



UNIVERSITÀ POLITECNICA DELLE MARCHE FACOLTÀ DI ECONOMIA “GIORGIO FUÀ”

Corso di Laurea Magistrale in Scienze Economiche e Finanziarie

A Network Approach to Cryptocurrency Portfolio Optimization

Un approccio di rete all'ottimizzazione di portafoglio nelle criptovalute

LM-16

Relatrice: Chiar.ma

Prof. Gloria Polinesi

Tesi di Laurea di:

Mototolea Daniel Emilian

Anno Accademico 2025– 2026

Summary

Summary	2
1. INTRODUCTION AND LITERATURE REVIEW	3
1.1. Basics: Modern Portfolio Theory	3
1.2. Critics	5
1.3. How the critics have been addressed	8
1.4. The choice of cryptocurrencies	15
2. DATA AND METHODS APPLIED	22
2.1. Data	22
2.2. Random Matrix Theory	25
2.3. Minimum Spanning Tree.....	27
2.4. Eigenvector Centrality.....	29
2.5. Alternative centrality measures	30
2.6. Network penalized portfolio optimization.....	33
2.7. Robustness and significance testing	34
3. RESULTS	36
3.1. Cryptocurrencies are single-factor	36
3.2. Residuality and MST	38
3.3. Portfolios performances	42
3.4. Robustness testing	46
4. CONCLUSIONS	50
5. BIBLIOGRAPHY	51

1. INTRODUCTION AND LITERATURE REVIEW

1.1. Basics: Modern Portfolio Theory

In 1952, Harry Markowitz laid the foundations for what is now called the Modern Portfolio Theory (MPT). First introduced in his seminal paper (Markowitz, 1952), and later expanded in his monograph (Markowitz, 1959), Markowitz's contribution shifted the paradigm of portfolio choice from qualitative, security by security assessment, to a formal mean variance optimization problem. In this framework, the expected return of a portfolio is computed as the weighted mean of the individual asset returns, while its risk is captured by the variance (or standard deviation) of the portfolio return. The key innovation was the explicit incorporation of the covariance structure among assets, making diversification both functional and measurable. The mean variance optimization produces in the (σ, μ) space an efficient set, also called the efficient frontier, that orders feasible allocations by their risk-return trade-off, enabling the investor to identify the portfolio that offers the highest expected return for a given level of risk.

Building on Markowitz's seminal contribution, a wide range of studies have developed this methodological approach. For instance, James Tobin, who derived the two-fund separation result, by introducing a risk-free asset in the mean-variance optimization framework (Tobin, 1958). The optimal allocation of portfolio resulting are combinations of this risk-free asset and one risky efficient portfolio, regardless of the

investor's risk preferences. Risk preference affects only the allocation between the safe and the risky fund, not the composition of the risky fund itself. Tobin's contribution provides a more parsimonious characterization of efficient choices under the existence of a safe asset. Since the composition of this risky portfolio is common across investors, Tobin's result effectively aggregates individual portfolio choice into a representative market demand. Under market clearing equilibrium assumptions, this separation logic leads to a pricing restriction: assets are priced according to how they co-move with the market portfolio. William Sharpe formalized an equilibrium model in which assets prices are determined assuming investors choose portfolios using the mean-variance criterion (Sharpe, 1964). Sharpe showed that the relevant risk for pricing is an asset's contribution to the risk of the market portfolio, rather than its idiosyncratic variance, which is not compensated because it can be diversified away in large portfolios. He derived a linear relationship between an asset's expected return and its market related risk exposure, also called *beta* or the market risk premium per unit of systematic exposure. In other words, assets more sensitive to market movements must offer higher expected returns. Beyond Sharpe's seminal formulation, the Capital Asset Pricing Model (CAPM) implication was independently obtained in two parallel equilibrium derivations (Lintner, 1965) and (Mossin, 1966), giving this literature contribution value in its robustness.

The above-mentioned framework is cast in a single-period setting, that is to say, the relevant risk-return trade-off is evaluated at one point in time and the opportunity set is treated as effectively fixed. Intertemporal portfolio theory makes the sequential nature of the problem explicit. In this regard, Samuelson's dynamic programming formulation implies that, in a multi-period setting, optimal portfolio choice should be described as

state and horizon dependent (Samuelson, 1969). Then Merton's continuous time framework further strengthens this intertemporal view by allowing the investment opportunity set to be stochastic: when expected returns and risk can vary through time, investors care not only about current trade-offs but also about hedging future changes in those opportunities (Merton, 1969), (1971) and (1973). The key step is to extend to continuous time with a stochastic investment opportunity set: when expected returns and risks vary through time, investors no longer care about current trade-offs but also about hedging future changes in those opportunities. This intertemporal perspective provides a natural rationale for the empirical use of rolling estimation and rebalancing in environments where dependence structures can shift rapidly: if the opportunity set is time varying, then relying on a single covariance estimate or a single period frontier is unlikely to provide a stable guide for allocation. Rolling windows pragmatically approximate the updating of beliefs about the underlying state variables, while periodic rebalancing operationalizes the dynamic adjustment of weights that the theoretical model prescribes. These methodological approaches are not only relevant from a theoretical perspective, but also constitute the core tools applied in the empirical analysis that follows.

1.2. Critics

The mean-variance framework provides the normative target for portfolio choice, as it formalizes diversification through the covariance matrix and defines an efficient frontier. Once the framework is taken to data though, the optimizer relies on plug-in estimates, the mean vector and the covariance matrix, which suffer in finite samples because the estimation noise can dominate the solution and destabilize weights. This

section reviews the theoretical and empirical literature that has identified and quantified this “estimation risk” problem.

Merton later warns about inference on the mean vector: estimating expected returns from realized returns alone is shown to be particularly fragile when the variance of market shifts across time (Merton, 1980). The core observation is simple: if one estimates the expected return as the sample average of T observations of a process with variance σ^2 , then the standard error of the estimate is $\sigma/\sqrt{T}$. Volatility is documented to vary strongly across subperiods, so naïve historical averaging can be an unstable estimator of the market return. The practical implication is that the mean vector needs a big amount of data, otherwise small perturbations can produce large swings in portfolio weights because the optimizer treats noisy estimates as if they were exact.

On the other hand, a distinctive contribution comes from Jobson and Korkie, proving how the covariance matrix is not immune. Even under standard distributional assumptions, the estimated frontier is a biased and poor proxy because it inherits noise from the sampling variation, that amplifies through inverting the covariance matrix, leading to unstable portfolio weights and poor out-of-sample performance (Jobson & Korkie, 1980). Since the optimizer seeks the minimum-variance portfolio for each target return, when the covariance matrix is inverted, noise in the small eigenvalues of the sample covariance matrix is amplified, because inversion maps small eigenvalues to large ones. The resulting portfolio weights are therefore dominated by the least precisely estimated eigenvectors, leading to extreme positions and instability. The estimated frontier appears to offer better risk-return trade-offs than are achievable in practice, or in other words suffers from in-sample overfitting; and the discrepancy between the estimated and the true frontier grows with the ratio N/T .

Furthermore, empirical analysis made by Bloomfield, Leftwich and Long reveal that the marginal benefits of adding assets to a portfolio diminish at much faster rate than previously assumed (Bloomfield et al., 1977). Consequently, this work suggests that for many investors, the transaction costs and estimation risks associated with maintaining a sophisticated rebalanced portfolio may outweigh the incremental returns, often making simpler strategies or index tracking more robust in the long run.

Along the same lines Frankfurter, Philips and Seagle via Monte Carlo Simulations conclude that portfolios selected via mean variance optimization are frequently no better, and sometimes worse, than a simple equally weighted portfolio. Again, stressing how sensitive the optimal weights, generated by the model, are to the input parameters (Frankfurter et al., 1971). Their experimental design generated asset returns from a known distribution, estimated the parameters from finite samples, solved the Mean Variance problem using the estimated parameters, and then evaluated the optimal portfolio using the true known distribution. The key insight was that the optimizer, in chasing small differences in estimated means, created extreme weights swings that were entirely a product of sampling noise. This early simulation study was among the first to stress how sensitive the optimal weights generated by the model are to the input parameters, foreshadowing the more formal analyses that would follow.

With respect to this issue, Richard Michaud famously coined the term '*estimation error maximizer*' referring to the Mean Variance optimizer. The optimization enigma is framed as a mismatch between theory and practice: unconstrained mean variance solutions can appear unintuitive and can perform poorly because the optimizer loads aggressively on assets whose estimated moments look favorable only due to noise, making in-sample optimality a weak guide to out-of-sample performance (Michaud,

1989). A later contribution extends this approach by benchmarking many allocation rules against the naïve portfolio. The study provides a stark requirement for sample length before mean-variance can beat $1/N$ portfolio (DeMiguel et al., 2009). None of the 14 optimization-based strategies consistently outperformed the naïve diversification, across 7 empirical datasets. A theoretical bound for minimum estimation window length was derived, to reliably beat $1/N$ diversification: the formula depends on the number of assets and the moments of the return distribution, and for typical parameters values for 25 assets, 250 years of monthly observations are required. This finding motivates the research design choice of using difficult universes (many assets, unstable connections) as stress tests, because they expose where plug-in optimization breaks first. The methodological message is that better dependence modelling and regularization are not merely incremental improvements, but conditions for feasibility of optimization-based allocation.

1.3. How the critics have been addressed

The critics reviewed in the previous section converge on a structural bottleneck: the covariance matrix estimated from finite samples is too noisy to serve as a reliable input to the optimizer, especially when the number of assets is large relative to the sample length. Two strategies have been proposed to address this problem. The first, drawn from statistical physics, uses Random Matrix Theory (RMT) to separate signal from noise in the eigenvalue spectrum of the correlation matrix. The second, drawn from network science, using the Minimum Spanning Tree (MST) to extract a parsimonious representation of the dependence structure. This section reviews the development of these approaches and their convergence into the framework adopted in this thesis.

A first response to estimation error comes from Random Matrix Theory. Early empirical applications showed that a substantial fraction of the eigenvalue spectrum of large financial correlation matrices is consistent with sampling noise rather than stable economic dependence (Laloux et al., 1998). The RMT literature compares the spectrum of empirical correlation matrices to theoretical benchmarks for random matrices, interpreting most eigenvalues in the bulk as noise, and only a few outliers as economically meaningful. The key portfolio implication is that the small eigenvectors in the bulk are the most noise sensitive.

Subsequent evidence extends this interpretation to both eigenvalues and eigenvectors, by systematically studying the properties of large financial correlation matrices, interpreting the largest eigenvalue as a collective market component and additional outliers as sectoral or group factors (Plerou et al., 2002). The central implication is that the bulk of the estimated correlation structure in a large equity universe is noise rather than signal, and that the blind use of empirical correlation matrices for risk management is problematic. A direct application to portfolio construction is replacing the noisy correlation matrix with a cleaned version that retains only the signal eigenvalues, producing portfolios with significantly better realized to predicted risk ratios.

Later work also shows that the signal-to-noise composition of financial correlation matrices depends on the observation horizon, with economically interpretable components becoming more visible at longer horizons, linking spectral structure to economic organization. At short horizons (minutes to hours), most eigenvalues fall between the bulk, reflecting a dominance of noise; at longer horizons (weeks to

months), more eigenvalues emerge as outliers, indicating that slower moving economic forces become visible (Potters et al., 2005).

The empirical approach provides a formal explanation to why plug-in mean variance procedures can exhibit systematic poor performance: by construction certain performance quantities are biased upward when N/T ratio is non negligible, namely in high dimensional data. Estimation noise distorts the lower part of the spectrum of $\hat{\Sigma}$, and matrix inversion magnifies errors associated with small eigenvalues, creating apparently attractive portfolio directions that may largely reflect estimation noise. The authors prove the mean variance procedure exhibits an overprediction bias in the sample-based estimate of the optimal return. (Bai et al., 2009). Denoting by R , the theoretical optimal return obtained from the population parameters and by $\hat{R}$ its sample-based counterpart:

$$\hat{R} \approx R / \sqrt{1 - y} \quad \text{for } y = N/T \in (0, 1)$$

which means that $\hat{R}$ is asymptotically inflated relative to its theoretical counterpart by a factor of $1/\sqrt{1 - y}$, where y is the ration of the dimension sample. Their work provides a practical remedy grounded in RMT through bootstrap corrected estimators for the optimal return and asset allocation that are consistent with their theoretical counterparts.

The study connects the Michaud “error maximizer” critique to the spectral properties of the sample covariance matrix, showing that the mechanism behind it admits a precise characterization in terms of the dimension-to-sample ratio. Estimation risk in the covariance component is thus, to a substantial extent, a random matrix phenomenon.

Spectral filtering addresses estimation noise, but it does not by itself provide a parsimonious representation of the dependence structure. Even after filtering, the

correlation matrix remains dense. This motivated a second strand of the literature, which represents financial dependence through correlation-based networks, by converting correlations into metric distances and applying graph-based filter, to extract the most informative connections.

In particular, a study introduced a correlation-based network representation in which pairwise correlations are transformed into distances and summarized through a the MST, providing a parsimonious representation that preserves the strongest links and supports the identification of the market's hierarchical organization (Mantegna, 1999). For a universe of N assets, the MST reduces the dense dependence structure to $N-1$ links while preserving representation of the market. The meaning of this is that branches that merge only at higher distances can be interpreted as groups driven by an idiosyncratic component, whereas early merging reflects stronger common factors. The key finding suggests that correlation-based filtering can reveal latent economic structure in return dependence.

The relevance of network filtering for portfolio construction is supported by evidence comparing different clustering techniques (these being single-linkage, average-linkage and PMFG) (Tola et al., 2008). The result being that clustering based covariance portfolios significantly improve the stability and out of sample performance of minimum-variance portfolios. The improvement is measured by comparing the ratio of realized risk to predicted risk: confirming the “error maximization” problem of Michaud (1989). The standard unfiltered mean variance portfolio construction produced realized to predicted risk ratios substantially above 1; whereas the clustering filtered estimators produce much closer to 1 risk predictions. This paper provides evidence that network-based filtering of the correlation matrix leads to better portfolio performance,

not just better network descriptions. The results are robust across different numbers of clusters, and the improvement is most pronounced in the high-dimensional regime where estimation error is most severe.

A dynamic extension of this approach constructs the MST over rolling windows, making Mantegna's taxonomy time-dependent, allowing to measure market structure and stability of its links; and its expansion or contraction (Onnela et al., 2003). The analysis characterizes the state of the tree through its normalized tree length, defined as the average length of its $N-1$ edges, together with the mean occupation layer, which measures the average topological distance of assets from a suitably defined central vertex. Since correlation-based distances decrease as correlations increase, a shorter normalized tree length corresponds to a more tightly connected market, whereas a longer tree is associated with greater diversification potential. The tree contracts markedly during market crises, when increasing market-wide correlations are accompanied by a sharp reduction in tree length. A complementary measure, edge survival across consecutive windows, drops abruptly around disruptive events, showing that dependence regimes reconfigure discontinuously rather than drifting smoothly. Although a large majority of connections survive between consecutive windows, pronounced drops in the single-step survival ratio indicate a substantial topological reorganization of the market during a crash. Thirdly, the authors relate the asset-tree topology to Markowitz portfolio selection. It was found that collective risk is concentrated in the core, motivating conditional allocation on evolving topology, or in other words, that portfolios composed of stocks from distant branches of the tree tend to be better diversified, with lower realized variance for a given expected return. This

provides evidence of a close relationship between peripheral network position and the diversification mechanism underlying minimum-variance portfolio selection.

To further support the topological properties of the MST built from return correlations we refer to evidence from equity markets, showing how different their structure is compared to randomly generated data (Bonanno et al., 2003). The authors build MSTs from a large set of stocks on the NYSE over a 12-year period and compare the resulting network structure with simulated data through purely random model (from a Wishart distribution) or single-factor model. The empirical MST shows a degree distribution with a fat tail: few highly connected hub nodes and many peripheral nodes with only one or two connections, a pattern entirely absent from the comparing models, which produced chain-like trees. The hub nodes correspond to major stocks, often sector leaders that serve as correlation intermediaries for their industry group. The analysis shows several topological measures and state that the empirical MST is significantly more star-like and heterogeneous than simulated MSTs. The implication for portfolio optimization is significant: a few hub nodes carry disproportionate network influence and their inclusion in a portfolio inflates variance through high-order interconnection effects that are not captured by the pairwise covariance matrix alone.

A direct link between spectral filtering and network topology is provided by comparisons of MSTs constructed before and after RMT filtering. By comparing the topological properties of the original MST, before applying RMT filtering, and the filtered MST, the study highlighted how its "cleaning" has a tangible impact on tree structure (EOM et al., 2009). The approach isolates the effect of noise on network structure. Their results show that the similarity between the estimated and original networks generally increases as additional eigenvalue components are included to the

filtered correlation matrix. However, this relationship varies according to the degree of the nodes in the original MST: stocks with different numbers of connections exhibit different responses to the progressive inclusion of spectral components, particularly around the eigenvalue interval predicted by RMT. When the largest eigenvalue was excluded, the resulting networks retained only 40-60% of the links found in the original network. More importantly, for the stocks with the highest degree, the survivor ratio fell to zero, indicating that none of their original connections were reproduced, meaning that as commonly interpreted as the market wide component, the largest eigenvalue is a major determinant of stock-network formation, although not sufficient on its own to explain the complete interaction among stocks. This evidence supports combining spectral cleaning with network extraction rather than treating the two procedures independently.

The literature establishes a formal theoretical bridge between network centrality and Markowitz portfolio optimization (Peralta & Zareei, 2016). The authors derive a negative relationship between eigenvector centrality and the optimal weight of an asset under the Markowitz framework. The explanation follows: securities that are strongly embedded in the correlation network greatly affect the market and their inclusion in a portfolio undermines the benefit of diversification, resulting in larger portfolio variance. Optimal portfolios must therefore overweight peripheral securities to avoid the variance inflation caused by including highly influential nodes. The authors use both in-sample and out-of-sample analysis, their work shows that this strategy outperforms standard benchmarks.

Building on the same insight, a separate study confirmed that asset allocation strategies incorporating the network structure of returns improve the portfolio's risk-

return profile relative to strategies that ignore it, thereby motivating the explicit use of network information at the optimization stage (Výrost et al., 2019).

This approach is further developed by embedding a network topology measure directly into the Markowitz objective function as a penalty term. Using both in sample and out of sample analysis, the authors showed that portfolios accounting for the network structure achieve better diversification and improved risk adjusted performance (Clemente et al., 2019).

These threads were synthesized into a single portfolio optimization framework by the work of Giudici, Polinesi and Spelta (2022). Their distinctive contributions are a parsimonious network derived from Random Matrix Theory filtering and an interconnectedness measure based on eigenvector centrality rather than the clustering coefficient. Their optimization problem augments the portfolio variance with a centrality penalty term γ . When γ equals 0, the problem reduces to a standard Markowitz on the RMT filtered covariance matrix. The parameter γ , can therefore be interpreted as an interconnectedness-aversion parameter: the larger its value, the stronger the penalty assigned to exposure to highly central assets and, consequently, to network-based systemic interconnectedness (Giudici et al., 2022).

It is this 3-stage pipeline, RMT filtering, MST construction and centrality penalized optimization, that the present thesis adopts and extends to the cryptocurrency domain.

1.4. The choice of cryptocurrencies

On 31 October 2008, the pseudonymous Satoshi Nakamoto published a white paper proposing a peer-to-peer electronic payment system that does not rely on a trusted third party (Nakamoto, 2008). Its core contribution solves the double spending problem through a decentralized network on nodes maintaining a shared, append only ledger, the

blockchain, in which each block is cryptographically linked to its predecessor and secured by a computationally expensive proof-of-work. The network went live on 3rd January 2009; the protocol caps at 21 million bitcoins, with the mining reward halving roughly every 4 years, so that scarcity is enforced by code rather than by sovereign authority. Following this model, thousands of alternative protocols have emerged, producing a heterogeneous asset class of over 17000 tokens, ranging from Bitcoin forks to smart-contracts platforms, DeFi tokens, and fiat-pegged stablecoins, whose combined market capitalization was, in November 2021, at its all-time high, at almost 3 trillion USD.

Whether cryptocurrencies satisfy the classical functions of money (W.S., 1875):

- medium of exchange;
- common measure of value;
- standard of deferred value;
- store of value;

remains contested. Their volatility undermines their use as a medium of exchange and unit of account, only stablecoins credibly serve as a standard of deferred payment, and even Bitcoin's status as a store of value is compromised by excessive drawdowns on several occasions. This thesis does not take a position in that debate. Rather than asking whether cryptocurrencies are money, it treats them pragmatically as one more investable asset class available in the financial toolkit and asks what a network-based allocation framework can extract from them. The empirical case for this stance is that cryptocurrencies behave statistically as a distinct asset class rather than as a currency, a point established by the asset pricing literature reviewed next.

The existing literature has documented that cryptocurrency returns have limited exposure to the factors commonly used to price traditional assets, including the market, size, value and momentum factors, as well as standard macroeconomic variables. Instead, returns are more closely related to crypto-specific variables such as network growth, investor attention and past performance, supporting the view that cryptocurrencies exhibit a distinct risk–return structure (Liu & Tsyvinski, 2018). This result was expanded by to the cross-section, showing with over 1800 cryptocurrencies that a three-factor structure, crypto market, size and momentum, captures cross-sectional expected returns, with these factors again distinct from their equity counterparts (Liu et al., 2022). Complementary evidence on the nature of this asset class comes from combining an analysis of return co-movements with direct evidence from the public Bitcoin ledger (Baur et al., 2018). Their findings are that Bitcoin returns are essentially uncorrelated with traditional asset classes in both calm and turbulent periods, implying diversification potential but ruling out a conventional safe-haven role, which would require negative correlation under stress; making it only a weak safe haven. Their transaction level analysis was more distinctive: approximately one third of Bitcoins were held by investors, particularly users who only received Bitcoin and never sent it onward and turnover rises with past returns and volatility; patterns characteristic of speculative holding rather than transactional use. The authors concluded that Bitcoin is used predominantly as a speculative investment rather than as a medium of exchange, reinforcing from the demand side the asset-pricing conclusion that cryptocurrencies are better understood as a distinct speculative asset class than as money.

Evidence from multivariate volatility models also indicates significant and time-varying return and volatility spillovers between Bitcoin and conventional financial

assets, while portfolio applications suggest potential diversification benefits from including Bitcoin (Guesmi et al., 2019). Empirically, the authors find that shocks to conventional financial assets, regardless of their sign, increase the conditional volatility of Bitcoin returns.

Complementary evidence is provided by applying the Diebold–Yilmaz (2012) generalized forecast-error variance decomposition to Bitcoin, Ripple and Litecoin alongside a range of conventional financial assets. The authors find that cryptocurrencies are strongly interconnected with one another while remaining relatively isolated from traditional markets on average (Corbet et al., 2018). However, both the intensity and direction of cross-market spillovers vary substantially over time, with temporary increases in connectedness occurring around major financial and cryptocurrency-specific events, including the Bitcoin flash crash, the Brexit referendum and the Bitfinex hack. Their frequency-domain analysis further indicates particularly weak connectedness between cryptocurrencies and conventional assets at short horizons, supporting potential short-horizon diversification benefits.

These findings should be interpreted alongside those of (Baur et al., 2018), who find that Bitcoin returns are essentially uncorrelated with traditional asset classes both under normal market conditions and during periods of financial turmoil. Rather than being contradictory, these results capture different dimensions of dependence. One examines contemporaneous return co-movements and their behaviour under extreme market conditions (2018), whereas the other use forecast-error variance decompositions to measure directional return and volatility spillovers (2018). Low contemporaneous correlation can therefore coexist with episodic and time-varying directional spillovers. Indeed, the first documents pronounced changes in connectedness around specific

events, while their DCC robustness analysis also identifies temporary increases in conditional correlations during some of these episodes. Taken together, the evidence suggests that the relationship between cryptocurrencies and traditional financial assets cannot be fully characterized by a single unconditional full-sample correlation estimate. This supports empirical approaches in which dependence structures are repeatedly re-estimated over time rather than assumed to remain stable throughout the sample.

The dependence structure within the crypto universe is the mirror image of its relation to traditional markets: whereas cryptocurrencies are relatively isolated from conventional assets, they are strongly interconnected with one another, and that internal structure is dominated by Bitcoin. The author quantifies this asymmetry directly through directional spillovers, Bitcoin price movements explain a large share of the forecast-error variance of Ripple and Litecoin, whereas the reverse influence of these altcoins on Bitcoin is small, identifying Bitcoin as the clear leader within the crypto market, particularly during periods of rapid price appreciation (2018).

The leadership is however strongly horizon dependent. Using ARDL models and bounds testing on daily data for Bitcoin and 16 altcoins over 2013-2016, the results suggest that in the short run fifteen of the sixteen altcoins are significantly affected by shocks to the Bitcoin price, a pervasive short-term dominance, yet in the long run only 4 altcoins remain cointegrated with Bitcoin (Ciaian et al., 2018); furthermore, over the long run, macro financial variables (the gold price, USD/EUR and CNY/USD exchange rates, and the 10 year treasury rate) determine altcoin prices to a slightly greater degree than Bitcoin itself. The crypto market is therefore not a monolithic bloc driven solely by Bitcoin, but a heterogenous ecosystem whose components are tied to Bitcoin at short horizons yet increasingly governed by independent drivers over the longer run. This

heterogeneity has a direct network counterpart. When correlation-based MSTs are constructed from cryptocurrency returns, Bitcoin and Ethereum (for the smart-contract ecosystem) typically emerge as a dominant hub, the analogue of sector leading stocks identified as hubs in equity MSTs. The degree of Bitcoin's centralization is not constant, however: it intensifies during speculative bull markets, when correlations across crypto assets rise, and relaxes during bear markets, when the topology becomes more distributed and secondary hubs emerge.

Cryptocurrency returns differ from traditional returns in ways that bear directly on the choice of optimization method. Their volatility is substantial: Bitcoin's annualized volatility has ranged from 50% to over 100%, and altcoin's far higher (Zdenka et al., 2019); their distributions are strongly leptokurtic, with tails substantially heavier than Gaussian, challenging modelling approaches that rely heavily on Gaussian distributional assumptions and motivating the use of methods designed to distinguish systematic dependence from sampling noise in high-dimensional settings, such as RMT-based spectral filtering.; they exhibit GARCH-type volatility clustering and asymmetric leverage effects, so that correlations rise during declines (Guesmi et al., 2019); they display strong short-horizon momentum (Liu et al., 2022); (Liu & Tsyvinski, 2018) implying a time-varying mean vector that must be re-estimated frequently; and they are subject to abrupt structural breaks, exchange failures, regulatory shocks, protocol forks, and deleveraging events such as the Terra/Luna collapse of 2022, that can rewrite the correlation structure overnight. The edge-survival analysis already showed that topological reconfigurations are abrupt rather than gradual (Onnela et al., 2003); in crypto they are more extreme still. Together these facts make the case for filtering, RMT to remove noise, MST to extract the dominant structure, and a centrality penalty

to limit exposure to systemic hubs, particularly compelling. The framework adopted in this thesis was in fact first applied to cryptocurrencies, in a smaller form, by Giudici, Pagnottoni and Polinesi (2020). Using daily returns of the ten largest cryptocurrencies over September 2017–October 2019, they applied the full RMT–MST–centrality pipeline and documented several crypto-specific features: the maximum link distance in the MST rose sharply during the 2017–2018 price hype, indicating a strongly misaligned correlation structure during bubbles, while the residuality coefficient fell as idiosyncratic variation collapsed and all tokens moved together, recovering after the bubble burst. On the portfolio side, their best-performing strategy during the bear phase was the RMT–MST-filtered Markowitz portfolio with the centrality penalty set to zero (their "Network Markowitz"), which more than halved losses relative to the equally weighted and the market benchmark; the authors were careful to note that they could not claim uniform outperformance over traditional approaches, since the advantage depended on the market phase and the risk-aversion parameter. The present thesis builds directly on this precursor. It applies the unified framework of Giudici, Polinesi and Spelta (2022) to a universe of 42 cryptocurrency assets over July 2021–September 2026, extending it across the major dislocations of the period (the Terra/Luna and FTX collapses).

The thesis is organised as follows. Chapter 1 has reviewed the literature, from the original mean-variance formulation and its estimation-risk critique through the random matrix and network approaches that address it. Chapter 2 sets out the data and the methodology, in the order in which the analysis proceeds: the construction of the cryptocurrency panel and the rolling-window design; the RMT filter used to separate systematic dependence from estimation noise; the MST that reduces the filtered

correlation matrix to a parsimonious network; the eigenvector centrality computed on that network, together with the alternative centrality measures used as benchmarks; the network-penalized optimization problem; and the resampling procedures used to assess significance and reliability. Chapter 3 presents the empirical results, beginning with the spectral and topological structure of the market, turning then to portfolio performance across values of the interconnectedness parameter and to the comparison with the CRIX benchmark, and closing with a set of robustness analyses. Chapter 4 concludes.

2. DATA AND METHODS APPLIED

This chapter describes the data and the methodological framework employed in the empirical analysis. The approach follows the following pipeline which integrates RMT, MST analysis and eigenvector centrality measures into a unified portfolio optimization framework. The methodology is applied here to a universe of cryptocurrency assets, where the statistical properties of the return distributions and the hub-dominated network topology make the approach particularly relevant.

2.1. Data

The dataset consists of daily closing prices for 42 cryptocurrency assets, denominated in US dollars, retrieved from publicly available exchange data. The assets are selected based on market capitalization and liquidity, requiring a minimum history of continuous trading over the sample period (07/2021-09/2026). The fixed universe requires continuous price availability over the sample period and is therefore conditional on asset survival. Each cryptocurrency is assigned a functional class, according to FTSE Grayscale Crypto Sector Index Series (Grayscale, 2024) reflecting

its primary use case within the digital asset ecosystem. This classification allows for the analysis of portfolio weight dynamics across functional segments of the market.

For each asset i , log-returns (r) are computed as:

$$r_i(t) = \log P_i(t) - \log P_i(t - 1) \quad (1)$$

where $P_i(t)$ denotes the closing price of asset i at time t . The use of logarithmic returns ensures additivity over time and symmetry in the treatment of gains and losses.




Figure 2.1 Distribution of cryptocurrencies returns grouped by classes.




Figure 2.2 Daily cryptocurrencies log-return, grouped by classes and pooled.

The boxplots in figure 2.2 provide a complementary view of the distribution of daily log-returns across crypto sectors, compared to the QQ plots in figure 2.1. Returns are generally centered around zero, with relatively narrow interquartile ranges compared with the overall dispersion of the observations. However, all sectors exhibit many observations beyond the whiskers, on both the positive and negative sides, highlighting the frequent occurrence of extreme returns. The dispersion and magnitude of these observations vary across sectors, with Smart Contract Platforms and Financials displaying particularly pronounced extremes. The fat tails are denoted by the S shape in the QQ plots.

The data is organized into a sequence of overlapping rolling windows. Each window has a total width of 250 trading days, partitioned into an in-sample (estimation) period 250 days and an out-of-sample (validation) period of 7 days. The windows advance with a step size of 7 trading days (one trading week), so that consecutive windows overlap by 250 days. For each window, the in-sample observations are used to estimate the correlation matrix C , apply the RMT filter, construct the MST, compute

centrality scores and solve the portfolio optimization problem. The resulting optimal weights are then applied to the out-of-sample period to compute realized portfolio returns.

This rolling-window design serves two purposes. First, it ensures a strict temporal separation between estimation and evaluation: no future information is used in the construction of the portfolio weights, thereby preventing look-ahead bias. Second, it allows the dependence structure among asset returns to evolve over time, capturing the dynamic nature of cryptocurrency correlations and network topology.

2.2. Random Matrix Theory

Random Matrix Theory (RMT) provides a principled method for separating the informative signal in an empirical correlation matrix from the estimation noise that arises when the number of assets N is large relative to the number of time-series observations T .

The approach compares the eigenvalue spectrum of the empirical correlation matrix with the theoretical spectrum that would be observed if the returns were drawn from independent, identically distributed random variables.

For each rolling window, the $N \times N$ sample correlation matrix C is computed from the in-sample log-returns using the Pearson correlation coefficient, where c_{ij} denotes the (i,j) -th element of the correlation matrix:

$$c_{ij} = [E(r_i r_j) - E(r_i)E(r_j)] / (\sigma_i \sigma_j) \quad (2)$$

while $E(\cdot)$ and $\sigma(\cdot)$ denote the expected value and standard deviation operators, respectively.

The eigenvalues of C are then compared with the Marchenko-Pastur distribution (Marcenko & Pastur, 1967). As $N \rightarrow \infty$ and $T \rightarrow \infty$ with fixed ratio $Q = T/N \geq 1$, the density of the eigenvalues of a random Wishart matrix converges to:

$$f(x) = (Q / 2\sigma^2 \pi) \sqrt{[(\lambda^+ - \lambda)(\lambda - \lambda^-)]} / \lambda \quad (3)$$

for $x \in (\lambda^-, \lambda^+)$, where the theoretical bounds are:

$$\lambda_{\pm} = \sigma^2 (1 + 1/Q \pm 2\sqrt{1/Q}) = \sigma^2 (1 \pm 1/\sqrt{Q})^2 \quad (4)$$

Eigenvalues that fall within the interval $[\lambda^-, \lambda^+]$ are consistent with the null hypothesis of random noise and cannot be statistically distinguished from sampling noise. Eigenvalues that exceed the upper bound λ^+ are interpreted as signal: they correspond to systematic factors that drive the co-movement among asset returns. It is important to note that the Marchenko–Pastur distribution is used here as a null benchmark for sampling noise. Although its standard theoretical formulation relies on assumptions such as independent observations and finite moments, it provides a useful reference against which the empirical eigenvalue spectrum can be assessed. Specifically, the null hypothesis assumes that the observed returns are independent, and that the resulting correlation matrix behaves as a random Wishart matrix; deviations from the corresponding Marchenko–Pastur bounds therefore provide evidence of non-random correlation structure. (Plerou et al., 2002).

The bound in (4) assumes $\sigma^2 = 1$, that is, that the entire matrix is noise. This is appropriate as a null benchmark, but it is not the correct scale against which to compare the residual spectrum: the leading eigenvalue absorbs a large share of the total variance, leaving the remaining $N-1$ directions with substantially less than unit variance each. On the present sample λ_1 accounts for between 47% and 72% of the trace (66% on average), and applying (4) with $\sigma^2 = 1$ leaves a single eigenvalue above threshold in every

window. The noise variance is therefore estimated from the residual spectrum as $\hat{\sigma}^2 = 1 - \lambda_1/N$ (Laloux et al., 1998), and the threshold becomes $\lambda^+ = \hat{\sigma}^2(1 + 1/\sqrt{Q})^2$. With $T = 250$ and $N = 42$ this lowers the bound from 1.99 to approximately 0.68 and retains between four and eight eigenvalues, which is the informative part of the spectrum this framework is designed to isolate.

The filtered correlation matrix C^* is obtained as

$$C^* = V \Lambda^* V^T \quad (5)$$

where V is the matrix whose columns are the eigenvectors, while Λ^* is the diagonal matrix that retains λ_k for eigenvalues above the threshold and replaces each eigenvalue below it with the average of the noise eigenvalues, $\delta = (N - \sum \lambda_k)/(N - k)$, with k the number of retained eigenvalues. This choice, eigenvalue clipping, preserves the trace of the correlation matrix and leaves C^* of full rank. Setting the noise eigenvalues to zero, as in the simplest formulation of the filter, would instead reduce the trace and yield a matrix of rank k ; on this sample, where $k = 1$ under the unadjusted threshold, the reconstructed matrix would be of unit rank and the induced distance matrix would collapse the MST into a degenerate star. Clipping is also the procedure that emerges from minimizing the deviation between the empirical spectrum and its random-matrix counterpart, which to first order amounts to replacing each group of eigenvalues by a single degenerate eigenvalue equal to the group mean (Laloux et al., 2000).

2.3. Minimum Spanning Tree

The filtered correlation matrix C^* provides a denoised representation of the pairwise dependence structure among all N assets, but with $N(N-1)/2$ distinct entries it remains difficult to interpret directly. To obtain a more parsimonious representation, the

correlation matrix is first converted into a distance matrix and then reduced to a tree structure via the MST algorithm (Mantegna, 1999; Bonanno et al., 2003).

The distance between any two assets i and j is defined as

$$d_{ij} = \sqrt{(2 - 2c_{ij}^*)} \quad (6)$$

where c_{ij} is the (i, j) element of the filtered correlation matrix C . Building the network from C^* rather than from the raw correlation matrix is possible because the clipping procedure of Section 2.2 preserves the full rank of the matrix: had the noise eigenvalues been set to zero, C^* would have been of rank k , and, under the unadjusted threshold, of unit rank, producing a degenerate distance matrix in which every asset is equidistant from every other, and a MST reduced to a star. Filtering the matrix from which the network is built is consistent with the evidence in that RMT filtering alters the topological properties of the resulting tree (EOM et al., 2009).

The metric mentioned above (6), satisfies the axioms of a Euclidean distance and maps correlations into non negative distances bounded between 0 and 2. The set of all pairwise distances forms the $N \times N$ distance matrix D . The MST is obtained by applying a single-linkage clustering algorithm (Kruskal, 1956) to the distance matrix D . The algorithm proceeds iteratively: at each step, the pair of clusters (initially, individual assets) with the smallest inter-cluster distance is merged, subject to the constraint that no cycles are formed. The process terminates when all N nodes are connected in a single tree with exactly $N-1$ edges. The resulting MST retains the most important connections while reducing the number of pairwise links from $N(N-1)/2$ to $N-1$. Single-linkage clustering is known to suffer from chaining, whereby elongated clusters absorb nearby points and the resulting hierarchy can be unstable (Carlsson & Memoli, 2010; Marti et al., 2021). The author cautions that alternative linkage criteria yield different grouping

structures without materially affecting portfolio performance, and the MST is retained here for comparability with the framework being replicated (Raffinot, 2017).

To monitor how the network topology evolves over time, the *residuality coefficient* (Res) is employed, following (Spelta & Araújo, 2012):

$$R = \frac{\sum_{d_{ij} > L} d_{ij}^{-1}}{\sum_{d_{ij} \leq L} d_{ij}^{-1}} \quad (7)$$

where d_{ij} is the distance between assets i and j as defined in (6), L is the maximum edge weight (highest distance) in the MST. The coefficient compares the aggregate strength of the connections lying outside the tree with those retained within it, and thus measures how much of the market's dependence structure a single spanning tree fails to capture. A low value of Res indicates a dependence structure that is compact and close to single-factor, well summarized by a single spanning tree. A high value, conversely, signals a richer, multidimensional structure that the tree cannot compress, one in which a substantial share of the connection strength lies outside it. Importantly, Res reflects the shape of the dependence rather than its overall strength: a market may be highly correlated yet still well described by a single tree, or only moderately correlated yet require more than one factor to be represented.

2.4. Eigenvector Centrality

The MST identifies the topology of the dependence network, but for portfolio construction it is necessary to quantify how central or peripheral each asset is within this structure. The eigenvector centrality measure (Bonacich, 1972) assigns a score to each node based on the principle that connections to highly scored nodes contribute more than connections to low-scored nodes.

Formally, for the i -th node, the centrality score x_i is defined by the recursive equation:

$$x_i = (1/\lambda) \sum_j \hat{A}_{ij} x_j \quad (8)$$

where $\hat{A}_{ij}$ is the (i, j) element of the adjacency matrix of the MST and λ is a constant. In matrix notation, this becomes the eigenvalue problem:

$$\hat{A} x = \lambda x \quad (9)$$

where x is the vector of node centralities.

There exist, in general, multiple eigenvalues λ for which a solution exists. However, the Perron-Frobenius theorem guarantees that for a connected network with non-negative entries, there exists a unique eigenvector with all positive components corresponding to the largest eigenvalue. This eigenvector provides the desired centrality scores.

In the context of cryptocurrency correlation networks, a high centrality score indicates that an asset is strongly connected to other highly connected assets, that is, it sits at a hub of the dependence structure. The literature generally established the negative relationship between eigenvector centrality and optimal Markowitz weights under the long-only constraint: assets that are central in the correlation network contribute disproportionately to portfolio variance through the dominant eigenvalue of the covariance matrix, and the optimizer therefore assigns them lower weights (Peralta & Zareei, 2016). Taken together, these results place the eigenvector-centrality penalty within a consistent body of evidence, from equities (Pozzi et al., 2013) to multi-asset portfolios (Giudici et al., 2022), showing that tilting away from central, highly connected assets is a robust route to diversification, an expectation that chapter 3 tests directly.

2.5. Alternative centrality measures

Eigenvector centrality is adopted as the baseline interconnectedness measure, in keeping with Giudici, Polinesi and Spelta (2022). The choice of centrality measure,

however, is not innocuous: different centrality measures encode different notions of how "central" a node is, and it is therefore worth examining how the strategy behaves when eigenvector centrality is replaced by each of them in turn. To assess this, the eigenvector centrality is replaced, in turn, by five alternative measures, grouped here into three conceptual families.

The first family comprises proximity-based measures, which rank a node by how close it is, on average, to all others in the tree. Closeness centrality (Freeman, 1977) is defined as the reciprocal of the sum of the shortest-path distances from a node to every other node:

$$Cl_i = (N - 1) / \sum_{j \neq i} d(i, j) \quad (10)$$

where Cl_i denotes the closeness centrality of node i , N is the total number of nodes in the network, and $d(i, j)$ is the length of the shortest path connecting nodes i and j . The normalization by $N-1$ makes the measure comparable across networks of different sizes. So, a node lying near the center of the tree, with short paths to the rest of the market, obtains a high score.

Harmonic centrality (Rochat, 2009) is a variant that sums the reciprocals of the distances rather than reciprocating their sum:

$$H_i = \sum_{j \neq i} 1 / d(i, j) \quad (11)$$

where H_i denotes the harmonic centrality of node i , and $d(i, j)$ is the shortest-path distance between nodes i and j ; which makes it well defined even on disconnected graphs and less sensitive to a small number of very distant nodes.

The second family is spectral, sharing with eigenvector centrality the recursive principle that a node is important if it is connected to important nodes. PageRank (Brin & Page, 1998) refines this idea through a damped random walk on the network: the

centrality of a node is the stationary probability that a random walker, who follows edges with probability α and jumps to a random node with probability $1 - \alpha$, is found at that node. The damping factor (here $\alpha = 0.85$) bounds the influence of any single dominant hub, yielding a smoother, more gradual ranking than the raw leading eigenvector.

Lastly, sparse measures assign appreciable scores only to a small number of well-placed nodes, though for different structural reasons. Betweenness centrality (Freeman, 1977) counts the fraction of shortest paths between all node pairs that pass through a given node,

$$B_i = \sum_{s \neq i \neq t} \sigma_{st}(i) / \sigma_{st} \quad (12)$$

where σ_{st} is the number of shortest paths from node s to node t and $\sigma_{st}(i)$ the number of those passing through i . On a tree, where the path between any two nodes is unique, betweenness reduces to counting the pairs separated by node i , and therefore, leaf nodes have zero betweenness, whereas internal nodes receive positive scores according to how many parts of the tree they connect.

The expected force (Lawyer, 2015), originally proposed as a measure of the epidemic influence of a node, quantifies the expected size and diversity of the clusters generated by a two-step diffusion process starting at that node, computed as the entropy of the normalized cluster forces. Like betweenness, it assigns appreciable values only to a small number of well-placed nodes. This measure was introduced into the portfolio setting by Výrost, Lyócsa and Baumöhl, from which the present comparison borrows it (2019).

This taxonomy anticipates the empirical finding of Chapter 3: the proximity- and spectral-based measures, which assign a graded score to every node, produce a

continuous and monotone peripheral tilt as γ increases, whereas the path-based measures, structurally sparse on a tree, being null on all leaf nodes, provide no continuous gradient and prove ineffective as penalty terms. This ineffectiveness is not specific to the present sample: in a systematic comparison of centrality measures for asset selection, the authors report that betweenness performs poorly precisely because a large share of nodes, nearly half in their networks, receive a betweenness score of zero, offering no basis for discrimination (Wang et al., 2024).

2.6. Network penalized portfolio optimization

The 3 preceding steps, RMT filtering, MST construction, eigenvector centrality computation, are synthesized in a single portfolio optimization framework that extends the classical Markowitz (1952) model by embedding network information directly into the objective function. In its classical form, the mean-variance problem selects weights that minimize portfolio variance for a given level of expected return, in this context though, Σ^* is introduced as the RMT-filtered covariance matrix, computed as the element-wise product of the outer product of asset standard deviations and the filtered correlation matrix ($\text{COV}^*_{ij} = \sigma_i \sigma_j c^*_{ij}$), and γ is a risk-tolerance parameter. The Giudici, Polinesi and Spelta framework (2022) extends this formulation by adding a centrality penalty term to the objective function:

$$\min^w w^T \Sigma^* w + \gamma \sum_i x_i w_i \quad (13)$$

subject to:

$$\sum_i w_i = 1, \quad \mu^p \geq \bar{\mu}, \quad w_i \geq 0 \text{ for all } i \quad (14)$$

where x_i is the eigenvector centrality of asset i computed from the MST; μ^p is the expected portfolio return and $\bar{\mu}$ the cross-sectional mean of individual expected returns,

so that the optimiser is required to do at least as well as an equally weighted allocation; the non-negativity constraint on w_i rules out short selling; and γ is the *interconnectedness aversion* parameter that governs the trade-off between first-order risk (portfolio variance) and high-order systemic risk (network centrality). The Markowitz model is expanded without modifying its underlying quadratic formula, through the term γ . When $\gamma=0$, the network-centrality penalty disappears and the model reduces to a Markowitz-type optimization based on the RMT-filtered covariance matrix.

2.7. Robustness and significance testing

The empirical results of the following chapter are subjected to two resampling procedures, addressing two distinct questions: whether the observed performance differences are statistically meaningful rather than artefacts of sampling variability, and whether the RMT filter genuinely improves the reliability of risk estimation.

The first procedure examines whether the risk-adjusted performance of the leading strategy differs from that of the market benchmark (the CRIX index¹). The object of the test is the pair of daily out-of-sample return series, of the strategy and that of the CRIX, over the full evaluation period.

Both series violate the assumptions of the classical Jobson and Korkie (1981) test of Sharpe-ratio equality, though in different ways: the benchmark exhibits the volatility clustering typical of financial returns, while the strategy series is serially dependent in the raw returns. The difference is therefore assessed with the studentised bootstrap test of Ledoit and Wolf (2008), built on a stationary block bootstrap that resamples blocks of consecutive days and so preserves dependence of either form. The mean block length is selected by the data-driven procedure of Politis and White (Politis & White, 2004),

¹ Roylton, 2020. *ROYALTON CRIX CRYPTO INDEX*.
Available at: <https://www.royalton-crix.com/>

with the correction of Patton, Politis and White, applied to the squared returns (Patton et al., 2009); this yields an optimal length of approximately 6 days. Robustness to this choice is verified over block lengths from two to ten days. The p-value is two-sided, and a 95% confidence interval for the annualized difference is obtained by inverting the studentized test.

The second procedure assesses the extent to which RMT filtering improves the estimation of portfolio risk, and how this improvement depends on the ratio $Q = T/N$ between the length of the estimation window and the number of assets. This ratio governs the regime in which estimation error dominates: as Q approaches one, the sample covariance matrix becomes increasingly noisy and unstable, and filtering is expected to yield the greatest benefit (DeMiguel et al., 2009) (Bai et al., 2009). Motivated by the risk-prediction reliability framework of Tola, Lillo, Gallegati and Mantegna, the empirical sample is treated as the data-generating process (2008): the in-sample returns are resampled in blocks, and for each replication the minimum-variance portfolio is constructed twice, once from the RMT-filtered covariance matrix and once from the sample covariance matrix. For each portfolio, the ratio between the realised out-of-sample risk and the risk predicted in sample is recorded; a ratio close to one indicates a reliable risk estimate, while a ratio above one signals that the strategy underestimates the true risk. By varying the length of the estimation window, the procedure traces this reliability ratio across a range of values of Q , isolating the regime in which the filtering advantage is largest. This design connects the empirical analysis directly to the estimation-risk theory reviewed in Chapter 1 and to the Marchenko–Pastur bounds introduced in Section 2.2.

3. RESULTS

All results are reported over the out-of-sample period from August 2021 to September 2026, with weekly rebalancing on rolling windows of 250 in-sample days over a universe of 42 cryptocurrencies.

3.1. Cryptocurrencies are single-factor

Throughout, the guiding finding is that the cryptocurrency market is, structurally, close to a single-factor system, essentially driven by the market mode, and this fact shapes both what the network approach can extract and how the strategy performs. Across the 230 estimation windows the leading eigenvalue of the correlation matrix accounts for between 50% and 75% of total variance, as shown in figure 3.1, a share far above what is observed in equity markets, where the market mode typically absorbs less.



Figure 3.1 Share of variance for the first eigenvalue in the estimation period.

This dominance has a direct methodological consequence. The Marchenko–Pastur bound in Equation (4) is derived under $\sigma^2 = 1$, that is, on the assumption that the entire matrix is noise. When one eigenvalue absorbs two thirds of the trace, the

remaining $N-1$ directions carry substantially less than unit variance each and comparing them against a bound calibrated for unit variance is the wrong test: applied unadjusted, it leaves a single eigenvalue above threshold in every window of the sample; just as in figure 3.2. Rescaling the bound by the residual variance, $\hat{\sigma}^2 = 1 - \lambda_1/N$, lowers to 0.68, and retains 4-8 eigenvalues.




Figure 3.2 First eigenvalue in the top time series, following eigenvalues in the second time series, and number of eigenvalues across the evaluation sample, in both cases with scaled λ_+ bound and unscaled.




Figure 3.3 Eigenvalue density for the evaluation period 09/2023-06/2024.

3.2. Residuality and MST

The reconstruction step matters equally. Setting the sub-threshold eigenvalues to zero, as in the simplest formulation of the filter, does not preserve the trace and returns a matrix of rank k . Under the unadjusted bound, where $k = 1$, the reconstructed matrix would be of unit rank: every pairwise distance would collapse to the same value and the MST would degenerate into a star, with no topology left to measure. Replacing them instead by their common mean, preserves the trace, leaves the matrix of full rank, and yields a tree whose maximum degree ranges from 5 to 17.




Figure 3.4 Residuality plot over time identifying distinct regimes not visible from correlation.

The residuality coefficient Res provides a topological counterpart to the spectral picture: across the sample Res and average pairwise correlation are only weakly related ($\rho = -0.21$), so that the two quantities describe different aspects of the dependence structure.

The contrast between the first half of 2023 and the first half of 2024 makes the point. In the earlier period average pairwise correlation is high, at 0.72, while residuality is at its lowest, 0.20: a highly synchronized market in which a parsimonious network backbone captures most of the dependence. A year later correlation has fallen to 0.54, the lowest of the sample, with a minimum in March 2024, yet residuality rises to 0.26. Lower market-wide co-movement therefore coexists with a larger share of dependence lying outside the tree. The reason is that Res measures how pairwise dependence is distributed relative to the connectivity threshold of the MST, not how strong it is on average: a market can be tightly coupled and still well summarized by a single spanning tree, or only moderately coupled and yet require more structure than a tree can hold.

The extremes bracket the same reading. Res reaches its maximum of 1.03 in the window ending May 2022, in which the Terra/Luna collapse falls at the end of the estimation period, and its minimum of 0.054 in December 2022, a phase of near-total synchronization, with average correlation at its sample high of 0.75, in which almost all of the dependence is absorbed by the tree. The sample closes with the lowest residuality of the period, indicating a market whose structure has become progressively more compressible into a single backbone.




Figure 3.5 The MST estimated on 12 November 2022 – 19 July 2023.




Figure 3.6 The MST estimated on 21 October 2023 – 26 June 2024.

Figure 3.5 and 3.6 shows representative MSTs at 2 different points in the examined period. After building the tree from the filtered correlation matrix, the trees exhibit genuine branching structure hubs, chains and sectoral sub-clusters, rather than the degenerate star that the filtered matrix would produce. Across the 230 rebalancing windows the highest-degree node rotates among mid-capitalisation assets, most frequently Polkadot, Flow and VeChain, with no single asset dominating. Bitcoin and Ethereum, despite their market capitalization, seldom occupy the center of the tree. The interpretation is that the MST identifies the assets that co-move most tightly with the rest of the market rather than the largest ones: over 2021–2026, Bitcoin and Ethereum

retain partially idiosyncratic dynamics that place them closer to the periphery than mid cap assets tracking the broad altcoin market.

3.3. Portfolios performances

Figure 3.7 reports the cumulative log-return of the eigenvector-centrality strategy for nine values of γ , alongside the CRIX index and the equally weighted portfolio. The ordering in γ is clear: the RMT-filtered Markowitz portfolio without centrality penalty ($\gamma = 0$) is at the ending of the period at a cumulative log-return above the CRIX and well above the equally weighted portfolio. A small tilt ($\gamma = 0.005$) remains close behind, while all larger values end the period negative. The degradation is continuous and substantial over the range examined, first up to ($\gamma = 0.02$) the cumulative log return is positive, then all the strategies end below zero at the end of the evaluation period.

The deterioration is accompanied by a progressive loss of diversification: higher values of γ raise portfolio concentration and reduce the number of assets receiving economically meaningful weights. The penalty ceases to act as a marginal diversification adjustment and increasingly constrains the optimizer toward a narrow

subset of peripheral assets.




Figure 3.7 Cumulative out-of-sample log-returns of the RMT-network portfolios for different values of γ , compared with the equally weighted portfolio and the CRIX benchmark.

The same ordering appears in the drawdown profile: $\gamma = 0$ and the CRIX show the shallowest drawdowns and the strongest recoveries, whereas the high- γ strategies remain pinned near -80% for most of the period, the direct consequence of concentration in a handful of peripheral, low-capitalisation assets whose idiosyncratic crashes are not diversified away.




Figure 3.8 Drawdown of the RMT-network portfolios for different values of γ , compared with the equally weighted portfolio and the CRIX benchmark.




Figure 3.9 Weights per asset for RMT-network portfolio with $\gamma=0$ and $\gamma=0.005$ to show concentration through rebalancing.

The results of the peripheral tilt to the choice of centrality measure is assessed by replacing eigenvector centrality, in turn, with closeness, harmonic, PageRank, betweenness and expected force, across the same grid of γ . Three distinct groups emerge.

The first group comprises eigenvector, closeness and harmonic centrality. These measures are strongly cross-correlated and produce an almost identical tilt: as γ increases, portfolio concentration rises smoothly and the three yield near-indistinguishable performance. Within this group, the peripheral tilt remains broadly consistent across alternative centrality measures, supporting the use of eigenvector centrality as a representative baseline.

PageRank forms a group of its own. Although it too modulates the portfolio continuously, it ranks assets very differently from the first group and produces a milder tilt. The reason is structural: on an undirected graph, the stationary distribution of a random walk is proportional to node degree (Newman, 2010). On the sparse structure of an MST, degree is a rough signal that ranks assets by their number of direct connections rather than by their global position, which accounts for both PageRank's low correlation with the proximity- and eigenvector-based measures and its gentler tilt. It is also the most damaging at high γ .

The third group comprises betweenness and expected force. Both are structurally sparse on a tree, non-zero only at the few internal nodes and null on every leaf, and therefore provide no continuous gradient for the optimizer: the concentration index remains essentially flat across all values of γ , and the corresponding strategies do not respond to the penalty. These measures act as an on/off switch rather than a continuous tilt. This ineffectiveness is not specific to the present sample: likewise find betweenness (and eccentricity) unsuitable for network-based portfolio optimization (Wang et al., 2024).

No measure at any value of γ exceeds the untilted portfolio, and none significantly outperforms the CRIX.

3.4. Robustness testing

While $\gamma = 0$ outperforms the CRIX in cumulative terms, cumulative return does not account for risk. Table 3 reports the annualized Sharpe ratio of each strategy alongside that of the CRIX, with the p-value of the Sharpe-ratio difference from the block bootstrap of Section 2.7.

γ	Sharpe	Sharpe CRIX	Difference	p-value	95% CI
0.	0.316	0.140	0.176	0.535	[-0.395, +0.742]
0.005	0.272	0.140	0.132	0.664	[-0.504, +0.738]
0.05	-0.358	0.140	-0.498	0.133	[-1.178, +0.160]
0.15	-0.371	0.140	-0.510	0.147	[-1.214, +0.176]
0.25	-0.435	0.140	-0.575	0.109	[-1.312, +0.105]
0.7	-0.551	0.140	-0.691	0.068	[-1.456, +0.032]
1.	-0.541	0.140	-0.681	0.077	[-1.468, +0.061]
2.	-0.540	0.140	-0.679	0.086	[-1.473, +0.070]
4.	-0.588	0.140	-0.728	0.071	[-1.531, +0.046]

Table 1 Bootstrap comparison of annualized Sharpe ratios between the eigenvector-centrality portfolio specifications and the CRIX benchmark.

The untilted portfolio achieves the highest Sharpe ratio, 0.316 against 0.140 for the CRIX, but the difference is not statistically distinguishable from zero. The confidence intervals make the reason plain: for $\gamma = 0$ the interval spans $[-0.395, +0.742]$, wide

enough to accommodate both moderate underperformance and substantial outperformance. A four-year daily sample does not deliver the power to separate the two hypotheses. The intervals do shift systematically with γ , centred above zero at $\gamma = 0$ and progressively below as the penalty grows, which reinforces the qualitative ordering even where individual tests are inconclusive. The deterioration is continuous and substantial, though not strictly monotone at the highest values, where the differences fall within the noise.

Across all 54 specifications (9 gamma values and 6 centrality measures), five reach significance at the 5% level. All five are PageRank at $\gamma \geq 0.25$, and all five are significantly *worse* than the benchmark.

Recentering the bootstrap distribution under the null is an approximation, since the distribution of the Sharpe-ratio difference is asymmetric. The studentised statistic of Ledoit and Wolf (2008) is asymptotically pivotal and enjoys faster-vanishing coverage error, at the cost of requiring a standard error for every replication. It is therefore applied to the leading specification alone, $\gamma = 0$ against the CRIX, using an inner bootstrap of 100 replications within each of 1,000 outer replications. Both series are taken gross of transaction costs here, since the CRIX is a market index that is not actively rebalanced. On this basis the strategy attains an annualised Sharpe ratio of 0.426 against 0.140 for the benchmark, an excess of 0.286 with a standard error of 0.27. The studentised test gives $t = 1.05$ and $p = 0.300$ at a block length of three days, with a 95% confidence interval of $[-0.219, +0.837]$. Repeating it for block lengths of two, five and ten leaves the conclusion unchanged, with p between 0.217 and 0.306 and intervals straddling zero throughout. The refinement lowers the p -value appreciably relative to the unstudentised procedure but does not alter the inference.

The mean block length is selected by the procedure of Politis and White (Politis & White, 2004) , with the correction of Patton, Politis and White (Patton et al., 2009), applied to the excess return series, which is the object of the test. This yields approximately six days. The value largely reflects the rebalancing design rather than a property of the market: portfolio weights are held fixed over each seven-day out-of-sample block, so the concatenated series is serially dependent by construction, a Ljung–Box test rejects the absence of autocorrelation in the strategy returns themselves at every lag from 5 to 30 ($p < 10^{-3}$), whereas the benchmark shows the more familiar pattern of uncorrelated returns and strongly autocorrelated squares ($p < 10^{-20}$). The estimate is nonetheless the relevant one, the purpose of the resampling scheme being to preserve whatever dependence is present in the series under test. Tests are reported for block lengths of two to ten days, bracketing the data-driven value.

The second robustness exercise assesses how much RMT filtering improves the reliability of risk estimation, and how this depends on the ratio $Q = T/N$ between the estimation-window length and the number of assets. Following the risk-prediction reliability criterion of (Tola et al., 2008), an out-of-sample experiment is designed and adapted to the rolling cryptocurrency setting: unlike their original design, which varies N and T over random portfolios and random start times, the in-sample returns are block-resampled within each rolling window, and for each resample the minimum-variance portfolio is constructed twice, from the RMT-filtered and from the sample covariance matrix. The reliability of each estimator is measured by the ratio between realized out-of-sample risk and predicted in-sample risk; a ratio of one denotes a perfectly reliable estimate, while a ratio above one signals that risk is underestimated.




Figure 3.1 Reliability of risk prediction across different values of $Q=T/N$.

Figure 3.9 reports this ratio as a function of Q , for estimation windows from $T = 45$ ($Q \approx 1.1$) to $T = 350$ ($Q \approx 8.3$). Two regularities emerge. First, the ratio declines monotonically in Q and crosses unity between $Q \approx 2.7$ and $Q \approx 3.6$: below that point in-sample optimisation underestimates the risk realised out of sample, by as much as 14% at the shortest windows, while above it risk is increasingly overestimated, reaching 0.70 at $T = 350$. The window adopted throughout, $T = 250$, corresponds to $Q \approx 6.0$ and a ratio of 0.86, and is therefore conservative in its risk assessment. Second the filtered covariance yields a ratio closer to one at every value of Q , and its advantage is concentrated in the low- Q regime: the gap between the two methods is largest, roughly 0.025, for Q between 1.0 and 1.7, where estimation noise dominates, and narrows to essentially zero as Q rises above three, where the sample covariance becomes reliable on its own. The RMT filter thus acts as a partial substitute for data abundance, improving the reliability of risk estimation precisely when the estimation window is

short relative to the number of assets. The magnitude of the improvement, while systematic and theoretically coherent, is nonetheless modest, and is reported as such.

4. CONCLUSIONS

This thesis has applied the network-based portfolio framework of Giudici, Polinesi and Spelta (2022) to 42 cryptocurrencies over July 2021 to September 2026, with weekly rebalancing on rolling windows of 250 in-sample days. Three findings emerge, and they share one explanation.

The market is close to a single-factor system. The leading eigenvalue absorbs between 47% and 72% of total variance, well above what is observed in equity markets. Beneath it the residual structure is thin and unstable: once the threshold is rescaled for the variance the market mode absorbs, four to eight eigenvalues survive the Marchenko–Pastur bound, but only the market mode keeps a stable composition across windows. The one exception is a nine-window episode spanning March 2024 to January 2025, in which a second factor separates payment- and enterprise-oriented tokens from the rest of the market.

The Network Markowitz portfolio achieves an annualized Sharpe ratio of 0.43 against 0.14 for the CRIX, but a studentized block bootstrap gives $p = 0.30$, with a confidence interval of $[-0.22, +0.84]$. Nor is it what its construction suggests: it holds two assets in effective terms, and imposing weight caps, while beneficial out of sample in the equity literature, degrades performance here. The mechanism is specific to the sample: a cap forces the optimizer to hold assets it would otherwise exclude, and over this period the broad altcoin universe lost heavily, as the equally weighted portfolio makes clear. Whether the concentrated portfolio reflects skill or the realized performance of two assets cannot be settled on four years of data.

The three findings are consequences of the same fact. Both components operate on the dependence structure that remains once the market mode is removed: the filter cleans it, the penalty exploits its topology. Where that structure is thin and unstable, as here, and as it is not in the equity markets for which the framework was designed, there is little for either to work with. The framework does not fail because it is misapplied; it produces no gain because the conditions for one are largely absent. This identifies a structural precondition for network-based portfolio methods, and a diagnostic for testing it before applying them: the share of variance carried by the leading eigenvalue, and the stability of the eigenvectors beneath it.

5. BIBLIOGRAPHY

Bai, Z., Liu, H. & Wing, K., 2009. ENHANCEMENT OF THE APPLICABILITY OF MARKOWITZ'S. *Mathematical Finance*.

Bai, Z., Liu, H. & Wong, W.-K., 2009. Enhancement of the Applicability of Markowitz's Portfolio Optimization by Utilizing Random Matrix Theory. *Mathematical Finance*.

Baur, D., Hong, K. & Lee, A., 2018. Bitcoin: medium of exchange or speculative assets?. *Journal of International Financial Markets, Institutions and Money*.

Bloomfield, T., Leftwich, R. & Long, J. J., 1977. Portfolio strategies and performance. *Journal of Financial Economics*.

Bonacich, P., 1972. Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology*.

Bonanno, G., Caldarelli, G., Lillo, F. & Mantegna, R., 2003. Topology of correlation based minimal spanning trees in real and model markets. *Physical Review E*.

- Bonanno, G., Caldarelli, G., Lillo, F. & Mantegna, R., 2003. Topology of correlation-based minimal spanning trees in real and model markets. *Physical Review E*.
- Brin, S. & Page, L., 1998. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*.
- Carlsson, G. & Memoli, F., 2010. Characterization, stability and convergence of hierarchical clustering method. *Journal of machine learning research*.
- Ciaian, P., Rajcaniova, M. & Kancs, d., 2018. Virtual relationships: Short- and long-run evidence from BitCoin and altcoin markets. *Journal of International Financial Markets, Institutions and Money*.
- Clemente, G. P., Grassi, R. & Hitaj, A., 2019. Asset allocation: new evidence through network approaches. *Annals of Operations Research*.
- Corbet, S. et al., 2018. Exploring the dynamic relationships between cryptocurrencies and other financial assets. *Economics Letters*.
- DeMiguel, V., Garlappi, L. & Uppal, R., 2009. Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?. *The Review of Financial Studies*.
- Diebold, F. & Yilmaz, K., 2012. Better to give than to receive: predictive directional measurement of volatility spillovers. *International Journal of Forecasting*.
- EOM, C. et al., 2009. Topological properties of stock network based on minimal spanning tree and random matrix theory in financial time series. *Physica*.
- Frankfurter, G. M., Phillips, H. E. & Seagle, J. P., 1971. Portfolio Selection: The Effects of Uncertain Means, Variances, and Covariances. *Journal of Financial and Quantitative Analysis*.
- Freeman, L., 1977. A set of measures of centrality based on betweenness. *Sociometry*.

Giudici, P., Pagnottoni, P. & Polinesi, G., 2020. Network models to enhance automated cryptocurrency portfolio management. *Frontiers in Artificial Intelligence*.

Giudici, P., Polinesi, G. & Spelta, A., 2022. Network models to improve robot advisory portfolios. *Journal of Financial Markets*.

Grayscale, F., 2024. *FTSE Grayscale Crypto Sector Index Series*. [Online] Available at: https://www.lseg.com/content/dam/ftse-russell/en_us/documents/ground-rules/ftse-grayscale-crypto-sector-index-series-ground-rules.pdf

[Accessed 10 02 2026].

Guesmi, K., Saadi, S., Abid, I. & Ftiti, Z., 2019. Portfolio diversification with virtual currency: evidence from Bitcoin. *International Review of Financial Analysis*.

Jobson, J. & Korkie, B., 1980. Estimation of Markowitz Efficient Portfolios. *Journal of the American Statistical Association*.

Jobson, J. & Korkie, B. M., 1981. Performance hypothesis testing with the Sharpe and Treynor measures. *Journal of Finance*.

K., G., S., S., I., A. & Z., F., 2019. Portfolio diversification with virtual currency:evidence from Bitcoin. *International review of Financial Analysis*.

Kruskal, J. B. J., 1956. On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem. *Proceedings of the American Mathematical Society*.

Laloux, L., Cizeau, P., Bouchaud, J. & Potters, M., 1998. Noise Dressing of Financial Correlation Matrices. *Physical Review Letters*.

Lawyer, G., 2015. Understanding the influence of all nodes in a network. *Scientific Reports*.

Ledoit, O. & Wolf, M., 2008. Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance*.

Ledoit, O. & Wolf, M., 2015. Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance*.

Lintner, J., 1965. The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *Review of Economics and Statistics*.

Liu, Y. & Tsyvinski, A., 2018. Risks and returns of cryptocurrency. *National Bureau of economic research*.

Liu, Y., Tsyvinski, A. & Wu, J., 2022. Common risk factors in cryptocurrency. *The Journal of Finance*.

Mantegna, R. N., 1999. Hierarchical Structure in Financial Markets. *The European Physical Journal*.

Marcenko, V. & Pastur, L., 1967. DISTRIBUTION OF EIGENVALUES FOR SOME SETS OF RANDOM MATRICES. *American Mathematical Society*.

Marčenko, V. & Pastur, L., 1967. *Distribution of eigenvalues for some sets of random matrices*. s.l.:Mathematics of the USSR.

Markowitz, H., 1952. Portfolio Selection. *The Journal of Finance*.

Markowitz, H., 1959. *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley.

Marti, G., Nielsen, F., B. M. & P., D., 2021. A review of two decades of correlations, hierarchies, networks and clustering. *Information Geometry: Theory and Applications*.

Merton, R. C., 1969. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, p. 247–257.

Merton, R. C., 1971. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory*.

Merton, R. C., 1973. An Intertemporal Capital Asset Pricing Model. *Econometrica*.

- Merton, R. C., 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics*, pp. 323-362.
- Michaud, R. O., 1989. The Markowitz Optimization Enigma: Is 'Optimized' Optimal?. *Financial Analysts Journal*.
- Mossin, J., 1966. Equilibrium in a Capital Asset Market. *Econometrica*.
- Nakamoto, S., 2008. *Bitcoin: A Peer-to-Peer Electronic Cash System*, s.l.: s.n.
- Newman, M. E. J., 2010. *Networks: An Introduction*. Oxford: Oxford University Press.
- Onnela, J.-P. et al., 2003. Dynamics of market correlations: Taxonomy and portfolio analysis. *Phys. Rev.*
- Patton, A. J., Politis, D. N. & White, H., 2009. Correction to 'Automatic Block-Length Selection for the Dependent Bootstrap. *Econometric Reviews*.
- Peralta, G. & Zareei, A., 2016. A network approach to portfolio selection. *Journal of empirical finance*.
- Peralta, G. & Zareei, A., 2016. A network approach to portfolio selection. *Journal of empirical finance*.
- Plerou, V. et al., 2002. A Random Matrix Approach to Cross-Correlations in Financial Data. *Physical Review*.
- Politis, D. N. & Romano, J. P., 1994. The stationary bootstrap. *Journal of the American Statistical Association*.
- Politis, D. N. & White, H., 2004. Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*.
- Politis, D. & Romano, J., 1994. The stationary bootstrap. *Journal of the American Statistical Association*.

Potters, M., Bouchaud, J.-P. & Laloux, L., 2005. Financial Applications of Random Matrix Theory: Old Laces and New Pieces. *Acta Physica Polonica*.

Pozzi, F., Di Matteo, T. & Aste, T., 2013. Spread of risk across financial markets: better to invest in the peripheries. *Scientific Reports*.

Raffinot, T., 2017. Hierarchical clustering-based asset allocation. *The Journal of Portfolio Management*.

Rochat, Y., 2009. Closeness centrality extended to unconnected graphs: The harmonic centrality index. *Applications of Social Network Analysis*.
[Accessed 10 02 2026].

S., C., A., M. & L., Y., 2018. Exploring the dynamic relationships between cryptocurrencies and other financial assets. *Economic Letters*.

Samuelson, P. A., 1969. Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, p. 239–246.

Sharpe, W. F., 1964. Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *The Journal of Finance*.

Spelta, A. & Araújo, T., 2012. The topology of cross-border exposures: Beyond the minimal spanning tree. *Physica A*.

Tobin, J., 1958. Liquidity Preference as Behavior Toward Risk. *The Review of Economic Studies*, pp. 65-86.

Tola, V., Lillo, F., Gallegati, M. & Mantegna, R. N., 2008. Cluster Analysis for Portfolio Optimization. *Journal of Economic Dynamics and Control*.

Tola, V., Lillo, F., Gallegat, M. & Mantegna, R., 2008. Cluster Analysis for Portfolio Optimizatio. *Journal of Economic Dynamics and Control*.

Tumminello, M., Mantegna, R. N., Aste, T. & Matteo, T. D., 2005. Correlation filtering in financial time series. *Noise and Fluctuations in Econophysics and Finance*.

Výrost, T., Lyócsa, Š. & Baumöhl, E., 2019. Network-based asset allocation strategies. *The North American Journal of Economics and Finance*.

W.S., J., 1875. Money and the mechanism of exchange. *D. Appleton and Company*.

Wang, G.-J., Huai, H., Zhu, Y. & Xie, C. & U. G. S., 2024. Portfolio optimization based on network centralities: Which centrality is better for asset selection during global crises?. *Journal of Management Science and Engineering*.

Wang, G.-J.et al., 2024. Portfolio optimization based on network centralities: Which centrality is better for asset selection during global crises?. *Journal of Management Science and Engineering*.

Zdenka, K., Nicole, S. & Cynthia, M., 2019. Bitcoin, Litecoin, and the Euro: an annualized volatility analysis. *Studies in Economics and Finance*.