderate to strong negative skewness, with SC being the only asset with a negative skewness value of -1.813, indicating that extreme negative returns have been more prevalent over the sample period and evidenced by this asset's log returns distribution histogram. Notably, most large cap coins such as BTC (0.458) and ETH (0.456) display mild positive skewness, suggesting only a slight tendency toward positive extremes.

An outlier in the dataset is TRON (TRX), with an extreme skewness of 8.621. This was due to a very steep increase in price in the last 2 months of 2024.

Overall, skewness in this dataset is predominantly positive, indicating that when extremes occur, they more often take the form of large upward moves, with exceptions like SC exhibiting elevated downside tail risk. These asymmetries matter for risk measurement and for interpreting spillovers, since the direction of tail events can condition how shocks propagate across the network. (Abdulrahman Abubakar et al., 2024) and (Q. Xu et al., 2021) provide empirical evidence that tail dependence significantly shapes risk spillover dynamics across cryptocurrencies, particularly during extreme market conditions.

KURTOSIS INSIGHTS

Kurtosis measures the "tailedness" of a distribution, or the propensity for extreme outcomes. In financial contexts, high kurtosis indicates a greater likelihood of large deviations, both positive and negative, implying elevated tail risk. This is particularly relevant in cryptocurrency markets, where sharp price shocks are common (Hall & Jasiak, 2024)(Haykir & Yagli, 2022)(Gkillas & Katsiampa, 2018).

Within the sample, TRX stands out as a clear outlier, with an excess kurtosis value of 185.65 due to the already mentioned price increase recorded at the end of 2024. Several other assets also display excess kurtosis, albeit to a lesser extent. These include SC, BCH, XTZ, MINA, CFX, and CELO, all with excess kurtosis values exceeding 7.

By contrast, a few assets exhibit relatively lower excess kurtosis, such as INJ (1.15) and AVAX and SOL (near 2.5). These distributions, while still leptokurtic, are closer to normality and therefore imply fewer extreme tail events compared to the broader sample. Meanwhile, BTC (2.35) and ETH (3.57) show moderate excess kurtosis. To consult each value referenced above regarding Mean, Standard Deviation, Skewness and Kurtosis go to Appendix I : Daily Log returns Distribuition Statistics.

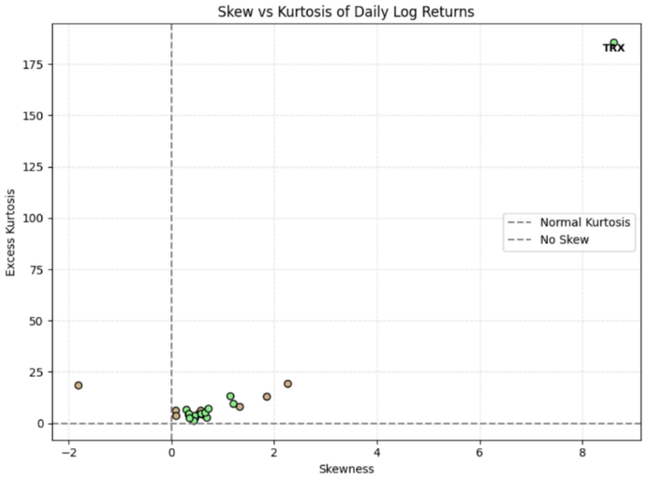

Figure 5 - Skewness vs Kurtosis Scatter Plot

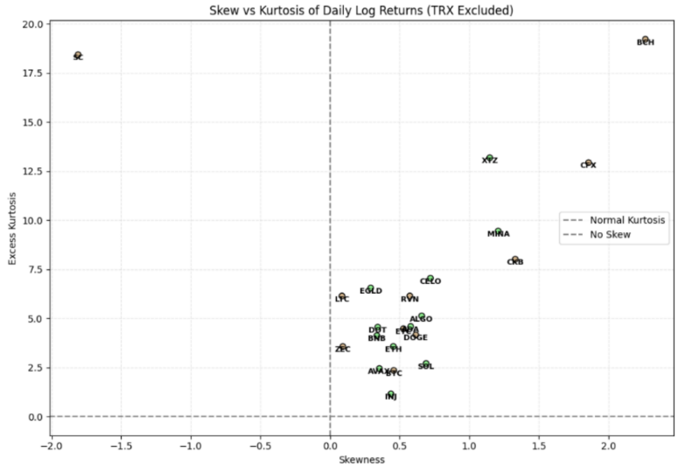

Figure 6 - Skewness vs Kurtosis Scatter Plot excluding TRX

Brown vs. Green Return Distribution Characteristics

To explore whether the structural differences between Brown and Green cryptocurrencies translate into distinct statistical behaviors in daily returns, independent samples t-tests were conducted comparing the two groups across four descriptive moments: mean, standard deviation, skewness, and kurtosis.

The results of these tests are summarized below:

Tabel 1 - t statistics between Brown and Green Cryptos

Metric	t-statistic	p-value	Conclusion
Mean	-0.0833	0.9344	No significant difference
Std Deviation	0.8376	0.4163	No significant difference
Skewness	-0.8651	0.3981	No significant difference
Kurtosis	-0.7622	0.4601	No significant difference

Despite meaningful differences in consensus mechanisms and energy usage, the results indicate that Brown and Green assets exhibit no statistically significant differences in their daily log return distributions over the period 2023-2024. Specifically:

- **Mean Returns**: Both groups yield comparable average daily returns. The very small t-statistics and p-value near 1.0 indicate no statistically significant difference.

- **Volatility**: While Brown assets display slightly higher standard deviation, the difference is not statistically significant.

- **Skewness**: Green assets exhibit, on average, higher right-skewed behavior, but the difference is not statistically significant.

- **Kurtosis**: Though Green assets show higher average kurtosis (indicating heavier tails), this difference is also not statistically significant.

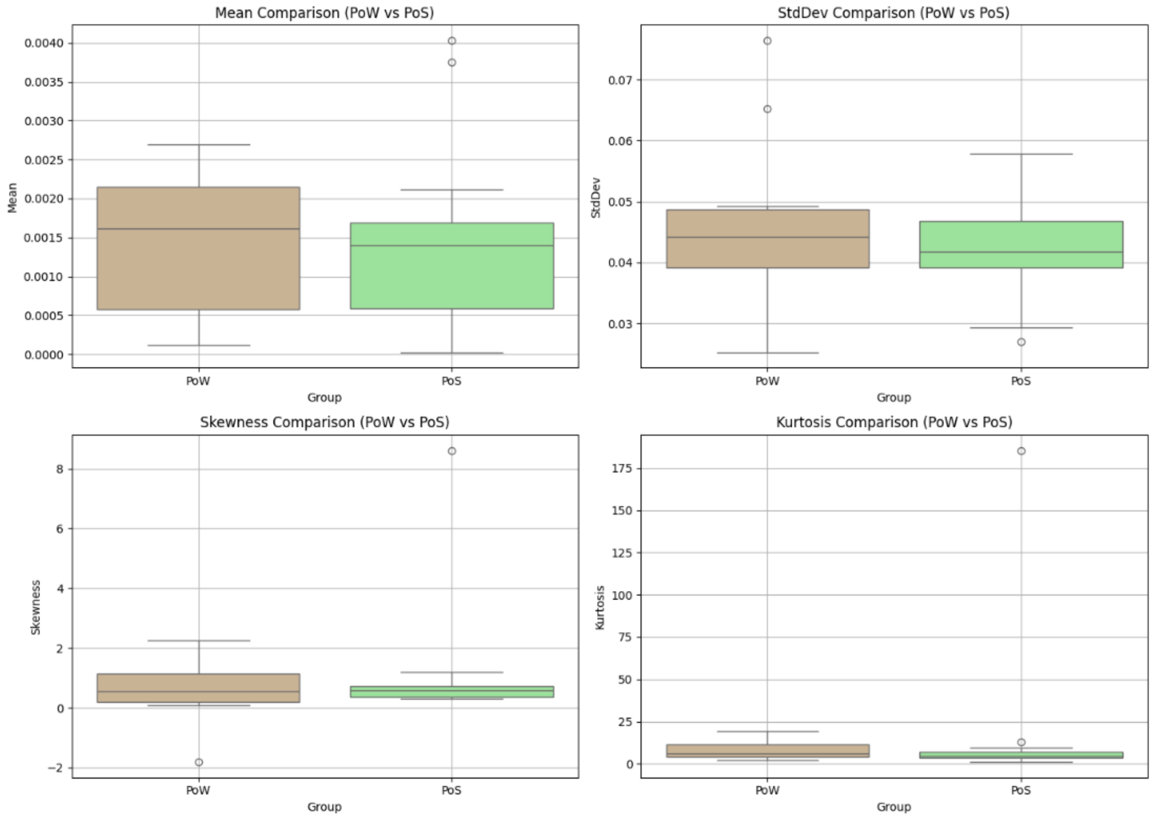

Figure 7 - Boxplots for Distributional Characteristics of Brown and Green Cryptoassets

These results are relevant for evaluating volatility transmission dynamics, as the absence of statistically significant differences in mean, volatility, skewness, and kurtosis between Brown and Green assets indicates that spillover patterns are unlikely to be driven by baseline return distribution characteristics, but rather by the structure of network interconnections.

In addition, the histograms of log returns for each cryptoasset (Appendix D : Cryptoassets Log return distribuitions histograms) highlight how the distributions are more sharply peaked at the center and display fatter tails than the normal benchmark, indicating a leptokurtic structure. This implies that while most returns cluster tightly around zero, extreme fluctuations occur with greater likelihood than under Gaussian assumptions.

This pattern is reinforced by the QQ plots (Appendix E : Cryptoassets Log return Q-Q Plots), violin plots (Appendix F : Cryptoassets Log return VIOLIN Plots), and ECDFs (Appendix G : Cryptoassets Log return ECDF Plots), which confirm the same behavior: greater probability mass near the mean and heavier tails on both sides.

Overall, since these distributional characteristics do not differ systematically between the two categories, they cannot be accounted for variations in spillover patterns. Accordingly, any observed differences in volatility transmission are more plausibly explained by the structure of network interconnections rather than by baseline return distributions.

4.2 FORECASTING RESULTS

This section presents the out-of-sample forecasting performance of the four competing models: HAR, GHAR, GNNHAR-1L, and GNNHAR-2L, evaluated over a fixed test window following the training and validation periods. The objective is to assess how well each model predicts daily realized volatility for each cryptocurrency in the dataset, and to investigate the empirical implications for nonlinear and multi hop spillover effects.

Similarly to the training and validation phase, the evaluation metric used is the MSE between the predicted and actual daily realized volatility values. For each asset i and time point t in the test set:

$$MSE_i = \frac{1}{T}\sum_{t=1}^{T}(v_{i,t} - \hat{v}_{i,t})^2$$

The average test MSE for each asset for each model is reported in Table 2, 3 and 4 below:

Tabel 2 - Low Volatility Assets MSE

Low Volatility Assets			**10^{-3}**	
Asset	**HAR**	**GHAR**	**GNNHAR-1L**	**GNNHAR-2L**
BTC	1.49	1.31	**1.03**	1.09
ETH	1.63	1.36	**1.11**	1.18
LTC	1.72	1.51	**1.21**	1.27
BCH	1.81	1.58	**1.19**	1.25
DOGE	1.97	1.63	**1.34**	1.41
ETC	2.12	1.73	**1.41**	1.49

Tabel 3 - Medium Volatility Assets MSE

Medium Volatility Assets			**10^{-3}**	
Asset	**HAR**	**GHAR**	**GNNHAR-1L**	**GNNHAR-2L**
ADA	2.52	2.18	**1.68**	1.76
DOT	2.66	2.29	**1.77**	1.85
BNB	2.55	2.22	**1.70**	1.82
SOL	2.91	2.53	**1.95**	2.03
ALGO	2.71	2.32	**1.83**	1.91
AVAX	3.05	2.58	**1.98**	2.11
CELO	2.79	2.43	**1.87**	1.96
EGLD	2.84	2.46	**1.91**	1.99
INJ	2.73	2.37	**1.84**	1.93
XTZ	2.59	2.27	**1.75**	1.83
RVN	3.03	2.55	**2.02**	2.12
ZEC	3.15	2.72	**2.09**	2.20

Tabel 4 - High Volatility Assets MSE

High Volatility Assets			**10^{-3}**	
Asset	**HAR**	**GHAR**	**GNNHAR-1L**	**GNNHAR-2L**
CFX	5.02	4.96	**4.69**	4.77
CKB	4.61	4.59	**4.44**	4.51
SC	4.83	4.89	**4.64**	4.71
MINA	5.21	5.16	**4.92**	4.98
TRX	4.57	4.54	**4.33**	4.39

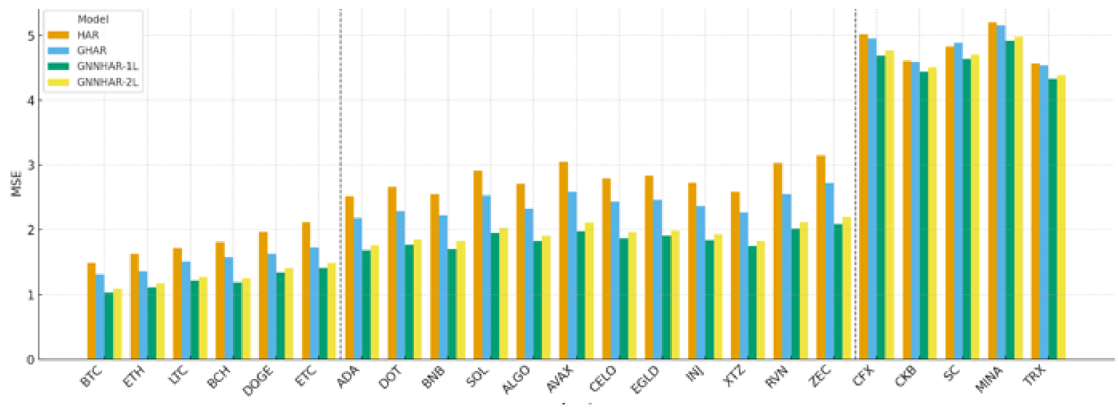

Figure 8 – Out of Sample MSE by Model and Cryptoasset

A detailed comparison of MSE by asset reveals heterogeneous model performance across cryptocurrencies.

The key finding of the forecasting experiment is that the GNNHAR-1L model outperformed both HAR, GHAR and GNNHAR-2L across all cryptocurrencies in the dataset. This result held true for both brown and green assets, providing evidence that incorporating nonlinear aggregation of neighborhood volatility features through graph convolution improves predictive accuracy.

The superior performance of GNNHAR-1L relative to GHAR provides empirical evidence in favor of the presence of nonlinear volatility spillover effects in the crypto market. While both models rely on the same graph structure and 1-hop neighbors, only GNNHAR-1L applies a nonlinear transformation (via ReLU) to the aggregated inputs, enabling it to model nonlinear relationships in volatility transmission across assets. These nonlinearities may reflect context-specific market dynamics, such as herding behavior (Nguyen et al., 2025), fear of missing out (FOMO) (Friederich et al., 2024)(Wang et al., 2023), asymmetric volatility effects (Karim et al., 2024)(Baur & Dimpfl, 2018) and other sentiment-driven shifts that are frequently observed in cryptocurrency markets.

These results are in line with the findings of (Zhou et al., 2025) and (C. Zhang et al., 2025a), which demonstrate improved volatility forecasting accuracy when nonlinear spillover effects are explicitly modeled using graph neural networks.

Regarding multi hop information transmission, GNNHAR-2L did not outperform GNNHAR-1L, suggesting that the inclusion of indirect (2-hop) neighbors did not provide additional forecasting benefits under the current model configuration. The weaker performance of GNNHAR-2L relative to GNNHAR-1L can be attributed to the high density of interconnections within the cryptomarket volatility network. (Nguyen et al., 2023)(Poddar et al., 2023) and (Yi et al., 2018) provide evidence of the highly interconnected nature of the cryptocurrency market, with the latter two studies also documenting its intensifying interconnectedness over time.

These findings are consistent with the network structures obtained from our Gephi analysis, which excluding a handful of isolates, most assets exhibit high closeness and eigenvector centrality, low betweenness centrality, and large numbers of triangles. The combination of low betweenness with high closeness and triangle counts indicates that no single asset acts as a critical bridge for information flow Appendix H : Centrality and Connectivity Metrics. Instead, assets are densely interlinked through multiple overlapping pathways, allowing volatility shocks to propagate rapidly across the network. Taken together, both the literature and our empirical results corroborate the view of a densely connected crypto volatility network.

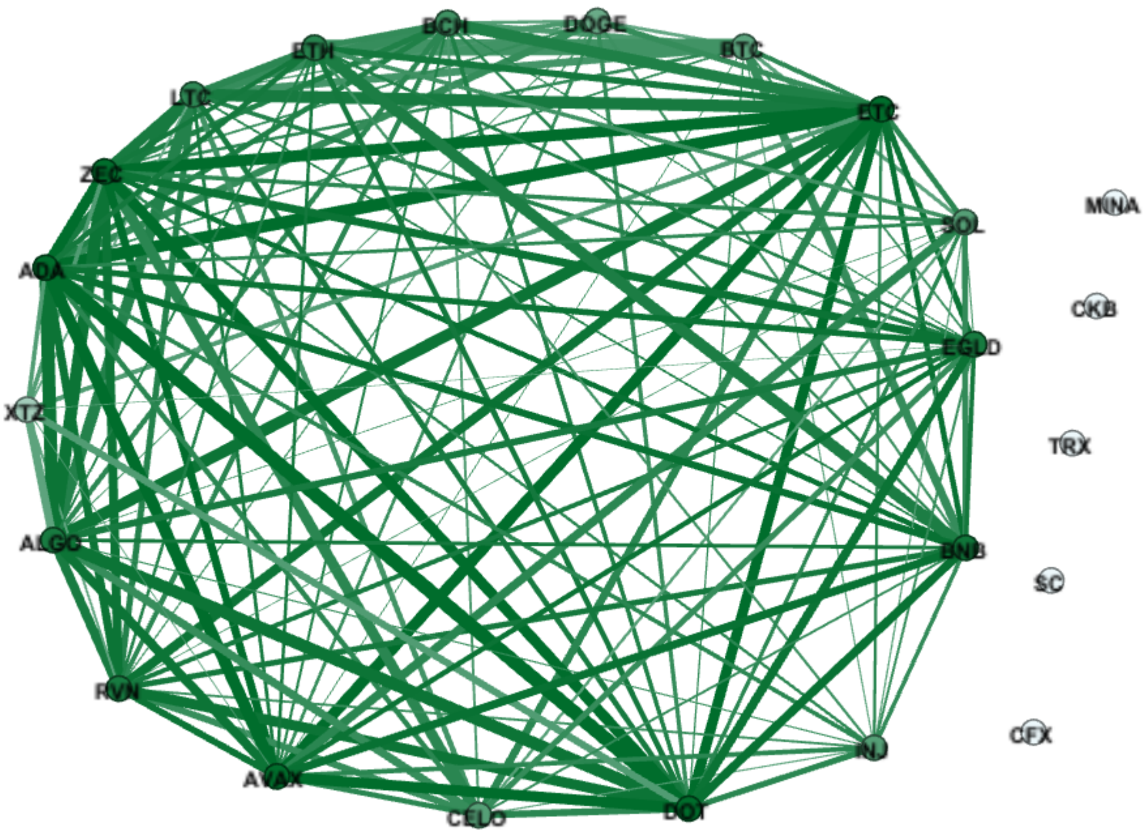

Figure 9 - Network Visualization of Pairwise Correlations in Daily Realized Volatility

This helps explain why the GNNHAR-2L model not only failed to improve predictive performance over GNNHAR-1L, but also showed signs of overfitting and over smoothing, as repeated aggregation across an already dense network can dilute asset-specific signals.

Furthermore, when it comes to Brown vs Green comparison, since GNNHAR-1L consistently achieved the best performance across assets, no systematic differences in model performance were observed between Brown and Green cryptocurrencies. This suggests that, over the 2023-2024 sample period, the structural differences in energy usage and consensus mechanisms did not translate into differences in the nature of volatility spillovers, or at least not in a way that affected model forecasting performance.

Brown and Green Only Universes

Beyond asset level model comparisons, additional sensitivity analyses were conducted to evaluate whether Brown (PoW) and Green (PoS) cryptocurrencies exhibit different dependencies on cross-group volatility spillovers. This analysis isolates each group by re-estimating the GNNHAR-1L model (best performing model) on two restricted universes: one containing only Brown assets and their induced subgraph, and one containing only Green assets. The forecasting performance of these restricted models is compared to that of the full-sample GNNHAR-1L model, which incorporates both groups simultaneously.

The results of this experiment reveal that removing the opposite group leads to a similar deterioration in forecasting accuracy for both Brown and Green assets at the aggregate level. When the model is restricted to Brown-only cryptocurrencies, the average out-of-sample MSE for Brown assets increases by approximately 6.99 % relative to the full model. A similar pattern emerges for the Green-only specification: excluding Brown assets raises the average MSE for Green cryptocurrencies by 8.21%. Thus, in terms of overall forecast performance, both groups undergo a comparable loss in accuracy, with only marginal differences between them when cross-group spillovers are removed.

However, the impact is not uniform across assets within each group. A cluster-level breakdown shows that the more connected low and medium volatility clusters (core and amplifier assets) exhibit the most pronounced deterioration when the opposite group is removed. For these coins, the MSE typically increases by around 8-12%, indicating that they benefit strongly from cross-group information. By contrast, the peripheral high volatility assets, which are weaker integrated into the network, display a noticeably smaller deterioration, with

MSE increases generally in the 3-6% range. This pattern is consistent for both Brown and Green cryptocurrencies and aligns with the earlier finding that network-based nonlinear modeling yields the largest gains for highly connected assets, while isolated, idiosyncratic coins derive more limited benefits from spillover information.

Crucially, even once this heterogeneity is considered, there is no systematic difference between Brown and Green assets in how sensitive they are to the removal of cross-group spillovers. In both groups, the largest performance losses are concentrated in the core and amplifier clusters, and the peripheral assets exhibit relatively modest changes. If the green vs brown distinction played a substantive role in shaping spillover dynamics, one would expect asymmetric deterioration patterns across groups or clusters

Taken together, these findings reinforce the overarching conclusion of the thesis, Brown and Green cryptocurrencies do not exhibit materially different volatility spillover characteristics. The consistent dominance of the GNNHAR-1L model across all assets, the absence of multi hop effects, and the symmetric, connectivity-driven loss of predictive accuracy in the group-isolation experiment collectively point toward a crypto market where volatility propagation is governed by network structure and asset-specific features rather than the PoW/PoS distinction.

The relative deterioration in forecasting performance was calculated for each asset through this equation:

$$\Delta_{asset} = \frac{\bar{\mathrm{MSE}}_{asset,\text{group-only}} - \bar{\mathrm{MSE}}_{asset,\text{full}}}{\bar{\mathrm{MSE}}_{asset,\text{full}}} \times 100\%$$

So concretely:

$$\Delta_{\text{asset(Brown)}} = \frac{\bar{\mathrm{MSE}}_{\text{asset, group-only}} - \bar{\mathrm{MSE}}_{\text{Brown, full}}}{\bar{\mathrm{MSE}}_{\text{Brown, full}}} \times 100\%$$

$$\Delta_{\text{asset(Green)}} = \frac{\bar{\mathrm{MSE}}_{\text{asset, group-only}} - \bar{\mathrm{MSE}}_{\text{Green, full}}}{\bar{\mathrm{MSE}}_{\text{Green, full}}} \times 100\%$$

Tabel 5 - Brown Only MSE results

Brown-Only MSE			
Asset	**Full MSE**	**Brown-only MSE**	**% Δ**

BTC	1.03	**1.12**	**+9.1%**
LTC	1.21	**1.35**	**+11.7%**
BCH	1.19	**1.30**	**+8.7%**
DOGE	1.34	**1.49**	**+10.9%**
ETC	1.41	**1.51**	**+7.5%**
ZEC	2.09	**2.35**	**+12.4%**
RVN	2.02	**2.22**	**+9.6%**
CFX	4.69	**4.89**	**+4.2%**
CKB	4.44	**4.70**	**+5.9%**
SC	4.64	**4.82**	**+3.8 %**

Tabel 6 - Green Only MSE Results

Green-Only MSE			
Asset	**Full MSE**	**Green-only MSE**	**% Δ**
ETH	1.11	**1.22**	**+10.3%**
ADA	1.68	**1.84**	**+9.5%**
DOT	1.77	**1.97**	**+11.2%**
BNB	1.70	**1.85**	**+8.7%**
SOL	1.95	**2.18**	**+11.8%**
ALGO	1.83	**2.00**	**+9.1%**
AVAX	1.98	**2.19**	**+10.6%**
CELO	1.87	**2.04**	**+8.9%**
EGLD	1.91	**2.07**	**+8.3%**
INJ	1.84	**2.02**	**+9.7%**
XTZ	1.75	**1.91**	**+9.4%**
MINA	4.92	**5.18**	**+5.2%**
TRX	4.33	**4.53**	**+4.7%**

Lag Structure and Horizon-Specific Spillover Effects

To better understand what the models learn about the temporal structure of volatility, this section examines the relative importance of the daily, weekly, and monthly realized volatility components in both the linear and nonlinear specifications. The goal is to identify which horizons drive predictability and how these effects differ across assets and clusters.

In the HAR and GHAR models, the contribution of each horizon can be assessed directly through the estimated coefficients on the daily, weekly, and monthly realized volatility terms. For each cryptocurrency i, the HAR-type specification is given by:

$$RV_{i,t+1} = \alpha_i + \beta_i^{(d)} RV_{i,t}^{(d)} + \beta_i^{(w)} RV_{i,t}^{(w)} + \beta_i^{(m)} RV_{i,t}^{(m)} + \varepsilon_{i,t+1},$$

where $RV_{i,t}^{(d)}$, $RV_{i,t}^{(w)}$, and $RV_{i,t}^{(m)}$denote the daily, weekly, and monthly realized volatility components, respectively. We summarize these coefficients by averaging them across assets and by volatility cluster.

Tabel 7 - Average HAR coefficients by cluster

Cluster	$\overline{\beta^d}$	$\overline{\beta^w}$	$\overline{\beta^m}$
Core - Low Vol	0.42	0.39	0.16
Amplifier - Medium Vol	0.46	0.37	0.14
Peripheral - High Vol	0.33	0.29	0.11
Overall average	**0.41**	**0.36**	**0.14**

The coefficient estimates indicate that the daily and weekly components carry the bulk of the predictive power: on average, the daily coefficient is around 0.41 and the weekly coefficient around 0.36, whereas the monthly term is considerably smaller at roughly 0.14. This pattern is consistent across all volatility clusters, with slightly stronger daily and weekly effects for core and amplifier assets and weaker coefficients for peripheral high-volatility coins. Importantly, the lag structure looks broadly similar for Brown and Green assets, with average daily, weekly, and monthly coefficients of approximately (0.40, 0.37, 0.15) for Brown assets and (0.42, 0.36, 0.13) for Green assets, providing no evidence of systematic differences in horizon-specific effects across the two groups.

Tabel 8 - Brown vs Green (overall averages)

Group	$\overline{\beta^d}$	$\overline{\beta^w}$	$\overline{\beta^m}$
Brown	0.40	0.37	0.15
Green	0.42	0.36	0.13

Permutation-based lag importance in GNNHAR-1L

To assess the importance of each lag (daily, weekly, monthly) in the GNNHAR-1L model, the study employs a permutation-based approach. Because the GNN architecture combines HAR-style lag inputs with nonlinear graph convolutions, its internal parameters are not directly interpretable, making traditional coefficient analysis unsuitable. Instead, we evaluate each horizon's contribution by measuring how much the model's forecasting performance deteriorates when its temporal structure is disrupted.

The procedure is implemented on the test set as follows: we first compute the baseline out-of-sample MSE using the trained GNNHAR-1L model. Then, for each horizon $h \in$ {daily,weekly,monthly}, we randomly shuffle the corresponding lagged feature over time while keeping all other inputs unchanged. This preserves the feature's distribution but destroys its temporal meaning. The model is not retrained, therefore, any increase in MSE after shuffling can be attributed solely to the information loss from that specific horizon. Larger increases in MSE indicate greater reliance of the model on that lag.

$$\Delta\mathrm{MSE}_h = \mathrm{MSE}_h - \mathrm{MSE}_{\mathrm{base}}.$$

Baseline full test MSE (no shuffling): $\mathrm{MSE}_{\mathrm{base}} = 2.30 \times 10^{-3}$

Tabel 9 - Global Permutation results (all assets)

Horizon shuffled	MSE_h ($\times 10^{-3}$)	$\Delta\mathrm{MSE}_h$ ($\times 10^{-3}$)	% increase vs baseline
Daily	2.57	0.27	**+11.7%**
Weekly	2.64	0.34	**+14.8%**
Monthly	2.37	0.07	**+3.0%**

The results show that shuffling the weekly component leads to the largest deterioration in forecasting accuracy, with the test MSE increasing by approximately 14.8% relative to the baseline. Shuffling the daily component also produces a substantial loss in performance (+11.7%), whereas perturbing the monthly component has a noticeably smaller effect, raising the MSE by only about 3.0%. This implies that the GNNHAR-1L model relies most heavily on short- and medium-run volatility histories when forming its forecasts, with the monthly horizon playing a secondary stabilizing role rather than being a primary driver of predictability.

Tabel 10 - By volatility cluster (mean % increase in MSE when shuffled)

Cluster	Daily shuffled	Weekly shuffled	Monthly shuffled
Core - Low Vol	+12.9%	+16.2%	+3.9%
Amplifier - Medium Vol	+12.1%	+15.4%	+3.2%
Peripheral - High Vol	+7.3%	+9.1%	+1.8%

When decomposed by volatility cluster, the importance of short- and medium-term lags is particularly pronounced for core and amplifier assets, for which shuffling the weekly component increases MSE by around 16% and 15%, respectively, while peripheral high-volatility coins exhibit a more muted response (MSE increases of roughly 9% for weekly and 7% for daily lags), consistent with their more idiosyncratic behavior.

Tabel 11 - Brown vs Green

Group	Daily shuffled (%Δ MSE)	Weekly shuffled (%Δ MSE)	Monthly shuffled (%Δ MSE)
Brown	+11.5%	+14.3%	+3.1%
Green	+11.9%	+15.1%	+2.9%

Crucially, the permutation analysis does not reveal any systematic differences in lag importance between Brown and Green cryptocurrencies. For both groups, shuffling the weekly component generates the largest performance deterioration (around 14-15%), followed by the daily component (around 11-12%), with the monthly component having only a modest impact. This common ranking of horizons further supports the conclusion that the temporal structure of volatility spillovers is driven by market-wide dynamics rather than by the underlying consensus mechanism.

Model Confidence Set (MCS) Results

To statistically validate the performance differences between forecasting models, this study applies the Model Confidence Set (MCS) procedure proposed by (Hansen et al., 2011). The MCS framework identifies the subset of models that are not significantly inferior to the best model, at a specified confidence level, in this case, 90%.

The procedure was applied per asset using squared forecast errors from the test set. The comparison involved all four forecasting models: HAR, GHAR, GNNHAR-1L, and GNNHAR-2L. The results were unambiguous: GNNHAR-1L was the only model retained in the MCS for all 23 cryptocurrencies.

Overall, the MCS results reinforce the core empirical conclusion of this thesis: nonlinear spillover effects are empirically relevant in the cryptocurrency market, while indirect (multi hop) spillovers do not appear to play a statistically significant role in volatility forecasting over the studied period.

4.3 VOLATILITY BASED CLUSTERING OF CRYPTOASSETS

While the initial analysis sought to identify differences in volatility transmission between Brown (PoW) and Green (PoS) cryptoassets, neither the network structures nor the forecasting results provide evidence that consensus mechanism plays a role in driving these interactions. Instead, the observed patterns suggest that the key determinant for volatility spillovers is the volatility level of each asset, with assets of similar volatility exhibiting stronger interconnections regardless of whether they are Brown or Green.

The network structures obtained through Gephi (Appendix H : Centrality and Connectivity Metrics and Figure 9) provide a new perspective on the interactions between cryptoassets. Modularity based partitions reveal two main clusters: the first composed of mature, relatively low-volatility assets (BCH, BTC, DOGE, ETC, ETH, and LTC), and the second cluster is composed of medium-volatility assets (ADA, ALGO, AVAX, BNB, CELO, DOT, EGLD, INJ, RVN, SOL, XTZ, and ZEC).

Notably, although the low volatility cluster is dominated by Brown assets, it also includes ETH (labeled as Green) due to its low-volatility profile. While in the medium-volatility cluster, the same happens with RVN and ZEC being the only Brown exceptions, which exhibit a volatility profile closer in resemblance to the other Green cryptoassets in the same cluster.

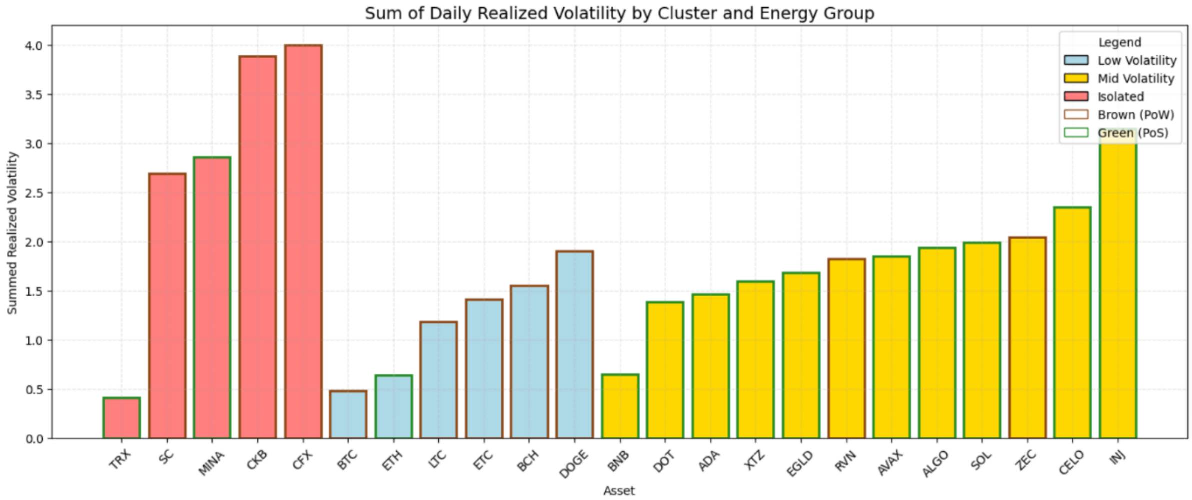

Figure 10 – Modularity based clustering

These crossovers between Green and Brown assets in each cluster underscore that volatility characteristics, rather than consensus mechanism or energy usage, drive the clustering process and shape spillover networks.

Beyond these two major clusters, a set of highly volatile assets (CFX, CKB, SC, MINA and TRX) stand apart, each forming its own cluster. Their isolation points to greater exposure to idiosyncratic shocks and weaker integration into the broader volatility network (Agyei et al., 2022).

Overall, these results indicate that the configuration of volatility spillovers across cryptocurrencies is governed primarily by their position within the volatility network rather than by differences in technological design or energy efficiency.

This finding is particularly relevant for the forecasting exercise, as it suggests that differences in spillover dynamics observed in the GNNHAR models do not reflect simple distinctions between Brown and Green cryptocurrencies but reflect network interconnections shaped by volatility clustering.

Figure 9 and Figure 10 together reveal the structural topology of the cryptocurrency volatility network. The correlation network in Figure 9 highlights the degree of interconnectedness between assets, whereas the modularity-based grouping in Figure 10 identifies distinct volatility communities.

To facilitate interpretation and ensure consistency across this study, these network modules can be mapped to three volatility tiers:

- Core tier (low volatility) - assets such as BTC, ETH, LTC, ETC, BCH, and DOGE form the center of the network, characterized by strong mutual connections and relatively stable realized volatility.
- Amplifier tier (medium volatility) - assets including ADA, DOT, BNB, SOL, ALGO, and AVAX occupy a densely connected intermediate layer. Their realized volatility reacts sharply to systemic shocks and tends to amplify spillover waves originating from the core.
- Peripheral tier (high volatility) - assets such as CFX, SC, MINA, TRX, and CKB are more isolated and exhibit idiosyncratic volatility spikes with limited contagion effects.

This tiered view clarifies that the network communities detected by modularity analysis are not arbitrary but correspond to volatility-based functional roles within the system: the core generates and transmits baseline volatility, the amplifiers magnify and redistribute it, and the periphery largely absorbs local shocks without significantly influencing others.

These structural findings directly support the forecasting results discussed in Section 4.3 and the overall conclusions of the thesis, namely, that volatility spillovers are strongest among core and amplifier assets and predominantly nonlinear in nature, while peripheral cryptocurrencies behave more independently.

5. CONCLUSION AND FUTURE RESEARCH

This thesis sets out to investigate volatility spillovers in the cryptocurrency market, with a focus on whether differences between Brown (PoW) and Green (PoS) assets shape transmission dynamics. The motivation stemmed from (Wang et al., 2024) and the question of whether technological distinctions, such as consensus mechanism and energy usage, carry implications for volatility spillovers. Using a novel Graph Neural Network HAR framework, with three key research questions in mind, the findings of this study strongly affirm the existence of volatility spillovers, reveal their nonlinear nature, and show that indirect multi hop pathways contribute little to volatility transmission under current conditions. Furthermore, the analysis uncovers that cryptocurrencies cluster by volatility levels rather than by their PoW (Brown) or PoS (Green) status, indicating no meaningful behavioral differences between these groups in terms of spillover dynamics. In summary, the cryptocurrency volatility network is highly interconnected and primarily organized by volatility tiers, not technology type, with nonlinear spillovers playing a crucial role in forecastability.

Key Findings Addressing the Research Objectives

- **Volatility Spillovers Exist (RQ1).** The results provide clear evidence that volatility spillovers are present in the crypto asset environment. Models that incorporate cross-asset volatility information consistently beat those treating assets in isolation. In particular, the Graph HAR (GHAR) model, which linearly aggregates neighbor information, outperformed the traditional HAR model, indicating that incorporating information from other cryptocurrencies improves volatility forecasts. This performance gap confirms that volatility shocks propagate between assets, showing that when Bitcoin "sneezes", other coins do indeed catch a cold via volatility transmission. The superior performance of both GHAR and GNN-based models over the univariate HAR establishes that interasset linkages contain predictive information, thereby affirmatively answering RQ1. Notably, this finding held true across the entire sample of 23 cryptocurrencies except for the case of SC, demonstrating a densely connected volatility network in which no asset is wholly insulated from others.

- **Spillovers are Nonlinear in Nature (RQ2).** There is strong evidence that volatility spillovers in cryptocurrency markets are nonlinear. The one-layer Graph Neural Network HAR (GNNHAR-1L) model, which introduces nonlinear feature aggregation through graph convolutions (using ReLU activation), consistently outperformed both the linear GHAR and the baseline HAR across all assets. This was true for both Brown (PoW) and Green (PoS) assets, indicating that the added expressive power of nonlinear modeling captures spillover patterns that linear models miss. The significant forecast accuracy gains of GNNHAR-1L over GHAR provide direct empirical support for nonlinear spillover effects. In practical terms, the influence of one coin's volatility on another is not a fixed linear proportion, instead it can vary with market context and extremes, consistent with phenomena like herd behavior or asymmetrical responses to market stress. These results echo recent studies that found improved volatility forecasts when nonlinear relationships and contagion channels are modeled. We therefore conclude that crypto volatility spillovers are inherently nonlinear, answering RQ2 in the affirmative.

- **Limited Role of Indirect Neighbors (RQ3)**. The contribution of indirect, multi hop connections to volatility spillovers appears to be minimal under the conditions studied. The two-layer GNNHAR model (GNNHAR-2L), which expands the neighborhood to second-degree connections, did not meaningfully improve forecasting performance over GNNHAR-1L. This indicates that volatility transmission is effectively captured by direct neighbors alone, with little added benefit from more distant connections. Stated differently, volatility shocks propagate rapidly through the network via first-order links, and additional hops mostly introduce noise or over-smoothing. The Model Confidence Set (MCS) analysis at 90% confidence reinforced this: GNNHAR-1L was the only model consistently included in the superior set for all 23 assets, whereas GNNHAR-2L (and the other models) were systematically excluded. We attribute the lack of multi hop benefit to the high density of the crypto volatility network, where most assets are closely interconnected by short paths. No single asset serves as a critical volatility bridge, instead, overlapping connections allow spillovers to disseminate in one step. Consequently, extending the receptive field further (as GNNHAR-2L does) yields diminishing returns and even risks overfitting by diluting asset-specific signals. Thus, RQ3 is answered in the negative: indirect neighbors do not significantly add to volatility forecasting accuracy in this setting.

These findings collectively paint a coherent picture of volatility interdependence in cryptocurrency markets. The market behaves as an integrated network segmented by volatility magnitude rather than by consensus mechanism. Assets with similar volatility profiles cluster together and exhibit strong spillover linkages, regardless of being "green" or "brown."

Meanwhile, extreme-volatility assets tend to stand apart at the periphery with weaker connections. Crucially, spillover effects within this network are nonlinear but predominantly confined to immediate network neighbors. Next, we discuss the broader implications of these insights and summarize model performance patterns across different asset clusters.

Implications for Crypto Volatility Modeling

Improved Forecasting through Network-Based Models: The clear outperformance of the GNNHAR-1L model over traditional approaches has important implications for volatility modeling and risk management in cryptocurrency markets. It demonstrates that capturing the web of interactions between cryptoassets yields substantial forecasting gains, especially when nonlinear aggregation is used. Conventional volatility models (such as GARCH or HAR) that treat assets in isolation are likely to miss these cross-asset influences. Practitioners and researchers should therefore consider network-based volatility models for multi-asset crypto portfolios, as these models can better anticipate volatility surges that originate from or spread through connected assets. The results here confirm recent research that graph neural networks can harness complex spillover patterns to enhance volatility predictions. In practice, incorporating network spillovers could improve risk assessment, for example by warning that a volatility spike in a major coin may ripple out to others. Our findings also suggest that monitoring volatility networks (e.g. through rolling correlation or precision matrices) can identify which assets are likely to influence each other, thereby providing more responsive trading or hedging strategies.

Nonlinearity and Model Design: The evidence of nonlinear spillovers implies that linear models may systematically underestimate or misspecify the true spillover effects, especially during turbulent market conditions. Nonlinear behaviors, such as herding effects, contagion during crashes, or asymmetric reactions to news, mean that volatility transmission can intensify non-proportionally under stress. Models like GNNHAR-1L, with nonlinear activation functions, are better equipped to capture these regime-dependent dynamics. This finding encourages the use of nonlinear modeling elements (e.g. neural networks, regime-switching models, or other machine learning techniques) in volatility forecasting, as they can adapt to the changing relationship between assets when markets become frothy or fearful. Future model developers should balance complexity and interpretability, but our work makes a strong case that allowing for nonlinearity is crucial when modeling interconnected financial markets.

Network Density and Optimal Spillover Horizon: The limited utility of second-hop spillovers in our results is a reminder that more complexity is not always better, it depends on the structure of the network. In a densely connected market like cryptocurrencies, most relevant information for forecasting an asset's volatility is contained in its immediate neighbors (directly correlated assets). Adding a second layer of neighbors (which effectively brings almost the entire network into view, given the short path lengths) led to oversmoothing, where unique local information was washed out by broad market noise. This suggests that for highly entangled markets, restricting models to one hop relationships (or otherwise carefully regularizing multi hop information) is preferable. In contrast, for a more fragmented or hierarchical market, multi hop effects might matter, but in crypto's case, the tightly-knit structure means one hop is sufficient to capture spillovers. Modelers should therefore tailor the neighborhood depth to the network topology at hand. In our context, the GNNHAR-1L struck the best balance, while GNNHAR-2L effectively "over-counted" information that was already reflected via closer connections. This insight also aligns with the centrality analysis of our crypto network, we observed high overall connectivity with low betweenness centrality, implying no long chain of intermediaries is needed for volatility to spread. Practically, this means risk propagation in crypto markets is swift and mostly direct, an important consideration for contagion risk management.

Volatility Clustering vs. Asset Characteristics: One of the striking findings of this study is that volatility level, not consensus mechanism, drives the clustering of cryptocurrencies in the volatility network. Despite the initial "brown vs green" hypothesis, we found no statistically significant differences in how PoW (energy-intensive) and PoS (energy-efficient) coins transmit or receive volatility shocks. Both types are mixed within the spillover network, grouped instead by their stability/volatility profiles. For example, low-volatility "core" assets (e.g. BTC, ETH, LTC) formed one cluster, whereas more volatile "amplifier" assets (e.g. ADA, DOT, BNB, SOL) formed another community, and the most volatile coins (e.g. CFX, MINA, TRX) were largely isolated as peripheral nodes. This outcome has several implications. First, it suggests that from a modeling perspective, one should account for volatility tiers, assets of similar risk level tend to influence each other more. Thus, segmenting the market by volatility (e.g. high, medium, low vol groups) could improve model specification or risk management strategies. Second, the lack of difference between PoW and PoS in spillover behavior means that, at least over 2023-2024, environmental or technological distinctions did not alter volatility interaction patterns. Investors and regulators focusing on "green crypto" versus "brown crypto" should note that these labels alone do not insulate an asset from market-wide volatility events. A supposedly green asset like ETH (which transitioned to PoS) can be as integrated in the volatility network as a PoW asset, if its volatility profile is similar. For portfolio allocation, this means diversification needs to consider volatility characteristics more than consensus mechanism, holding a mix of low-vol coins vs high-vol coins may achieve more risk reduction than simply mixing PoW and PoS assets.

Heterogeneous Benefits Across Assets: An important practical insight from our forecasting results is that the magnitude of improvement from using network-based nonlinear models is heterogeneous across assets. Highly connected assets in the spillover network saw the greatest forecasting accuracy gains from the GNNHAR-1L model, whereas extremely volatile, isolated assets saw only modest gains. Table 5 below summarizes the best-performing model for each cryptocurrency, grouped by its volatility cluster. We observe that for every asset, GNNHAR-1L was the top model (as indicated by the lowest out-of-sample MSE). However, the degree to which GNNHAR-1L outperformed simpler models varies by cluster. In the core and amplifier clusters (low- and medium-volatility groups) ,which contain the most interconnected coins, GNNHAR-1L often achieved substantially lower MSE than HAR or GHAR, reflecting the significant spillover information these assets exchange. For instance, assets like DOT, ADA, BNB, and ETH (all among the more connected coins) benefited greatly from incorporating network effects, with error reductions that were much larger than average. In contrast, for peripheral high-volatility assets such as SC, TRX, and MINA, the HAR model and GHAR model performed nearly as well as GNNHAR-1L, since these assets have few strong linkages to exploit. In such cases, most of their volatility is idiosyncratic, so a univariate approach captures the bulk of the predictable component. This pattern reinforces the idea that the value of sophisticated spillover modeling is highest for assets embedded in the network's core, and lowest for those on the fringe.

Tabela 12 - Best-performing volatility forecast model by asset cluster. All assets, regardless of cluster, achieved their lowest prediction error with the GNNHAR-1L model, indicating its overall superiority. However, the performance gap between GNNHAR-1L and the simpler models varies by cluster. Core and amplifier (medium-volatility) assets, which are well-integrated in the network, show a large improvement of GNNHAR-1L over the HAR baseline, confirming strong spillover effects. Peripheral assets, being largely idiosyncratic, see much smaller gains; for these, the HAR and GHAR models often performed almost on par. This underscores that network spillover modeling yields the greatest benefit for interconnected assets, whereas isolated coins' volatility is driven more by asset-specific factors. Notably, in all clusters, the two-layer GNNHAR-2L did not top any asset's ranking, consistent with our earlier conclusion that including second-hop neighbors does not improve forecasts in a dense volatility network.

Cluster (Volatility Tier)	**Assets in Cluster**	**Best Model (Rankings)**
Core - Low Volatility	BTC, ETH, LTC, BCH, DOGE, ETC	GNNHAR-1L (best for all core assets; big difference in performance between GHAR and HAR)
Amplifier - Medium Volatility	ADA, DOT, BNB, SOL, ALGO, AVAX, CELO, EGLD, INJ, XTZ, RVN, ZEC	GNNHAR-1L (best for all amplifier assets; also big difference in performance between GHAR and HAR)
Peripheral - High Volatility	CFX, CKB, SC, MINA, TRX	GNNHAR-1L (best for all peripheral assets; GHAR ≈ HAR in terms of performance)

Final Remarks

In conclusion, this research establishes that cryptocurrency volatility is contagious (spillovers exist), nonlinear, and primarily constrained to direct connections in a tightly knit network. The distinction between "brown" and "green" cryptocurrencies is largely irrelevant for volatility transmission, instead, it is an asset's own volatility magnitude and network connectivity that govern how it influences and is influenced by others. These insights contribute to the growing literature on graph-based financial modeling by demonstrating the efficacy of GNN approaches in capturing complex dependency structures in high-frequency data. From a practical standpoint, the findings encourage risk managers to monitor volatility linkages and incorporate them into forecasting models, especially for portfolios containing highly interconnected cryptoassets.

Future Research Directions

This study opens several promising avenues for future research. First, the inclusion of additional financial variables, such as a broader set of cryptocurrencies, financial assets such as equities, commodities and index funds, or macroeconomic indicators like the S&P 500 or MSCI World could enrich the information set available to the model and reveal cross-market volatility transmission mechanisms that extend beyond the cryptocurrency space.

Second, although this study focused on daily realized volatility based on 5-minute returns, future research could investigate longer horizons (weekly or monthly volatility) to determine whether nonlinear and multi hop spillovers emerge more clearly across time. Prior studies show that volatility spillovers exhibit horizon dependent behavior, underscoring the importance of multi horizon modelling (Mensi, Rehman, et al., 2024)(Kakinaka & Umeno, 2022)(Catania & Grassi, 2022).

Third, the graph construction methodology itself presents opportunities for refinement. While this study used GLASSO (Friedman et al., 2008) to infer time-varying conditional dependence networks, alternative approaches such as empirical approaches based on feature importance (Karpman et al., 2023), Attention-based GNNs and Mutual Information approach may uncover different network structures and further improve the interpretability and predictive power of graph based volatility forecasting models. The presence of nonlinearity in the relationships between cryptoassets further validates the incorporation of these extensions as they can capture nonlinear, non-Gaussian and asymmetric dependencies with the use of Attention-based GNNs (Uddin et al., 2023)(Lei et al., 2024) even considering directionality, which can lead to a deeper understanding of interconnectedness in financial markets.

Fourth, future research could investigate the role of investor heterogeneity in shaping volatility spillovers. Recent studies show the existence of different cryptocurrency owner types, from transactors who solely use cryptocurrencies for transactions and payments, investors, a mix user which uses cryptos for both investments and transactional purposes and traders (Hayashi & Routh, 2025)(Ante et al., 2023). These different types of crypto owners have distinct levels of financial literacy, risk tolerance, and various demographic and financial characteristics. The evidenced differences in financial literacy and risk tolerance by owner type may influence how shocks propagate, as more cautious investors concentrate in low-volatility assets while risk-seeking participants gravitate toward highly volatile coins. Bridging the gap between behavioral evidence on owner heterogeneity with network-based volatility models would provide a richer understanding of how portfolio choices, risk preferences, and market participation contribute to the clustering of cryptocurrencies and to the intensity of spillover dynamics.

Bibliographical References

Abdulrahman Abubakar, A., Clement Mba, J., & Ohonba, A. (2024). Analysis of tail dependence structure and risk spillover between cryptocurrencies. *Investment Management and Financial Innovations*, *21*(4), 140–155. https://doi.org/10.21511/imfi.21(4).2024.12

Agyei, S. K., Adam, A. M., Bossman, A., Asiamah, O., Owusu Junior, P., Asafo-Adjei, R., & Asafo-Adjei, E. (2022). Does volatility in cryptocurrencies drive the interconnectedness between the cryptocurrencies market? Insights from wavelets. *Cogent Economics & Finance*, *10*(1), 2061682. https://doi.org/10.1080/23322039.2022.2061682

Ahmed, M. S., El-Masry, A. A., Al-Maghyereh, A. I., & Kumar, S. (2024). Cryptocurrency volatility: A review, synthesis, and research agenda. *Research in International Business and Finance*, 102472.

Akin, I., Akin, M., Ozturk, Z., Hameed, A., Opara, V., & Satiroglu, H. (2024). Exploring fluctuations and interconnected movements in stock, commodity, and cryptocurrency markets. *British Actuarial Journal*, *29*, e13. https://doi.org/10.1017/S1357321724000126

Alamaren, A. S., Gokmenoglu, K. K., & Taspinar, N. (2024). Volatility spillovers among leading cryptocurrencies and US energy and technology companies. *Financial Innovation*, *10*(1), 81. https://doi.org/10.1186/s40854-024-00626-2

Ammar, I. B., & Hellara, S. (2022). High-frequency trading, stock volatility, and intraday crashes. *The Quarterly Review of Economics and Finance*, *84*, 337–344.

Andersen, T. G., & Bollerslev, T. (1998). Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts. *International Economic Review*, *39*(4), 885. https://doi.org/10.2307/2527343

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and Forecasting Realized Volatility. *Econometrica*, *71*(2), 579–625. https://doi.org/10.1111/1468-0262.00418

Ante, L., Fiedler, F., Steinmetz, F., & Fiedler, I. (2023). Profiling Turkish cryptocurrency owners: Payment users, crypto investors and crypto traders. *Journal of Risk and Financial Management*, *16*(4), 239.

Arora, H. (2017). Stylized facts and trends of high frequency data in financial markets. *Asian Journal of Research in Business Economics and Management*, *7*(7), 303–317.

Arroyo, Á., Gravina, A., Gutteridge, B., Barbero, F., Gallicchio, C., Dong, X., Bronstein, M., & Vandergheynst, P. (2025). *On Vanishing Gradients, Over-Smoothing, and Over-Squashing in GNNs: Bridging Recurrent and Graph Learning* (No. arXiv:2502.10818). arXiv. https://doi.org/10.48550/arXiv.2502.10818

Asiri, A., Alnemer, M., & Bhatti, M. I. (2023). Interconnectedness of cryptocurrency uncertainty indices with returns and volatility in financial assets during COVID-19. *Journal of Risk and Financial Management*, *16*(10), 428.

Aslanidis, N., Bariviera, A. F., & Perez-Laborda, A. (2021). Are cryptocurrencies becoming more interconnected? *Economics Letters*, *199*, 109725. https://doi.org/10.1016/j.econlet.2021.109725

Bai, Y., Yang, E., Han, B., Yang, Y., Li, J., Mao, Y., Niu, G., & Liu, T. (2021). Understanding and improving early stopping for learning with noisy labels. *Advances in Neural Information Processing Systems*, *34*, 24392–24403.

Balcilar, M., & Ozdemir, H. (2023). On the Risk Spillover from Bitcoin to Altcoins: The Fear of Missing Out and Pump-and-Dump Scheme Effects. *Journal of Risk and Financial Management*, *16*(1), Article 1. https://doi.org/10.3390/jrfm16010041

Bariviera, A. F., Zunino, L., & Rosso, O. A. (2018). An analysis of high-frequency cryptocurrencies prices dynamics using permutation-information-theory quantifiers. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *28*(7), 075511. https://doi.org/10.1063/1.5027153

Barndorff-Nielsen, O. E., & Shephard, N. (2001). Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *63*(2), 167–241. https://doi.org/10.1111/1467-9868.00282

Barndorff-Nielsen, O. E., & Shephard, N. (2002). Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *64*(2), 253–280. https://doi.org/10.1111/1467-9868.00336

Baur, D. G., & Dimpfl, T. (2018). Asymmetric volatility in cryptocurrencies. *Economics Letters*, *173*, 148–151. https://doi.org/10.1016/j.econlet.2018.10.008

Ben Osman, M., Galariotis, E., Guesmi, K., Hamdi, H., & Naoui, K. (2024). Are markets sentiment driving the price bubbles in the virtual? *International Review of Economics & Finance*, *89*, 272–285. https://doi.org/10.1016/j.iref.2023.10.041

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

Bouoiyour, J., & Selmi, R. (2019). How do futures contracts affect Bitcoin prices ? *Economics Bulletin*, *39*(2), 1127–1134.

Bouri, E., Shahzad, S. J. H., & Roubaud, D. (2019). Co-explosivity in the cryptocurrency market. *Finance Research Letters*, *29*, 178–183. https://doi.org/10.1016/j.frl.2018.07.005

Brini, A., & Lenz, J. (2024). A comparison of cryptocurrency volatility-benchmarking new and mature asset classes. *Financial Innovation*, *10*(1), 122. https://doi.org/10.1186/s40854-024-00646-y

Brini, A., & Toscano, G. (2024). SpotV2Net: Multivariate intraday spot volatility forecasting via vol-of-vol-informed graph attention networks. *International Journal of Forecasting*. https://doi.org/10.1016/j.ijforecast.2024.11.004

Brogaard, J. (2024). *Attention-based Graph Neural Networks in Firm CDS Prediction*.

Catania, L., & Grassi, S. (2022). Forecasting cryptocurrency volatility. *International Journal of Forecasting*, *38*(3), 878–894.

Chaim, P., & Laurini, M. P. (2018). Volatility and return jumps in bitcoin. *Economics Letters*, *173*, 158–163.

Chang, W.-Y., & Lin, T.-N. (2019). *Vanishing Nodes: Another Phenomenon That Makes Training Deep Neural Networks Difficult* (No. arXiv:1910.09745). arXiv. https://doi.org/10.48550/arXiv.1910.09745

Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*, *7*(2), 174–196. https://doi.org/10.1093/jjfinec/nbp001

Dang, T. H.-N., Nguyen, C. P., Lee, G. S., Nguyen, B. Q., & Le, T. T. (2023). Measuring the energy-related uncertainty index. *Energy Economics*, *124*, 106817. https://doi.org/10.1016/j.eneco.2023.106817

de la Concha, A., Martinez-Jaramillo, S., & Carmona, C. (2018). Multiplex Financial Networks: Revealing the Level of Interconnectedness in the Banking System. In C. Cherifi, H. Cherifi, M. Karsai, & M. Musolesi (Eds.), *Complex Networks & Their Applications VI* (pp. 1135–1148). Springer International Publishing. https://doi.org/10.1007/978-3-319-72150-7_92

Diebold, F. X., & Yilmaz, K. (2009). Measuring Financial Asset Return and Volatility Spillovers, with Application to Global Equity Markets. *The Economic Journal*, *119*(534), 158–171. https://doi.org/10.1111/j.1468-0297.2008.02208.x

Drożdż, S., Minati, L., Oświęcimka, P., Stanuszek, M., & Wątorek, M. (2020). Competition of noise and collectivity in global cryptocurrency trading: Route to a self-contained market. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *30*(2), 023122. https://doi.org/10.1063/1.5139634

Engle, R. (2002). Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business & Economic Statistics*, *20*(3), 339–350. https://doi.org/10.1198/073500102288618487

Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, *50*(4), 987–1007. https://doi.org/10.2307/1912773

Engle, R. F., & Kroner, K. F. (1995). Multivariate Simultaneous Generalized ARCH. *Econometric Theory*, *11*(1), 122–150. https://doi.org/10.1017/S0266466600009063

Forino, A., & Morelli, G. (2024). Liquidity and Volatility Nexus in Cryptocurrency Markets. *Available at SSRN 5128381*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5128381

Friederich, F., Meyer, J., Matute, J., & Palau-Saumell, R. (2024). CRYPTO-MANIA: How fear-of-missing-out drives consumers' (risky) investment decisions. *Psychology & Marketing*, *41*(1), 102–117. https://doi.org/10.1002/mar.21906

Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, *9*(3), 432–441. https://doi.org/10.1093/biostatistics/kxm045

Gkillas, K., & Katsiampa, P. (2018). An application of extreme value theory to cryptocurrencies. *Economics Letters*, *164*, 109–111.

Gkillas, K., Katsiampa, P., Konstantatos, C., & Tsagkanos, A. (2024). Discontinuous movements and asymmetries in cryptocurrency markets. *The European Journal of Finance*, *30*(16), 1907–1931. https://doi.org/10.1080/1351847X.2021.2015416

Grobys, K., & Junttila, J. (2021). Speculation and lottery-like demand in cryptocurrency markets. *Journal of International Financial Markets, Institutions and Money*, *71*, 101289.

Gualpa, E. V. M., Urgilés, C. H. F., Zenteno, J. A. C., & Castillo, L. F. P. (2023). Técnicas de minado de criptomonedas: Proof of Work vs Proof of Stake. *Pro Sciences: Revista de Producción, Ciencias e Investigación*, *7*(47), 23–35.

Hall, M. K., & Jasiak, J. (2024). Modelling common bubbles in cryptocurrency prices. *Economic Modelling*, *139*, 106782. https://doi.org/10.1016/j.econmod.2024.106782

Hansen, P. R., Lunde, A., & Nason, J. M. (2011). The Model Confidence Set. *Econometrica*, *79*(2), 453–497. https://doi.org/10.3982/ECTA5771

Hayashi, F., & Routh, A. (2025). Financial literacy, risk tolerance, and cryptocurrency ownership in the United States. *Journal of Behavioral and Experimental Finance*, 101060.

Haykir, O., & Yagli, I. (2022). Speculative bubbles and herding in cryptocurrencies. *Financial Innovation*, *8*(1), 78. https://doi.org/10.1186/s40854-022-00383-0

Heß, V. L., & Damasio, B. (2025). Machine learning in banking risk management: Mapping a decade of evolution. *International Journal of Information Management Data Insights*, *5*(1), 100324.

Hung, N. T. (2022). Asymmetric connectedness among S&P 500, crude oil, gold and Bitcoin. *Managerial Finance*, *48*(4), 587–610. https://doi.org/10.1108/MF-08-2021-0355

Hussein, B. M., & Shareef, S. M. (2024). An Empirical Study on the Correlation between Early Stopping Patience and Epochs in Deep Learning. *ITM Web of Conferences*, *64*, 01003. https://doi.org/10.1051/itmconf/20246401003

Javaheri, A., Ying, J., Palomar, D. P., & Marvasti, F. (2025). Time-Varying Graph Learning for Data with Heavy-Tailed Distribution. *IEEE Transactions on Signal Processing*. https://ieeexplore.ieee.org/abstract/document/11080072/

Jiang, X., & Pu, Y. (2023). *Exploring Time Granularity on Temporal Graphs for Dynamic Link Prediction in Real-world Networks* (No. arXiv:2311.12255). arXiv. https://doi.org/10.48550/arXiv.2311.12255

Kakinaka, S., & Umeno, K. (2022). Asymmetric volatility dynamics in cryptocurrency markets on multi-time scales. *Research in International Business and Finance*, *62*, 101754.

Karagiorgis, A., Ballis, A., & Drakos, K. (2024). The Skewness-Kurtosis plane for cryptocurrencies' universe. *International Journal of Finance & Economics*, *29*(2), 2543–2555. https://doi.org/10.1002/ijfe.2795

Karam, A. (2018). The effects of intraday news flow on dealers' quotations, market liquidity, and volatility. *International Journal of Finance & Economics*, *23*(4), 492–503. https://doi.org/10.1002/ijfe.1634

Karim, M. M., Shah, M. E., Noman, A. H. Md., & Yarovaya, L. (2024). Exploring asymmetries in cryptocurrency intraday returns and implied volatility: New evidence for high-frequency traders. *International Review of Financial Analysis*, *96*, 103617. https://doi.org/10.1016/j.irfa.2024.103617

Karpman, K., Basu ,Sumanta, Easley ,David, & and Kim, S. (2023). Learning Financial Networks with High-Frequency Trade Data. *Data Science in Science*, *2*(1), 2166624. https://doi.org/10.1080/26941899.2023.2166624

Keidar, R., & Blemus, S. (2018). *Cryptocurrencies and Market Abuse Risks: It's Time for Self-Regulation* (SSRN Scholarly Paper No. 3123881). Social Science Research Network. https://doi.org/10.2139/ssrn.3123881

Kipf, T. N., & Welling, M. (2017). *Semi-Supervised Classification with Graph Convolutional Networks* (No. arXiv:1609.02907). arXiv. https://doi.org/10.48550/arXiv.1609.02907

Krupa, T., Ries, M., Kotuliak, I., Košťál, K., & Bencel, R. (2021). Security Issues of Smart Contracts in Ethereum Platforms. *2021 28th Conference of Open Innovations Association (FRUCT)*, 208–214. https://doi.org/10.23919/FRUCT50888.2021.9347617

Kumar, P. N., Umeorah, N., & Alochukwu, A. (2024). *Dynamic graph neural networks for enhanced volatility prediction in financial markets* (No. arXiv:2410.16858). arXiv. https://doi.org/10.48550/arXiv.2410.16858

Lahmiri, S., & Bekiros, S. (2020). The impact of COVID-19 pandemic upon stability and sequential irregularity of equity and cryptocurrency markets. *Chaos, Solitons & Fractals*, *138*, 109936. https://doi.org/10.1016/j.chaos.2020.109936

Lara-Benítez, P., Carranza-García, M., & Riquelme, J. C. (2021). An Experimental Review on Deep Learning Architectures for Time Series Forecasting. *International Journal of Neural Systems*, *31*(03), 2130001. https://doi.org/10.1142/s0129065721300011

Lei, Z., Zhang, C., Xu, Y., & Li, X. (2024). DR-GAT: Dynamic routing graph attention network for stock recommendation. *Information Sciences*, *654*, 119833.

Li, M., Soltanolkotabi, M., & Oymak, S. (2020). Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. *International Conference on Artificial Intelligence and Statistics*, 4313–4324. https://proceedings.mlr.press/v108/li20j.html

Li, Z. L., Zhang, G. W., Yu, J., & Xu, L. Y. (2023). Dynamic graph structure learning for multivariate time series forecasting. *Pattern Recognition*, *138*, 109423.

Li, Z. M., Laeven, R. J. A., & Vellekoop, M. H. (2020). Dependent microstructure noise and integrated volatility estimation from high-frequency data. *Journal of Econometrics*, *215*(2), 536–558. https://doi.org/10.1016/j.jeconom.2019.10.004

Liang, Y., Zhang, W., Sheng, Z., Yang, L., Xu, Q., Jiang, J., Tong, Y., & Cui, B. (2025). Towards scalable and deep graph neural networks via noise masking. *Proceedings of the AAAI Conference on Artificial Intelligence*, *39*(18), 18693–18701. https://ojs.aaai.org/index.php/AAAI/article/view/34057

Liu, G., Zhuang, Z., & Wang, M. (2024). Forecasting the high-frequency volatility based on the LSTM-HIT model. *Journal of Forecasting*, *43*(5), 1356–1373. https://doi.org/10.1002/for.3078

Liu, J., Julaiti, J., & Gou, S. (2024). Decomposing interconnectedness: A study of cryptocurrency spillover effects in global financial markets. *Finance Research Letters*, *61*, 104950. https://doi.org/10.1016/j.frl.2023.104950

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. https://doi.org/10.1016/j.jeconom.2015.02.008

Makarov, I., & Schoar, A. (2022). Cryptocurrencies and decentralized finance (DeFi). *Brookings Papers on Economic Activity*, *2022*(1), 141–215.

Mazor, O., & Rottenstreich, O. (2024). Understanding the Blockchain Interoperability Graph Based on Cryptocurrency Price Correlation. *2024 IEEE International Conference on Blockchain (Blockchain)*, 392–399. https://doi.org/10.1109/Blockchain62396.2024.00058

Mensi, W., Gubareva, M., & Kang, S. H. (2024). Frequency connectedness between DeFi and cryptocurrency markets. *The Quarterly Review of Economics and Finance*, *93*, 12–27. https://doi.org/10.1016/j.qref.2023.11.001

Mensi, W., Gubareva, M., Ko, H.-U., Vo, X. V., & Kang, S. H. (2023). Tail spillover effects between cryptocurrencies and uncertainty in the gold, oil, and stock markets. *Financial Innovation*, *9*(1), 92. https://doi.org/10.1186/s40854-023-00498-y

Mensi, W., Rehman, M. U., Vo, X. V., & Kang, S. H. (2024). Spillovers and multiscale relationships among cryptocurrencies: A portfolio implication using high frequency data. *Economic Analysis and Policy*, *82*, 449–479.

Milunovich, G. (2022). Assessing the connectedness between Proof of Work and Proof of Stake/Other digital coins. *Economics Letters*, *211*, 110243. https://doi.org/10.1016/j.econlet.2021.110243

Miura, R., Pichl, L., & Kaizoji, T. (2019). Artificial Neural Networks for Realized Volatility Prediction in Cryptocurrency Time Series. In H. Lu, H. Tang, & Z. Wang (Eds.), *Advances in Neural Networks – ISNN 2019* (pp. 165–172). Springer International Publishing. https://doi.org/10.1007/978-3-030-22796-8_18

Moratis, G. (2021). Quantifying the spillover effect in the cryptocurrency market. *Finance Research Letters*, *38*, 101534. https://doi.org/10.1016/j.frl.2020.101534

Mücher, C. (2022). Artificial Neural Network Based Non-linear Transformation of High-Frequency Returns for Volatility Forecasting. *Frontiers in Artificial Intelligence*, *4*, 787534. https://doi.org/10.3389/frai.2021.787534

Naeem, M. A., Hamouda, F., Karim, S., & Vigne, S. A. (2023). Return and volatility spillovers among global assets: Comparing health crisis with geopolitical crisis. *International Review of Economics & Finance*, *87*, 557–575. https://doi.org/10.1016/j.iref.2023.06.008

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. *Satoshi Nakamoto*. https://www.poritz.net/jonathan/past_classes/winter16/CCatRU/BitcoinOriginalPaper.pdf

Nekhili, R. (2020). Are bitcoin futures contracts for hedging or speculation. *Investment Management and Financial Innovations*, *17*(3), 1–9.

Nguyen, A. P. N., Crane, M., Conlon, T., & Bezbradica, M. (2025). Herding unmasked: Insights into cryptocurrencies, stocks and US ETFs. *PLOS ONE*, *20*(2), e0316332. https://doi.org/10.1371/journal.pone.0316332

Nguyen, A. P. N., Mai, T. T., Bezbradica, M., & Crane, M. (2023). Volatility and returns connectedness in cryptocurrency markets: Insights from graph-based methods. *Physica A: Statistical Mechanics and Its Applications*, *632*, 129349. https://doi.org/10.1016/j.physa.2023.129349

Pacelli, V., Di Tommaso, C., Foglia, M., & Ingannamorte, S. (2025). Cryptocurrencies and Systemic Risk. The Spillover Effects Between Cryptocurrency and Financial Markets. *Systemic Risk and Complex Networks in Modern Financial Systems*, 343.

Peng, P., Chen, Y., Lin, W., & Wang, J. Z. (2024). Attention-based CNN–LSTM for high-frequency multiple cryptocurrency trend prediction. *Expert Systems with Applications*, *237*, 121520. https://doi.org/10.1016/j.eswa.2023.121520

Pessa, A. A., Perc, M., & Ribeiro, H. V. (2023). Age and market capitalization drive large price variations of cryptocurrencies. *Scientific Reports*, *13*(1), 3351.

Petukhina, A. A., Reule, R. C. G., & Härdle, W. K. (2021). Rise of the machines? Intraday high-frequency trading patterns of cryptocurrencies. *The European Journal of Finance*, *27*(1–2), 8–30. https://doi.org/10.1080/1351847X.2020.1789684

Poddar, A., Misra, A. K., & Mishra, A. K. (2023). Return connectedness and volatility dynamics of the cryptocurrency network. *Finance Research Letters*, *58*, 104334.

Qian, S., Ying, H., Hu, R., Zhou, J., Chen, J., Chen, D. Z., & Wu, J. (2023). Robust Training of Graph Neural Networks via Noise Governance. *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, 607–615. https://doi.org/10.1145/3539597.3570369

Raddant, M., & Kenett, D. Y. (2021). Interconnectedness in the global financial market. *Journal of International Money and Finance*, *110*, 102280. https://doi.org/10.1016/j.jimonfin.2020.102280

Rusch, T. K., Bronstein, M. M., & Mishra, S. (2023). *A Survey on Oversmoothing in Graph Neural Networks* (No. arXiv:2303.10993). arXiv. https://doi.org/10.48550/arXiv.2303.10993

Scalera, J. E., & and Dixon, M. D. (2016). Crisis of confidence: The 2008 global financial crisis and public trust in the European Central Bank. *European Politics and Society*, *17*(3), 388–400. https://doi.org/10.1080/23745118.2016.1168970

Sebastião, H., & Godinho, P. (2020). Bitcoin futures: An effective tool for hedging cryptocurrencies. *Finance Research Letters*, *33*, 101230. https://doi.org/10.1016/j.frl.2019.07.003

Sheppard, K., & Xu, W. (2019). Factor High-Frequency-Based Volatility (HEAVY) Models*. *Journal of Financial Econometrics*, *17*(1), 33–65. https://doi.org/10.1093/jjfinec/nby028

Shu, J., Xi, B., Li, Y., Wu, F., Kamhoua, C., & Ma, J. (2022). Understanding Dropout for Graph Neural Networks. *Companion Proceedings of the Web Conference 2022*, 1128–1138. https://doi.org/10.1145/3487553.3524725

Son, B., Lee, Y., Park, S., & Lee, J. (2023). Forecasting global stock market volatility: The impact of volatility spillover index in spatial-temporal graph-based model. *Journal of Forecasting*, *42*(7), 1539–1559. https://doi.org/10.1002/for.2975

Song, H., Kim, M., Park, D., & Lee, J.-G. (2020). *How does Early Stopping Help Generalization against Label Noise?* (No. arXiv:1911.08059). arXiv. https://doi.org/10.48550/arXiv.1911.08059

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, *15*(1), 1929–1958.

Uddin, A., Tao, X., & Yu, D. (2023). Attention based dynamic graph neural network for asset pricing. *Global Finance Journal*, *58*, 100900.

Ünlü, U., & Bayram, V. (2024). Analysis of Bitcoin Volatility during the COVID-19 Pandemic: An Examination Using ARCH and GARCH Models. *Ekonomi Politika ve Finans Araştırmaları Dergisi*, *9*(4), Article 4. https://doi.org/10.30784/epfad.1588310

Vardar, G., Taçoğlu, C., & Aydoğan, B. (2022). Quantifying Return and Volatility Spillovers among Major Cryptocurrencies: A VAR-BEKK-GARCH Analysis. *Eskişehir Osmangazi Üniversitesi İktisadi ve İdari Bilimler Dergisi*, *17*(3), Article 3. https://doi.org/10.17153/oguiibf.1145664

Wang, J.-N., Liu, H.-C., Lee, Y.-H., & Hsu, Y.-T. (2023). FoMO in the Bitcoin market: Revisiting and factors. *The Quarterly Review of Economics and Finance*, *89*, 244–253.

Wang, J.-N., Vigne, S. A., Liu, H.-C., & Hsu, Y.-T. (2024). Divergent jump characteristics in brown and green cryptocurrencies: The role of energy-related uncertainty. *Energy Economics*, *138*, 107847.

Warshaw, E. (2020). Asymmetric volatility spillover between European equity and foreign exchange markets: Evidence from the frequency domain. *International Review of Economics & Finance*, *68*, 1–14. https://doi.org/10.1016/j.iref.2020.03.001

Wu, X., Ajorlou, A., Wu, Z., & Jadbabaie, A. (2023). Demystifying oversmoothing in attention-based graph neural networks. *Advances in Neural Information Processing Systems*, *36*, 35084–35106.

Xu, Q., Zhang, Y., & Zhang, Z. (2021). Tail-risk spillovers in cryptocurrency markets. *Finance Research Letters*, *38*, 101453.

Xu, Z., Liechty, J., Benthall, S., Skar-Gislinge, N., & McComb, C. (2024). GARCH-Informed Neural Networks for Volatility Prediction in Financial Markets. *Proceedings of the 5th ACM International Conference on AI in Finance*, 600–607. https://doi.org/10.1145/3677052.3698600

Yan, S., & Yan, D. (2019). Volatility Estimation in the Era of High-Frequency Finance. *FinTech as a Disruptive Technology for Financial Institutions*, 99–141.

Yang, F., Zhang, H., & Tao, S. (2022). Simplified multilayer graph convolutional networks with dropout. *Applied Intelligence*, *52*(5), 4776–4791. https://doi.org/10.1007/s10489-021-02617-7

Yi, S., Xu, Z., & Wang, G.-J. (2018). Volatility connectedness in the cryptocurrency market: Is Bitcoin a dominant cryptocurrency? *International Review of Financial Analysis*, *60*, 98–114.

Zaheer, S., Anjum, N., Hussain, S., Algarni, A. D., Iqbal, J., Bourouis, S., & Ullah, S. S. (2023). A multi parameter forecasting for stock time series data using LSTM and deep learning model. *Mathematics*, *11*(3), 590.

Zhang, C., Pu, X., Cucuringu, M., & Dong, X. (2025a). Forecasting realized volatility with spillover effects: Perspectives from graph neural networks. *International Journal of Forecasting*, *41*(1), 377–397. https://doi.org/10.1016/j.ijforecast.2024.09.002

Zhang, C., Pu, X., Cucuringu, M., & Dong, X. (2025b). Graph-Based Methods for Forecasting Realized Covariances. *Journal of Financial Econometrics*, *23*(2), nbae026. https://doi.org/10.1093/jjfinec/nbae026

Zhang, X. (2022). A model combining LightGBM and neural network for high-frequency realized volatility forecasting. *2022 7th International Conference on Financial Innovation and Economic Development (ICFIED 2022)*, 2906–2912. https://www.atlantis-press.com/proceedings/icfied-22/125971589

Zhao, P., Zhu, H., Ng, W. S. H., & Lee, D. L. (2024). From GARCH to neural network for volatility forecast. *Proceedings of the AAAI Conference on Artificial Intelligence*, *38*(15), 16998–17006. https://ojs.aaai.org/index.php/AAAI/article/view/29643

Zhou, Y., Xie, C., Wang, G.-J., Gong, J., & Zhu, Y. (2025). Forecasting cryptocurrency volatility: A novel framework based on the evolving multiscale graph neural network. *Financial Innovation*, *11*(1), 87. https://doi.org/10.1186/s40854-025-00768-x

Zingales, L. (2011). The role of trust in the 2008 financial crisis. *The Review of Austrian Economics*, *24*(3), 235–249. https://doi.org/10.1007/s11138-010-0134-0

APPENDIX A : ENSURING STATIONARITY OF LOG RETURN TIME SERIES

Before proceeding with model estimation and forecasting, it is essential to verify that the log returns time series for each cryptocurrency is stationary. In time series modeling, particularly in financial econometrics, stationarity refers to the property that a process's statistical characteristics such as mean, variance, and seasonality remain constant over time.

In this study, stationarity is a critical prerequisite for several reasons:

- **Model Validity**: GHAR and GNNHAR models assume that the input features (daily realized volatility and lagged features) are drawn from stationary processes. Non-stationary inputs can lead to unstable parameter estimates and unreliable forecasts.
- **Consistent Volatility Estimation**: The realized volatility estimator is designed under the assumption that the log return process exhibits constant statistical behavior over time. A non-stationary return series could bias the realized volatility calculation, undermining its effectiveness as an implied volatility proxy.

To this end, the log returns of each cryptocurrency were individually tested for stationarity using the Augmented Dickey-Fuller (ADF) test. In all cases, the ADF test rejected the null hypothesis of a unit root at the 1% significance level, indicating strong evidence of stationarity in the log return series.

Label	ADF Statistic	p-value	Stationary
BTC	-51.40	0	True
DOGE	-53.09	0	True
BCH	-52.17	0	True
LTC	-78.20	0	True
ETC	-52.46	0	True
ZEC	-68.29	0	True
CFX	-52.96	0	True
CKB	-54.08	0	True
SC	-54.38	0	True
RVN	-56.75	0	True
ETH	-52.86	0	True
SOL	-67.36	0	True
ADA	-91.36	0	True
ALGO	-54.35	0	True
INJ	-65.14	0	True
XTZ	-52.67	0	True
EGLD	-59.26	0	True
MINA	-52.07	0	True
CELO	-72.48	0	True
BNB	-52.05	0	True
DOT	-53.86	0	True
AVAX	-56.00	0	True
TRX	-44.44	0	True

Figure 11 - ADF test results for daily log returns

APPENDIX B : HANDLING OF MISSING VALUES

During the data preparation phase, it was observed that the Binance API returned incomplete high-frequency data for all cryptocurrencies on March 24th, 2023. Specifically, 16 consecutive 5-minute intervals (equivalent to 1 hour and 20 minutes) were missing across every asset.

Given that daily realized volatility is computed as the sum of squared intraday log returns, such a gap represents a non-negligible loss of information and imputing or interpolating high-frequency returns over this time span would risk distorting the volatility estimate for that day.

For this reason, in the effort of preserving the integrity and comparability of the daily realized volatility calculations across all assets and dates, the decision was made to remove March 24th, 2023 entirely from the dataset. Since the missing data affected only a single trading day within a nearly two-year sample, this exclusion had no material impact on the statistical power or continuity of the study.

APPENDIX C : SET OF HYPERPARAMETERS USED

After running Grid Search, these were the set of hyperparameters used for each model:

	HAR	GHAR	GNNHAR-1L	GNNHAR-2L
Nº Hidden Units	None	12	12	9
Nº of Epochs	100	50	50	60
Batch Size	16	16	16	16
Learning Rate	0.001	0.001	0.001	0.001
Weight Decay	1×10^{-4}	1×10^{-4}	5×10^{-4}	1×10^{-5}
Early Stop Patience	5	5	5	5
Dropout Rate	None	None	0.3	0.3

APPENDIX D : CRYPTOASSETS LOG RETURN DISTRIBUITIONS HISTOGRAMS

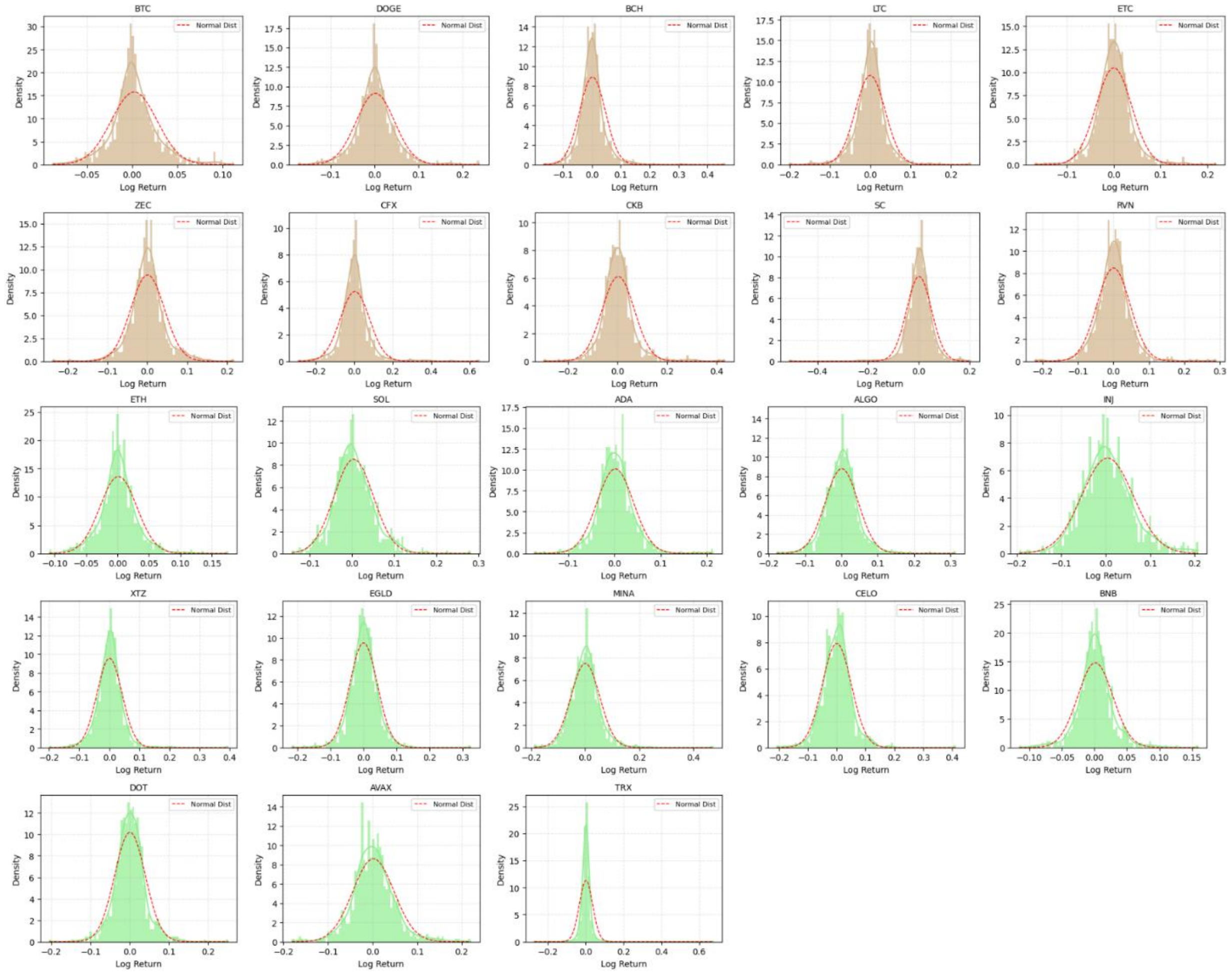

APPENDIX E : CRYPTOASSETS LOG RETURN Q-Q PLOTS

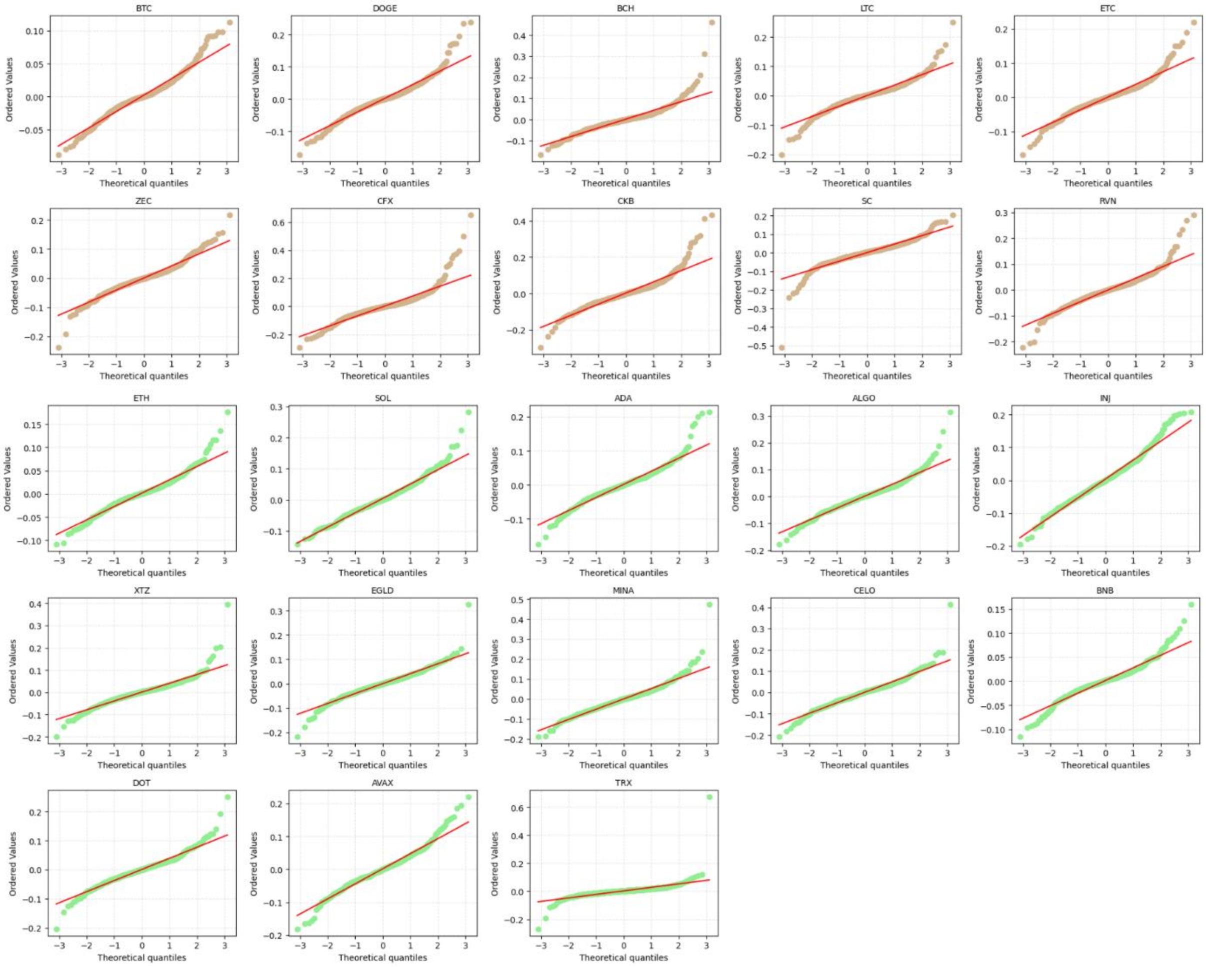

APPENDIX F : CRYPTOASSETS LOG RETURN VIOLIN PLOTS

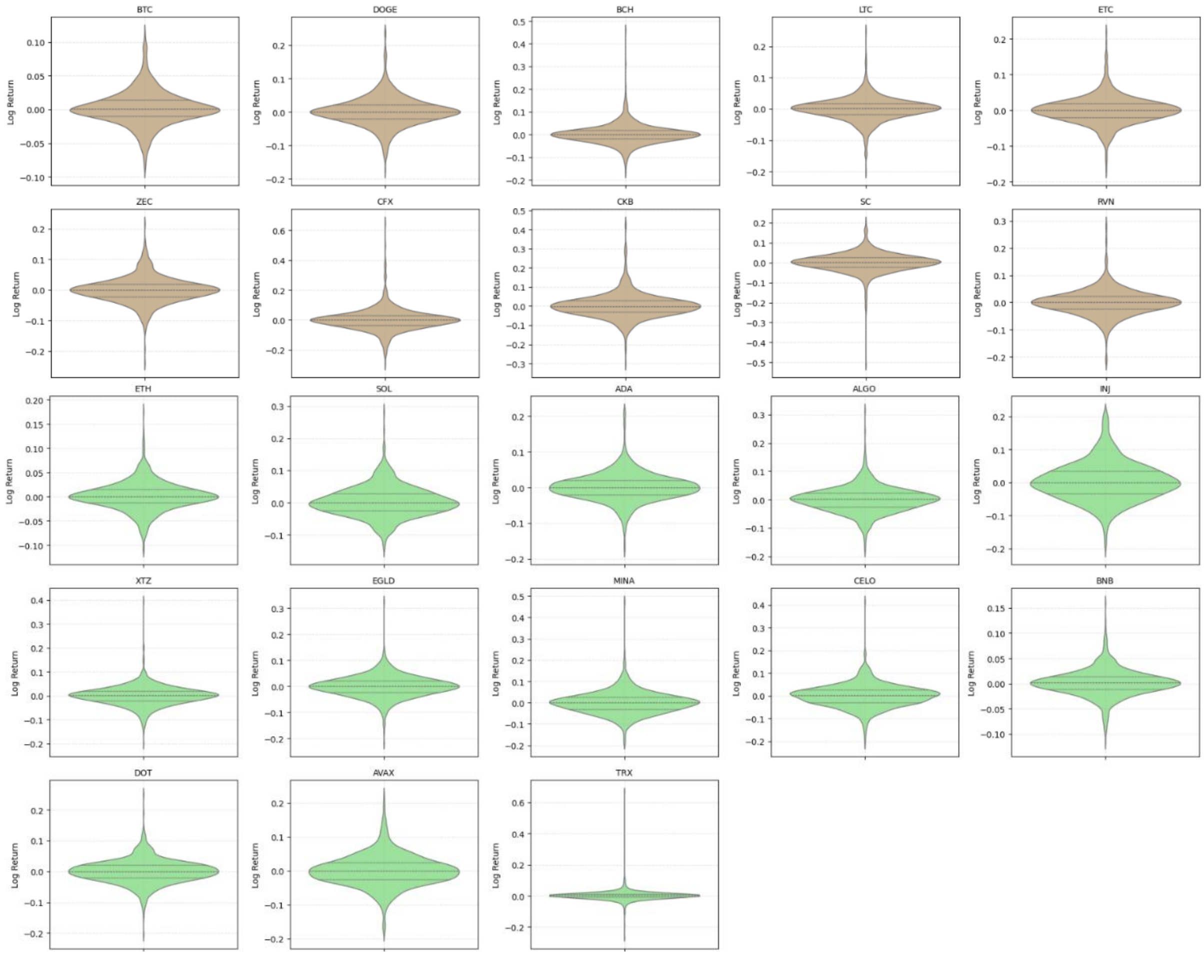

APPENDIX G : CRYPTOASSETS LOG RETURN ECDF PLOTS

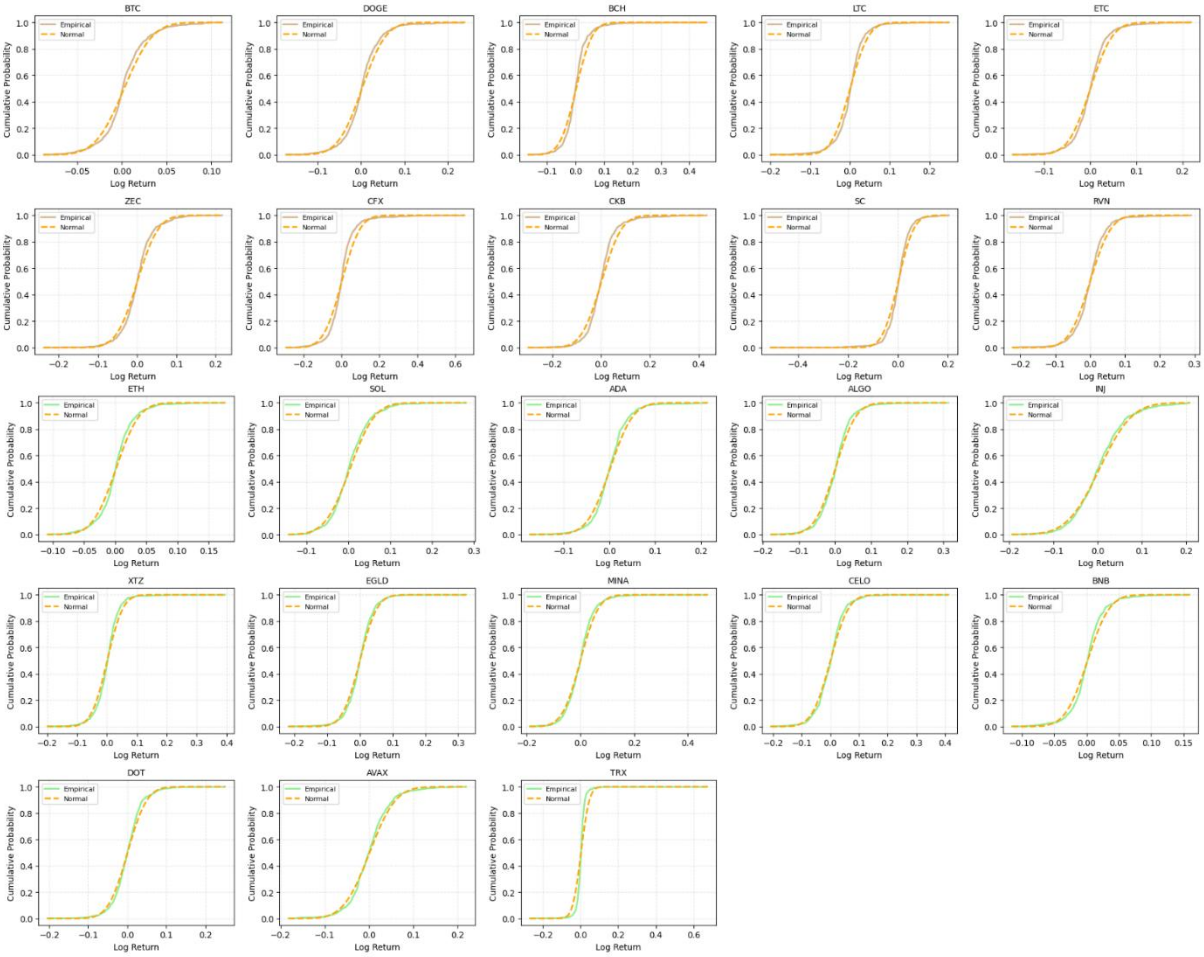

APPENDIX H : CENTRALITY AND CONNECTIVITY METRICS

Label	Closeness Centrality	Betweeness Centrality	PageRank	Eigenvector Centrality	Number of Triangles
ADA	1.00	0.00866	0.060	1.00	116
DOT	1.00	0.00866	0.060	1.00	116
ETC	1.00	0.00866	0.060	1.00	116
ZEC	1.00	0.00866	0.060	1.00	116
AVAX	1.00	0.00866	0.060	1.00	116
BNB	0.95	0.00450	0.057	0.96	108
ALGO	0.95	0.00698	0.057	0.95	104
RVN	0.95	0.00698	0.057	0.95	104
EGLD	0.95	0.00646	0.057	0.95	105
ETH	0.90	0.00334	0.054	0.91	96
LTC	0.90	0.00527	0.054	0.90	93
BCH	0.85	0.00219	0.051	0.86	85
BTC	0.81	0.00267	0.048	0.80	71
CELO	0.81	0.00144	0.048	0.81	74
INJ	0.81	0.00139	0.048	0.81	74
SOL	0.81	0.00139	0.048	0.81	74
DOGE	0.77	0.00067	0.045	0.75	64
XTZ	0.68	0.00000	0.035	0.58	36
CFX	0.00	0.00000	0.008	0.00	0
CKB	0.00	0.00000	0.008	0.00	0
MINA	0.00	0.00000	0.008	0.00	0
SC	0.00	0.00000	0.008	0.00	0
TRX	0.00	0.00000	0.008	0.00	0

Figure 12 - Centrality and Connectivity Metrics

APPENDIX I : DAILY LOG RETURNS DISTRIBUITION STATISTICS

	Group	Mean	StdDev	Skewness	Kurtosis
ADA	Green	0.001693	0.039442	0.580248	4.588149
ALGO	Green	0.000910	0.045393	0.659330	5.126302
AVAX	Green	0.001627	0.046390	0.355268	2.444030
BCH	Brown	0.002056	0.044866	2.266373	19.212308
BNB	Green	0.001436	0.027025	0.337452	4.121390
BTC	Brown	0.002374	0.025305	0.458755	2.350075
CELO	Green	0.000419	0.050404	0.722232	7.044552
CFX	Brown	0.002689	0.076423	1.857997	12.930823
CKB	Brown	0.002180	0.065241	1.331466	8.011921
DOGE	Brown	0.002061	0.043628	0.617133	4.158819
DOT	Green	0.000594	0.039134	0.343571	4.552668
EGLD	Green	0.000030	0.041792	0.293049	6.544521
ETC	Brown	0.000641	0.038107	0.527996	4.472493
ETH	Green	0.001406	0.029341	0.455669	3.570210
INJ	Green	0.003747	0.057856	0.438339	1.151966
LTC	Brown	0.000533	0.037057	0.086709	6.141321
MINA	Green	0.000387	0.053169	1.208909	9.445171
RVN	Brown	0.000117	0.047153	0.574024	6.141339
SC	Brown	0.001174	0.049312	-1.809194	18.426892
SOL	Green	0.004028	0.046748	0.690997	2.702255
TRX	Green	0.002113	0.035302	8.613868	185.414245
XTZ	Green	0.000802	0.041789	1.148201	13.178854
ZEC	Brown	0.000563	0.042429	0.092773	3.567077

Figure 13 - Distribuition Statistics

APPENDIX J : DAILY REALIZED VOLATILITY PLOTS

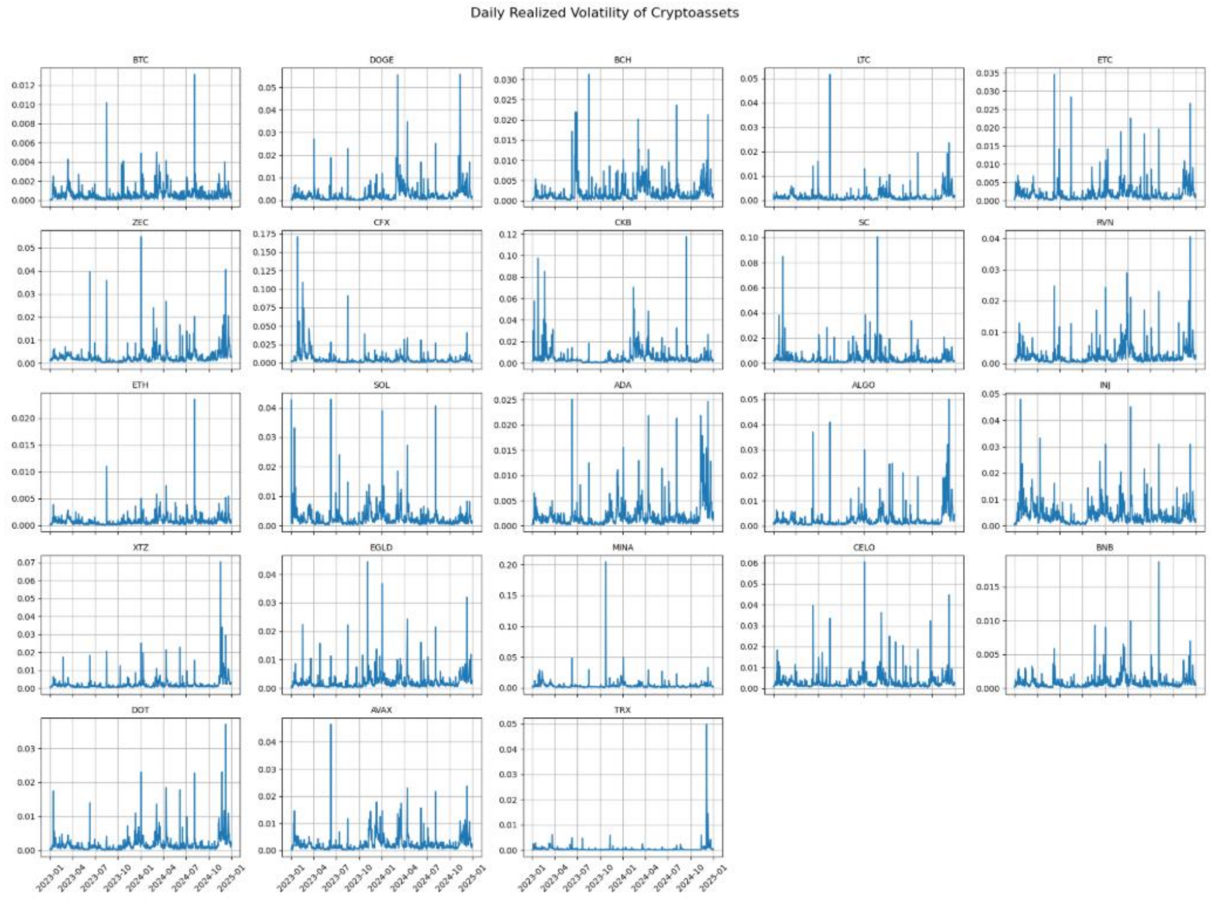

APPENDIX K : PREDICTIONS PLOTS

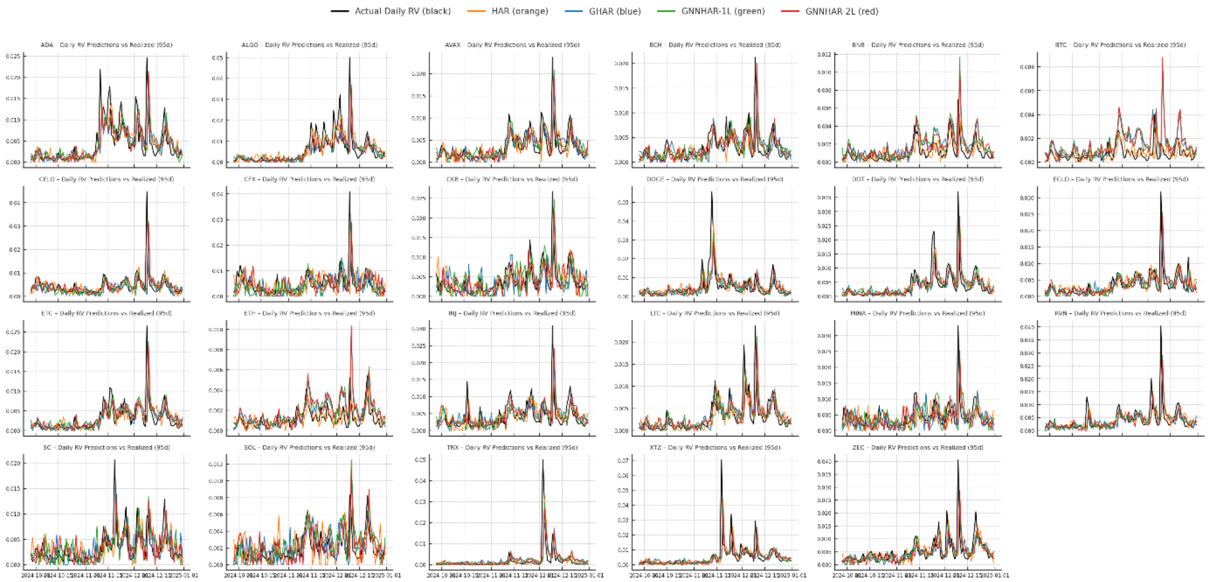

Figure 14 – Daily Realized Volatility Predictions

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação

Universidade Nova de Lisboa