

Entropy-regularized portfolio optimization and financial network entropy: diversification and concentration in a multi-asset ETF portfolio

Dilawar Ahmad Bhat^{a,*}, Sajad Ahmad Sheikh^b, Manzoor Ahmad Khanday^c

^a*Symbiosis School of Banking and Finance, Symbiosis International (Deemed) University, Pune, India*

^b*Department of Mathematics, University of Kashmir (South Campus), Anantnag, India*

^c*Department of Statistics, Lovely Professional University, Punjab, India*

Abstract

This study examines whether Shannon entropy regularization can mitigate the concentration associated with classical minimum-variance portfolio optimization while preserving risk-adjusted performance. Using a twelve-ETF global multi-asset universe covering equities, fixed income, precious metals, and real estate, we compare an equal-weight (EW) benchmark, a long-only fully invested minimum-variance (MV) portfolio, and an entropy-regularized minimum-variance (ERMV) portfolio. All portfolio aggregation, wealth compounding, and performance ratios use simple returns, with the risk-free rate expressed on the same basis. The MV optimizer allocates 61.87% to intermediate U.S. Treasury bonds (IEF) and 38.13% to high-yield corporate bonds (HYG), with effectively zero weight on the remaining assets. At $\lambda = 0.01$, the ERMV solution is close to equal weighting, with all twelve weights between 0.0828 and 0.0841. A sensitivity analysis over $\lambda \in [10^{-5}, 1]$ shows a systematic transition from concentration toward the maximum-entropy equal-weight allocation. Full-sample Sharpe ratios under a 4.5% annual risk-free rate are 0.2979 (EW) and 0.2966 (ERMV), and a circular block-bootstrap test indicates no statistically distinguishable difference. A rolling out-of-sample evaluation yields the same qualitative pattern. Correlation-based minimum spanning tree (MST) analysis reveals a core-periphery structure, while rolling spectral entropy identifies the March 2020 COVID-19 sell-off as the period of most concentrated cross-asset dependence. The findings position entropy regularization primarily as a concentration-control mechanism rather than a universal source of superior investment performance.

DOI: [10.46481/jnsps.2027.3778](https://doi.org/10.46481/jnsps.2027.3778)

Keywords: Portfolio diversification, Shannon entropy, Minimum spanning tree, Financial networks, Exchange-traded funds

Article History:

Received: 02 August 2026

Received in revised form: 29 September 2026

Accepted for publication: 29 September 2026

Available online: 09 October 2026

© 2027 The Author(s). Published by the [Nigerian Society of Physical Sciences](#) under the terms of the [Creative Commons Attribution 4.0 International license](#). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Communicated by: B. J. Falaye

1. Introduction

Exchange-traded funds (ETFs) have become an important vehicle through which investors obtain diversified exposure to broad asset classes, geographical markets, and investment styles while retaining the liquidity and transparency associated

with exchange-traded securities [1, 2]. Their portfolio structure also makes ETFs particularly useful for empirical studies of cross-asset dependence: rather than focusing on individual securities, a relatively small set of ETFs can provide exposure to economically distinct segments of the financial system. This feature has contributed to growing interest in the statistical dependence, efficiency, liquidity, and risk-management properties of ETF markets [3–9].

A persistent difficulty in portfolio construction is that sta-

*Corresponding Author Tel. No.: +91-206-193-6283.

Email address: dilawar.bhat@ssbf.edu.in (Dilawar Ahmad

Bhat )

tistical risk minimization need not coincide with economically broad diversification. The classical minimum-variance (MV) prescription of Markowitz [10] minimizes the variance of portfolio returns subject to investment constraints, but the resulting solution can be highly concentrated when the optimizer exploits differences in estimated variances and covariances across assets [11–14]. Such concentration is not necessarily a computational defect: conditional on the estimated covariance matrix and the imposed constraints, a concentrated portfolio may be the mathematically optimal solution. The concern is instead that the statistical optimum can differ substantially from an economically diversified allocation and can be sensitive to estimation error in the covariance matrix.

The concentration tendency of variance-based optimization has motivated alternative portfolio-construction approaches that incorporate explicit diversification preferences. Among these, entropy-based methods are particularly attractive because Shannon entropy provides a direct measure of the dispersion of portfolio weights. For a long-only portfolio, entropy is maximized by the equal-weight allocation and decreases as portfolio mass becomes concentrated in fewer assets. Incorporating an entropy term into the minimum-variance objective therefore provides a continuous mechanism through which the optimizer can trade statistical variance reduction against dispersion of capital across assets [15–20]. The resulting entropy-regularized minimum-variance (ERMV) portfolio does not replace the covariance structure; rather, it modifies the optimization criterion so that extreme concentration carries an explicit penalty.

The magnitude of this regularization is consequential. Because portfolio variance and Shannon entropy are expressed on different numerical scales, the regularization parameter cannot be interpreted independently of the scale of the variance objective and the return data. A single value of λ may therefore generate a solution that is only weakly different from classical minimum variance in one application but very close to equal weighting in another. For this reason, examining the ERMV solution at one arbitrarily selected value of λ is insufficient to establish the role of entropy regularization. The present study consequently considers the regularization path over a broad range of λ values and evaluates how portfolio concentration, entropy, effective number of assets, volatility, and risk-adjusted performance change as the strength of the entropy penalty increases.

A second, complementary perspective is provided by financial-network analysis. Financial markets can be represented as networks in which assets constitute nodes and statistical dependence between their returns determines the links between them [21–27]. Correlation-based distance measures and minimum spanning trees (MSTs) provide a parsimonious representation of the dominant dependence structure while avoiding the visual and analytical complexity of a fully connected network. Previous research has used MSTs and related network measures to characterize market organization, asset centrality, and changes in financial connectedness during periods of market stress [21–23, 28–32]. Such representations are useful in the present context because portfolio diversification depends not only on the individual volatility of assets but also on the

way in which their returns co-move.

Network entropy provides an additional way of characterizing changes in this dependence structure. In particular, entropy calculated from the eigenvalue spectrum of a correlation matrix summarizes how evenly the system's variance is distributed across its principal components. A reduction in spectral entropy can therefore be interpreted as evidence of a more concentrated dependence structure, or a lower effective dimensionality of common variation. This interpretation is descriptive rather than prescriptive: network entropy does not itself constitute a portfolio optimization criterion. Rather, it provides a complementary measure with which to examine whether the dependence structure underlying the portfolio problem changes during periods of market stress.

The empirical setting of this study is a twelve-ETF universe covering U.S. large-cap, technology, and small-cap equities; emerging and developed international equities; gold and silver; U.S. Treasury securities; corporate bonds; and U.S. real estate. Daily adjusted closing prices are used to construct daily *simple* returns, and all portfolio aggregation, wealth compounding, and risk-adjusted performance ratios are computed on this consistent return basis (see Sections 3.2 and 4.10). Three portfolio strategies are examined: an equal-weight (EW) benchmark, a long-only fully-invested classical minimum-variance portfolio, and an entropy-regularized minimum-variance portfolio. The portfolio analysis is complemented by a correlation-based MST and a rolling spectral-entropy measure calculated from the same return data.

The empirical analysis addresses four related questions. First, how does a long-only, fully-invested minimum-variance optimizer allocate capital across the heterogeneous ETF universe, and does the resulting allocation exhibit substantial concentration? Second, to what extent does Shannon entropy regularization alter that concentration, and how does the answer vary with the regularization parameter λ ? Third, does entropy regularization preserve risk-adjusted performance relative to the equal-weight benchmark, both in the full sample and under a rolling out-of-sample evaluation, when returns and the risk-free rate are expressed on the same basis? Fourth, how does the dependence structure of the ETF universe evolve through time, and does a spectral measure of network entropy identify the March 2020 market-stress episode as a period of unusually concentrated common variation?

These questions are deliberately separated because portfolio optimization and network analysis address different objects. The portfolio procedures determine capital allocations from the covariance structure and portfolio constraints, whereas the MST and spectral-entropy analysis characterizes the structure and evolution of cross-asset dependence. We therefore do not interpret the network measures as additional portfolio constraints, nor do we claim that peripheral network position mechanically produces lower portfolio variance. Instead, the network analysis provides a complementary description of the dependence environment within which diversification decisions are made.

The empirical results provide several useful findings. The corrected long-only minimum-variance solution is highly con-

centrated in intermediate U.S. Treasury and high-yield corporate bond ETFs. This concentration substantially reduces portfolio volatility but is accompanied by relatively weak realized return and negative excess-return risk-adjusted performance under the 4.5% risk-free-rate specification. Entropy regularization largely eliminates this concentration: at the principal regularization value considered in the analysis, the ERMV weights are very close to equal weighting, while the λ -sensitivity analysis reveals a smooth transition from the concentrated minimum-variance solution toward the equal-weight allocation as the entropy penalty increases. Importantly, the evidence does not support interpreting ERMV as a strategy that systematically improves Sharpe performance over equal weighting. Rather, the full-sample and rolling out-of-sample results indicate that the diversification induced by entropy regularization can be obtained while preserving performance that is very close to the equal-weight benchmark.

The network analysis yields a complementary result. The full-sample correlation structure exhibits distinct equity/real-estate, fixed-income, and precious-metal groupings, while the MST provides a parsimonious representation of the dominant cross-asset dependence structure. More importantly, the rolling spectral entropy reaches its sample minimum during the acute phase of the March 2020 market disruption. This behavior is consistent with a temporary concentration of common variation and a reduction in the effective dimensionality of the correlation structure. Because the portfolio and network analyses are conducted separately, this finding is interpreted as evidence concerning the time variation of the market dependence environment rather than as evidence that the portfolios themselves are network-optimized.

The contribution of the paper is therefore primarily methodological and empirical. Methodologically, it combines a constrained entropy-regularized portfolio objective with complementary network characterization within a common ETF return universe, using a fully consistent simple-return measurement basis throughout. Empirically, it documents how the concentration of a classical minimum-variance solution changes when an explicit entropy preference is introduced, traces the resulting solution across the regularization parameter, and evaluates the portfolio comparison both in sample and through a rolling out-of-sample framework. The network component further illustrates that the dependence structure underlying these diversification decisions is not static and can undergo substantial contraction during periods of market stress. Taken together, the results support a more cautious interpretation of entropy regularization: its principal benefit in this setting is not demonstrable superior performance, but the ability to substantially reduce concentration while retaining risk characteristics close to a simple and difficult-to-beat equal-weight benchmark.

The remainder of the paper is organized as follows. Section 2 reviews the literature. Section 3 describes the ETF universe, data construction, and return measurement. Section 4 presents the portfolio optimization procedures, entropy measures, MST construction, rolling network-entropy methodology, and performance evaluation framework. Section 5 reports the empirical results. Section 6 discusses robustness and sensi-

tivity analyses, Section 7 discusses the findings, and Section 8 concludes.

2. Literature review

2.1. Portfolio diversification and the limits of minimum-variance optimization

The modern portfolio framework established by Markowitz [10] treats portfolio risk as a function of the covariance structure of asset returns and formalizes diversification as a reduction in portfolio variance through the combination of assets whose returns are not perfectly correlated. The minimum-variance (MV) portfolio is the limiting case in which the investor places primary emphasis on minimizing estimated return variance subject to the relevant portfolio constraints. Although the theoretical objective is straightforward, the empirical solution can differ markedly from an economically diversified portfolio because the optimizer exploits the entire estimated covariance structure rather than imposing a direct preference for dispersion across assets.

This distinction becomes particularly important under long-only constraints. When short selling is prohibited, the optimizer must distribute the available capital among non-negative positions, but it is not required to distribute that capital broadly. If one or a small number of assets provide a sufficiently attractive combination of low variance and covariance with the rest of the investment universe, the variance-minimizing solution can assign most or even effectively all of the portfolio to those assets. Consequently, a long-only MV portfolio can be highly concentrated even though it is fully diversified in the formal mean-variance sense of allowing multiple assets in the opportunity set. The distinction is therefore between diversification as an available investment opportunity and diversification as an actual allocation outcome.

A related literature has emphasized the instability and estimation sensitivity of optimized portfolios. The sample covariance matrix is itself an estimated object, and optimization can amplify relatively small sampling differences in variances and correlations when the objective is minimized without an explicit penalty on concentration [11–14]. Aw *et al.* [12], for example, examine the trade-off between portfolio concentration and the number of securities held by active managers. Hou and Holmes [13] demonstrate that higher-order characteristics of return distributions can matter for minimum-variance-type hedging decisions. Ehling and Ramos [14] show that the relevant dimension of diversification can itself depend on whether diversification is evaluated geographically, industrially, or along another economically meaningful dimension. These studies collectively suggest that a statistically efficient allocation need not correspond to an allocation that is robust to alternative definitions of diversification.

The distinction between in-sample optimization and out-of-sample robustness is particularly relevant in this context. Choi and Min [11] show that well-diversified portfolios need not dominate in-sample according to conventional mean-variance criteria, while diversification can provide greater robustness

when portfolio performance is evaluated outside the estimation sample. The implication is not that equal weighting or any other simple allocation is universally superior to optimization. Rather, it motivates examining whether explicit diversification controls can moderate the sensitivity of optimized portfolios without abandoning the information contained in the covariance matrix.

In the present study, this distinction is operationalized using a long-only, fully invested minimum-variance portfolio. Because the optimization imposes both the full-investment constraint and non-negativity constraints, the empirical benchmark is not an unconstrained minimum-variance portfolio. The comparison with an entropy-regularized alternative therefore asks a more specific question: how does an explicit preference for dispersion of portfolio weights alter the solution generated by the same covariance information and investment constraints?

2.2. Entropy-based portfolio optimization

Entropy provides an information-theoretic framework for quantifying the dispersion of a probability distribution. Pola [15] defines entropy as

$$H(p) = - \sum_{i=1}^N p_i \ln p_i, \quad (1)$$

where p_i denotes the probability assigned to state i . When portfolio weights are non-negative and sum to one, the same mathematical structure can be applied to the vector of portfolio weights. Portfolio entropy is then

$$H(w) = - \sum_{i=1}^N w_i \ln w_i, \quad w_i \geq 0, \quad \sum_i w_i = 1, \quad (2)$$

with the convention $0 \ln 0 = 0$. The measure is maximized by the equal-weight portfolio, for which $w_i = 1/N$ and $H(w) = \ln N$, and declines as capital becomes increasingly concentrated in a smaller number of assets. Entropy therefore provides a direct and interpretable measure of allocation dispersion that complements variance, which measures the volatility of the resulting portfolio return rather than the concentration of the portfolio itself.

The use of entropy in portfolio construction builds on the broader information-theoretic literature on diversification and asset allocation [16–20]. Entropy-based approaches introduce a diversification preference without requiring the investor to impose an arbitrary hard constraint on the maximum weight of an individual asset. Instead, concentration is penalized continuously. Information-theoretic ideas also have a longer history in asset-pricing estimation, where entropic and linear projection methods have been connected in a common statistical framework [33].

In the present application, the entropy-regularized minimum-variance (ERMV) problem is expressed as

$$\min_w \{w^\top \Sigma w - \lambda H(w)\} \quad \text{subject to} \quad \mathbf{1}^\top w = 1, \quad w_i \geq 0, \quad (3)$$

where Σ is the estimated covariance matrix and $\lambda \geq 0$ controls the strength of the entropy preference. The parameter λ determines the trade-off between statistical variance minimization and dispersion of capital across the available assets.

An important implication is that entropy regularization should not be interpreted as a free reduction in risk. Increasing λ generally places greater value on diversification of weights, but the resulting portfolio need not have lower return variance than the unregularized MV solution. The economic question is therefore not whether entropy can simultaneously maximize entropy and minimize variance, but whether the trade-off produces a portfolio with a more desirable balance between concentration and risk-adjusted performance.

This trade-off also explains why the regularization parameter λ requires particular care. Unlike a dimensionless preference parameter, λ multiplies a dimensionless entropy term while the variance term is expressed in squared-return units of the *simple* returns used here. The numerical magnitude of λ therefore depends on the scale of the variance objective and on the frequency at which returns are measured. For this reason, a single reported value of λ is insufficient to demonstrate the economic role of entropy regularization. A more informative approach examines the solution path as λ varies. At $\lambda = 0$, Equation (3) reduces to the classical constrained minimum-variance problem. As the entropy term becomes increasingly important, the optimizer has greater incentive to distribute capital across the investment universe. Under the long-only full-investment constraints used here, sufficiently strong entropy regularization approaches the equal-weight allocation, for which $w_i = 1/N$ and $H(w) = \ln N$. This limiting relationship provides a useful diagnostic for the empirical regularization path.

Entropy-based diversification also has an important conceptual distinction from performance optimization. The entropy term is explicitly concerned with the distribution of portfolio weights; it does not directly maximize expected return or Sharpe ratio. Out-of-sample evaluation is therefore necessary if the regularized portfolio is to be interpreted as a potentially useful allocation rule rather than merely as a different full-sample portfolio construction.

The present approach is also related to broader regularization strategies in portfolio optimization. Penalties based on norms, transaction costs, turnover, sparsity, and other structural preferences have been proposed to control the instability associated with lightly constrained optimization [34, 35]. Entropy differs from many such penalties in that its natural reference point under long-only full-investment constraints is the maximally dispersed equal-weight portfolio. Alternative information-theoretic measures, including Rényi, Tsallis, Yager, and transfer-entropy formulations, have also been considered in portfolio and diversification applications [36–39], and entropy regularization has been applied within reinforcement-learning agents for allocation recommendation [40]. The present study focuses on Shannon entropy because of its tractability, interpretability, and direct connection between maximum entropy and equal weighting.

2.3. Financial networks and entropy measures

The portfolio literature emphasizes the covariance matrix as an input to allocation decisions, whereas the financial-network literature provides tools for representing the same dependence structure as a system of interconnected assets. In a financial network, assets are represented as nodes and statistical dependence is represented through links or distances between nodes [21–27]. This representation allows researchers to examine not only the magnitude of pairwise dependence but also the topology of the broader system, including centrality, clustering, path structure, and the emergence of core–periphery patterns.

One of the most widely used approaches is the minimum spanning tree (MST). Correlations can be transformed into a metric distance according to

$$d_{ij} = \sqrt{2(1 - \rho_{ij})}, \quad (4)$$

where ρ_{ij} is the correlation between the returns of assets i and j . Assets with stronger positive correlation have smaller distance. The MST retains the set of links required to connect all nodes while minimizing the total distance. It provides a parsimonious representation of the dominant dependence relationships without requiring the full $N(N - 1)/2$ set of pairwise connections to be interpreted visually. Related filtered-graph and network-based allocation methods have also been used to construct portfolios directly from dependence topology [21, 25, 41].

The MST literature has demonstrated that financial markets frequently exhibit economically meaningful structures in which assets belonging to similar sectors, geographical regions, or asset classes cluster together [21–27]. Central nodes can represent assets that connect otherwise separated portions of the network. Such topology can be informative about dependence and potential diversification opportunities, but network position should not be equated mechanically with portfolio variance: a peripheral node is not necessarily a low-volatility asset, and peripherality by itself is not an optimization criterion. The contribution of network analysis in the present study is therefore descriptive: it characterizes the dependence environment rather than imposing network topology directly on the portfolio optimization problem.

Entropy provides a second connection between portfolio optimization and network analysis. In portfolio construction, entropy is calculated from the distribution of capital weights and measures allocation dispersion. In network analysis, entropy can instead be calculated from a distribution describing the network itself. These two uses of entropy share an information-theoretic foundation but measure different forms of dispersion and should not be conflated.

A particularly useful network measure is spectral entropy derived from the eigenvalues of a correlation matrix. Let $\lambda_1, \dots, \lambda_N$ denote the eigenvalues of the correlation matrix, normalized such that $p_i = \lambda_i / \sum_j \lambda_j$. The corresponding spectral entropy is

$$H_{\text{net}} = - \sum_{i=1}^N p_i \ln p_i. \quad (5)$$

Because the eigenvalues describe the amount of variation represented by the principal components of the correlation matrix, the entropy of their normalized distribution indicates how evenly the system's variation is distributed across common factors. When variance is spread across many components, spectral entropy is relatively high; when a small number of components account for a disproportionately large share of variation, entropy falls, indicating a more concentrated or effectively lower-dimensional dependence structure [28–32].

This property makes spectral entropy particularly relevant to the study of market stress. During periods of financial disruption, correlations among assets can increase substantially, causing a larger proportion of total variation to be captured by common market components. A corresponding decline in spectral entropy can therefore signal a contraction in the effective dimensionality of the dependence structure. Such a measure does not imply that all pairwise correlations must increase simultaneously, nor does it establish a causal effect of stress on network topology. Rather, it provides a compact statistical descriptor of changes in the distribution of system-wide correlation variance.

The network-entropy and connectedness literature has used related measures to investigate changes in financial-market organization around episodes of instability, including tail-risk spillovers and synchronization among ETFs [28–32, 42]. These studies generally support the view that periods of market stress can be associated with greater synchronization, stronger dominance of common factors, and a reduction in effective diversification. This perspective complements portfolio optimization because an allocation that appears diversified under a static correlation structure may provide less effective diversification when cross-asset dependence becomes concentrated during a stress episode. A parallel strand documents that the dependence between asset classes—particularly equities and bonds—is itself dynamic and asymmetric [43], and that the diversification value of ETFs can vary across wavelet scales and market-turbulence regimes [44, 45]. Multifractal evidence on the efficiency of sector ETF indices further motivates scale-aware dependence analysis [9].

The present study applies this idea using a rolling correlation matrix over a fixed historical window. Rolling spectral entropy is used to identify changes in the dependence structure through time, while the MST provides a static representation of the principal cross-asset relationships over the full sample. The rolling measure is not used to alter the portfolio weights directly. Instead, it provides an independent empirical characterization of the market environment within which the portfolio allocation problem is embedded. Recent work has likewise examined novel procedures for MST construction and their graph-theoretic properties [46].

Overall, the literature suggests three related but distinct propositions that motivate the empirical analysis. First, variance-based optimization can produce concentrated allocations because it has no intrinsic preference for the number or dispersion of portfolio holdings [11–14]. Second, Shannon entropy can introduce an explicit and continuous preference for diversified weights, with the regularization parameter governing the trade-off between variance minimization and allocation

Table 1: ETF universe and economic exposure.

Asset class	Ticker	ETF exposure
U.S. large-cap equity	SPY	S&P 500
U.S. technology equity	QQQ	Nasdaq-100
U.S. small-cap equity	IWM	Russell 2000
Emerging-market equity	EEM	Emerging markets
Developed intl. equity	VEA	Developed ex-U.S.
Gold	GLD	Gold bullion
Silver	SLV	Silver bullion
Long-term govt. bonds	TLT	Long-term Treasuries
Interm. govt. bonds	IEF	Intermediate Treasuries
High-yield corp. bonds	HYG	U.S. high yield
Inv.-grade corp. bonds	LQD	U.S. investment grade
Real estate	VNQ	U.S. REITs

tion dispersion [15–20]. Third, financial markets exhibit time-varying dependence structures that can be represented using network methods and summarized through measures such as spectral entropy, particularly around periods of market stress [21, 28–32]. The present study brings these strands together empirically while maintaining a clear distinction between portfolio optimization and network characterization.

3. Data and variable construction

3.1. ETF universe and data source

The empirical analysis is conducted using a cross-asset universe of twelve U.S.-listed exchange-traded funds. The universe is designed to capture economically distinct sources of return and risk while maintaining a common and readily reproducible data source. The selected ETFs represent U.S. equities, international and emerging-market equities, precious metals, government bonds, corporate bonds, and U.S. real estate. Although all instruments are U.S.-listed, several provide exposure to markets outside the United States. We therefore refer to the universe as a globally diversified multi-asset ETF universe, rather than as a set of securities representing global markets in an unrestricted geographical sense.

The twelve instruments are SPY, QQQ, IWM, EEM, VEA, GLD, SLV, TLT, IEF, HYG, LQD, and VNQ. Their respective economic exposures are summarized in Table 1.

The price data are obtained from Yahoo Finance. Adjusted closing prices for all twelve ETFs are retrieved over the fixed window from 1 January 2015 through 1 August 2026. The analysis subsequently aligns the series on common trading dates and removes observations for which one or more assets are unavailable. This procedure produces a balanced panel for the multivariate portfolio and network calculations. The exact final trading date included in the aligned dataset is 31 July 2026, and the aligned panel contains 2,908 daily price observations (2,907 daily return observations). Because an open-ended “latest available” query would produce results that change when the code is executed at a later date, all reported results in this version are tied to this fixed extraction window.

The use of adjusted prices is important because the objective is to measure the realized investment return associated with holding the ETF rather than the change in its unadjusted quoted price alone. The resulting price series therefore incorporate the adjustments provided by the data source for distributions and other relevant corporate actions. Because all subsequent portfolio and network calculations are based on the same aligned dataset, the analysis uses a common observation set across assets.

3.2. Return construction

Let $P_{i,t}$ denote the adjusted closing price of ETF i on trading day t . Consistent with the portfolio aggregation and wealth compounding performed below, all analysis is conducted on *simple* (arithmetic) daily returns,

$$r_{i,t} = \frac{P_{i,t}}{P_{i,t-1}} - 1, \quad i = 1, \dots, N, \quad N = 12, \quad (6)$$

and the return vector on day t is $\mathbf{r}_t = (r_{1,t}, r_{2,t}, \dots, r_{N,t})^\top$.

Simple returns are used because the portfolio return of a fixed-weight, rebalanced portfolio is exactly the weighted average of the constituent simple returns, $r_{p,t} = \mathbf{w}^\top \mathbf{r}_t$, and because simple returns compound multiplicatively through the wealth identity $W_t = \prod_{s \leq t} (1 + r_{p,s})$. Logarithmic returns do not share the first property: a weighted average of log returns is not the log return of the rebalanced portfolio, and compounding log returns with a product of $(1 + r_{p,s})$ terms mixes return bases. Using simple returns throughout therefore keeps the portfolio return, the cumulative wealth index, maximum drawdown, and the Sharpe ratio on a single, internally consistent measurement basis. This construction is also used in the rolling out-of-sample analysis, where weights are estimated using only information available during the preceding estimation window and are subsequently applied to the following evaluation period.

3.3. Descriptive characteristics of the ETF universe

The heterogeneous composition of the ETF universe is important for interpreting the optimization results. Equity and real-estate exposures generally exhibit greater return dispersion than intermediate and long-duration government bonds, while precious-metal ETFs provide additional sources of variation that are not identical to either the equity or fixed-income segments. Corporate bond ETFs occupy an intermediate position because their returns combine interest-rate exposure with credit-risk exposure.

The descriptive statistics generated from the simple-return series are reported in Table 2. For each ETF, the table reports the mean daily simple return, standard deviation, minimum return, and maximum return. Portfolio weights are determined subsequently from the estimated covariance matrix rather than from univariate volatility rankings alone. IEF has the lowest daily standard deviation in the sample, approximately 0.00413, while SLV has the highest, approximately 0.01999. The equity and real-estate ETFs generally occupy the higher-volatility portion of the universe, whereas government and corporate bond ETFs exhibit lower unconditional dispersion.

Table 2: Descriptive statistics of daily simple returns.

Ticker	Mean	Std. Dev.	Minimum	Maximum
SPY	0.000571	0.011120	-0.109417	0.105027
QQQ	0.000778	0.013855	-0.119789	0.120031
IWM	0.000459	0.014158	-0.132669	0.091489
EEM	0.000347	0.013278	-0.124792	0.080527
VEA	0.000394	0.011017	-0.111804	0.088948
GLD	0.000458	0.010160	-0.102742	0.063586
SLV	0.000627	0.019993	-0.285402	0.091426
TLT	0.000002	0.009330	-0.066683	0.075196
IEF	0.000051	0.004130	-0.025073	0.026417
HYG	0.000182	0.005134	-0.054965	0.065459
LQD	0.000111	0.005243	-0.050030	0.073918
VNQ	0.000303	0.012871	-0.177276	0.089967

3.4. Sample alignment and reproducibility

All multivariate calculations are performed after aligning the twelve price series on common trading dates. This requirement is particularly relevant for covariance and correlation estimation because pairwise calculations based on different effective samples could otherwise produce internally inconsistent dependence matrices. The same aligned return matrix is used for the full-sample correlation analysis, covariance-based portfolio optimization, and the principal-component-based network entropy measure.

The computational implementation is conducted in Python using publicly available packages, including numpy, pandas, scipy, matplotlib, and networkx. The portfolio optimization is solved numerically using convex optimization, while the network calculations use the correlation matrix derived from the same return observations. All optimized weight vectors are checked to satisfy $\sum_i w_i \approx 1$ and $w_i \geq 0$, and the optimized objective is independently verified from $w^T \Sigma w$.

3.5. Asset-class interpretation and scope

The ETF universe is intentionally heterogeneous. SPY, QQQ, and IWM represent different segments of the U.S. equity market, while EEM and VEA provide emerging-market and developed-market international exposure, respectively. GLD and SLV represent precious metals; TLT and IEF provide different maturities of U.S. Treasury exposure; HYG and LQD capture high-yield and investment-grade corporate bonds; and VNQ represents U.S. real-estate exposure.

This construction permits the empirical analysis to examine diversification across broad economic exposures rather than within a single equity sector. At the same time, the use of U.S.-listed ETFs imposes a clear boundary on the interpretation of the results. The sample should not be interpreted as a complete representation of global financial markets. Instead, it represents a U.S.-listed investment universe containing both domestic assets and international exposures. This distinction is maintained throughout the analysis.

4. Methodology

The empirical framework combines three related but distinct components. First, constrained portfolio optimization is used to compare equal-weight, classical minimum-variance, and entropy-regularized minimum-variance allocations. Second, a correlation-based MST is constructed to characterize the dominant cross-sectional dependence structure of the ETF universe. Third, a rolling spectral-entropy measure is used to examine how the dimensionality of the correlation structure changes through time. The portfolio and network analyses use the same underlying return data but are not combined into a network-constrained optimization problem.

4.1. Covariance and correlation estimation

Let $\mathbf{r}_t$ denote the vector of daily *simple* returns for the $N = 12$ ETFs on trading day t , let $\boldsymbol{\mu}$ denote the vector of sample mean daily returns, and let Σ denote the sample covariance matrix,

$$\Sigma = \frac{1}{T-1} \sum_{t=1}^T (\mathbf{r}_t - \boldsymbol{\mu})(\mathbf{r}_t - \boldsymbol{\mu})^T, \quad (7)$$

where T is the number of aligned daily observations in the relevant estimation sample. Pearson correlations are calculated as $\rho_{ij} = \text{Cov}(r_i, r_j) / (\sigma_i \sigma_j)$. For the full-sample portfolio analysis, Σ is estimated from the complete in-sample return matrix; for the rolling out-of-sample analysis, Σ is re-estimated at each portfolio formation date using only observations contained in the preceding estimation window.

4.2. Equal-weight benchmark

The equal-weight (EW) portfolio provides the principal diversified benchmark: $w_i^{\text{EW}} = 1/N$. It satisfies $\sum_i w_i^{\text{EW}} = 1$ and $H(w^{\text{EW}}) = \ln N$; for the present twelve-ETF universe the maximum attainable portfolio entropy is $H_{\max} = \ln 12 \approx 2.4849$. The EW portfolio is not assumed to be theoretically optimal; it is a transparent benchmark representing maximal dispersion of capital under the long-only full-investment constraints.

4.3. Long-only fully invested minimum-variance portfolio

The classical minimum-variance (MV) portfolio solves

$$\min_w w^T \Sigma w \quad \text{subject to} \quad \mathbf{1}^T w = 1, \quad w_i \geq 0. \quad (8)$$

The objective is quadratic and the feasible set is convex, so the problem is a convex quadratic program. The solution minimizes the variance implied by the estimated covariance matrix without imposing any direct preference for the number or dispersion of portfolio holdings. Concentration in the solution is a consequence of the stated objective and constraints, not a computational failure.

4.4. Shannon portfolio entropy

For a feasible long-only portfolio, Shannon entropy is defined as in Equation (2), with $0 \leq H(w) \leq \ln N$. The upper bound is attained uniquely by the equal-weight allocation. An effective number of assets can be defined as

$$N_{\text{eff}} = \exp\{H(w)\}, \quad (9)$$

so that $N_{\text{eff}} = N$ under equal weighting, while a highly concentrated portfolio has a substantially smaller effective number of assets.

4.5. Entropy-regularized minimum-variance portfolio

The ERMV portfolio augments the variance objective with a Shannon entropy reward:

$$\min_w \left\{ w^\top \Sigma w + \lambda \sum_{i=1}^N w_i \ln w_i \right\} \quad \text{subject to} \quad \mathbf{1}^\top w = 1, \quad w_i \geq 0, \quad (10)$$

which is Equation (3) with the entropy term expanded. Because Shannon entropy is concave on the non-negative simplex, the term $-\lambda H(w)$ is convex for $\lambda \geq 0$, and the complete objective is convex. The parameter λ controls the relative importance of diversification: at $\lambda = 0$ the problem reduces to the classical constrained MV problem, and as λ grows the solution approaches the equal-weight allocation—the maximum-entropy point in the *interior* of the long-only simplex. Increasing λ can reduce concentration while increasing portfolio variance relative to the unregularized MV solution; the purpose of the ERMV framework is to examine whether this trade-off can be achieved without a material deterioration in risk-adjusted performance.

4.6. Regularization parameter and sensitivity analysis

The numerical value of λ cannot be interpreted independently of the scale of the variance objective. The ERMV solution is therefore evaluated over a logarithmically spaced grid of 21 values covering $\lambda \in [10^{-5}, 1]$. For each value of λ the sensitivity analysis records portfolio volatility, annualized mean return, excess-return Sharpe ratio, Shannon entropy, effective number of assets, maximum portfolio weight, and the L_2 distance from the equal-weight portfolio. The principal specification uses $\lambda = 0.01$, interpreted as an application-specific regularization choice rather than a universally optimal parameter.

4.7. Correlation-based financial network and MST

The ETF universe is represented as an undirected weighted financial network in which each ETF is a node and the correlation-based distance (4) defines edge weights. The MST is the connected, acyclic subgraph containing all nodes that minimizes the total edge distance,

$$G_{\text{MST}} = \arg \min_{G' \subseteq G} \sum_{(i,j) \in E(G')} d_{ij}. \quad (11)$$

The MST contains exactly $N - 1$ edges and is used as a descriptive network representation rather than as an additional portfolio constraint. Node degree, centrality, or peripheral position is

not interpreted mechanically as implying lower or higher portfolio variance; portfolio weights are obtained independently from Equations (8) and (10).

4.8. Static network measures

Three summary measures are calculated for the full-sample MST. The average clustering coefficient is mechanically zero because an MST is acyclic; it is reported as a structural property of the MST, not as evidence about the underlying correlation network. The average shortest path length,

$$L = \frac{1}{N(N-1)} \sum_{i \neq j} d_G(i, j), \quad (12)$$

summarizes the compactness of the filtered dependency backbone. Eigenvector centrality $\mathbf{c} = (c_1, \dots, c_N)$ is normalized as $\tilde{c}_i = c_i / \sum_j c_j$, and a static network entropy is defined as

$$H_{\text{net}} = - \sum_{i=1}^N \tilde{c}_i \ln \tilde{c}_i. \quad (13)$$

Higher values indicate a more even distribution of eigenvector centrality across the MST nodes. This measure is conceptually distinct from portfolio entropy: portfolio entropy describes the distribution of capital weights, whereas H_{net} describes the distribution of network centrality.

4.9. Rolling spectral network entropy

The static MST provides a cross-sectional description of dependence but does not capture changes in market structure through time. A rolling spectral entropy measure is therefore calculated using a 126-trading-day window. For each admissible end date t , let C_t denote the $N \times N$ correlation matrix estimated from the preceding 126 trading days, with eigenvalues $\eta_{1,t}, \dots, \eta_{N,t}$. The implementation verifies the eigenvalue spectrum numerically before constructing the entropy measure; eigenvalues are examined for materially negative values, and any numerical treatment is restricted to values within a prespecified tolerance. For the positive eigenvalues, define $p_{k,t} = \eta_{k,t}^+ / \sum_j \eta_{j,t}^+$, with $\eta_{k,t}^+ = \max(\eta_{k,t}, 0)$. The rolling spectral entropy is then

$$H_{\text{roll},t} = - \sum_{k=1}^N p_{k,t} \ln p_{k,t}. \quad (14)$$

When the eigenvalue spectrum is relatively evenly distributed, spectral entropy is higher; when a small number of principal components account for a disproportionately large fraction of the correlation structure, entropy declines. A decline in $H_{\text{roll},t}$ is interpreted as evidence of a more concentrated, lower-dimensional dependence structure. The measure is not used to determine portfolio weights or to select λ .

4.10. Portfolio return, wealth, and risk measures

For a fixed weight vector w , the daily portfolio return is exactly

$$r_{p,t} = w^\top \mathbf{r}_t, \quad (15)$$

because w is fixed (rebalanced) over the evaluation period. The cumulative wealth index and maximum drawdown are computed from simple returns as

$$W_t = \prod_{s=1}^t (1 + r_{p,s}), \quad DD_p = \min_t \left(\frac{W_t}{\max_{s \leq t} W_s} - 1 \right). \quad (16)$$

Following the conventional 252-trading-day annualization, the annualized mean return and volatility are

$$\bar{r}_p^{\text{ann}} = 252 \bar{r}_p, \quad \sigma_p^{\text{ann}} = \sqrt{252} \sigma(r_p), \quad (17)$$

where $\bar{r}_p$ is the arithmetic mean daily simple portfolio return. The terminology “annualized mean return” is used deliberately: Equation (19) annualizes the arithmetic mean of daily simple returns and is not a compounded growth rate.

4.11. Risk-free-adjusted Sharpe ratio

The Sharpe ratio is computed on the same simple-return basis as the portfolio returns. The annual risk-free rate r_f^{ann} is converted to a daily simple rate as $r_f^{\text{ann}}/252$, so that the annualized excess return is $252(\bar{r}_p - r_f^{\text{ann}}/252) = \bar{R}_p^{\text{ann}} - r_f^{\text{ann}}$ and

$$SR_p = \frac{\bar{R}_p^{\text{ann}} - r_f^{\text{ann}}}{\sigma_p^{\text{ann}}}. \quad (18)$$

All quantities in Equation (20) are simple (arithmetic) rates of return, eliminating the previous mixing of log-return means with simple risk-free rates. The principal specification uses $r_f^{\text{ann}} = 4.5\%$; robustness to 0% and 2% is reported in Section 6.

4.12. Rolling out-of-sample evaluation

Full-sample optimization does not by itself establish out-of-sample robustness. A rolling out-of-sample procedure is therefore employed. At each portfolio formation date t , the covariance matrix is estimated using the preceding 252 trading days, $\hat{\Sigma}_t = \text{Cov}(r_{t-251}, \dots, r_t)$. The EW, MV, and ERMV portfolios are constructed using information available up to that formation date, and the resulting weights are held for the subsequent 21-trading-day evaluation period: $r_{p,t+h}^{\text{OOS}} = w_{p,t}^\top r_{t+h}$, $h = 1, \dots, 21$. After the evaluation period, the estimation window is advanced and the portfolio is re-estimated, generating a sequence of genuinely out-of-sample daily portfolio returns.

4.13. Bootstrap inference for Sharpe-ratio differences

Because differences between the EW and ERMV Sharpe ratios can be very small, sampling uncertainty is evaluated using a circular block bootstrap, which accounts for temporal dependence in financial returns. Let $\Delta SR = SR_A - SR_B$ denote the Sharpe-ratio difference between two strategies. The circular block bootstrap repeatedly samples contiguous return blocks of 21 trading days with wrap-around at the sample boundary and

reconstructs bootstrap samples of the same length as the original series. For each of the 1,000 replications, the Sharpe ratio is recalculated for both strategies on the simple-return basis of Section 4.11, and the empirical distribution of ΔSR^* is used to construct percentile confidence intervals. A confidence interval for $SR_{\text{ERMV}} - SR_{\text{EW}}$ that contains zero indicates that the observed difference is not statistically distinguishable from zero.

4.14. Numerical validation and reproducibility

Several numerical checks are applied to the optimization and network calculations. First, optimized weights are checked to satisfy $\sum_i w_i \approx 1$ and $w_i \geq 0$. Second, the optimized objective is independently verified from $w^\top \Sigma w$. Third, the eigenvalue spectrum of the rolling correlation matrices is examined for numerical validity. Finally, all portfolio weights, performance statistics, network measures, and figures are generated from a single executed computational script (`replication.py`), which guarantees that the mapping between the optimized numerical weight vectors and the ETF ticker ordering is identical throughout the analysis.

5. Empirical results

The empirical analysis proceeds from the dependence structure of the ETF universe to portfolio allocation and performance. We first examine the cross-sectional correlation structure and its minimum spanning tree, then the temporal behaviour of network spectral entropy, and finally the portfolio allocations and performance of equal weighting, constrained minimum-variance optimization, and entropy-regularized minimum variance. The results are interpreted primarily in terms of the trade-off between portfolio concentration and risk: the evidence does not support the stronger claim that entropy regularization universally improves investment performance.

5.1. Cross-sectional correlation structure

The estimated correlation matrix reveals substantial heterogeneity in the degree of co-movement across the twelve ETFs. The equity ETFs form the most cohesive portion of the investment universe, reflecting their common exposure to broad movements in global equity markets. Real estate also exhibits meaningful dependence with the equity segment, whereas government bonds and precious metals occupy relatively more peripheral positions in the correlation structure. These differences matter because minimum-variance optimization operates on the complete covariance matrix and therefore evaluates each asset jointly with all other assets.

5.2. Minimum spanning tree

Figure 1 presents the minimum spanning tree constructed from the full-sample correlation-distance matrix. The network displays a comparatively cohesive equity-oriented core, with fixed-income and precious-metal exposures occupying more peripheral positions. The full-sample MST contains eleven edges, as required for a spanning tree over twelve nodes, and




Figure 1: Minimum spanning tree of the twelve-ETF correlation network. Node size and colour indicate eigenvector centrality. Shorter edges correspond to stronger (closer) correlation between ETFs.

its average clustering coefficient is zero—a mechanical consequence of the MST construction rather than evidence about the underlying correlation network. The average shortest-path length is 3.2273, and the network entropy based on normalized eigenvector centrality is 2.2537, indicating that network importance is distributed relatively broadly across nodes. This network entropy is conceptually distinct from the Shannon entropy of portfolio weights: the former measures the distribution of network centrality, the latter the distribution of invested capital.

5.3. Time variation in network spectral entropy

Figure 2 reports the rolling spectral entropy calculated using a 126-trading-day window. The rolling entropy series has a full-sample mean of approximately 1.505 and a standard deviation of approximately 0.118. Its maximum is approximately 1.723, while the minimum is approximately 1.093 and occurs on 17 March 2020, coinciding with the acute phase of the COVID-19-related market disruption. Calendar-year averages reinforce this temporary contraction: rolling entropy averages approximately 1.53 in 2015, falls to approximately 1.38 in 2020, and recovers to approximately 1.63 in 2021, with subsequent annual averages broadly around 1.43–1.50.

The decline in spectral entropy during 2020 is consistent with a concentration of the correlation matrix’s variance into fewer dominant components. From a portfolio perspective, diversification based on historically distinct asset movements can become less effective when cross-asset dependence increases during systemic stress. This result is a statistical association, not evidence that entropy causes or predicts market stress, and the rolling spectral entropy measure is not used to determine the portfolio weights.




Figure 2: Rolling spectral entropy of the correlation matrix using a 126-trading-day window. The horizontal dashed line marks the full-sample mean. The vertical marker indicates 17 March 2020, the date of the sample minimum.

Table 3: Portfolio weights under equal weighting (EW), long-only minimum variance (MV), and entropy-regularized minimum variance (ERMV, $\lambda = 0.01$).

Ticker	EW	MV	ERMV
SPY	0.0833	0.0000	0.0829
QQQ	0.0833	0.0000	0.0836
IWM	0.0833	0.0000	0.0837
EEM	0.0833	0.0000	0.0841
VEA	0.0833	0.0000	0.0828
GLD	0.0833	0.0000	0.0839
SLV	0.0833	0.0000	0.0829
TLT	0.0833	0.0000	0.0828
IEF	0.0833	0.6187	0.0831
HYG	0.0833	0.3813	0.0840
LQD	0.0833	0.0000	0.0831
VNQ	0.0833	0.0000	0.0831

5.4. Portfolio allocation

Table 3 reports the portfolio weights generated by the three allocation rules. The constrained MV portfolio is highly concentrated: it allocates 61.87% to the intermediate U.S. Treasury ETF (IEF) and 38.13% to the high-yield corporate bond ETF (HYG), while the remaining ten assets receive numerically negligible weights that are effectively zero within floating-point precision. The optimizer is not required to spread capital across the investment universe; it selects the combination of assets that minimizes estimated covariance-based portfolio variance subject to full investment and non-negativity. In this sample that solution is concentrated in two fixed-income exposures whose low individual volatilities and modest covariance with the equity cluster generate an attractive contribution to portfolio variance minimization.

It is worth noting—purely descriptively—that IEF and HYG occupy comparatively peripheral positions in the MST of Figure 1, reflecting their weak co-movement with the equity-dominated core. We emphasize, consistent with Section 2.3, that this observation is a cross-sectional description of the dependence structure, not an explanation of the optimized weights: network peripherality is not an optimization criterion, and we do not claim that it causes or justifies the concentration of the MV portfolio.

Table 4: Full-sample portfolio performance. All quantities are computed from daily simple returns. “Sharpe” is the reward-to-volatility ratio ($r_f = 0\%$); see Table 6 for the risk-free-rate sensitivity. “Max. DD” is the maximum drawdown of the simple-return wealth index $W_t = \prod_{s \leq t} (1 + r_{p,s})$.

Portfolio	Ann. Ret.	Ann. Vol.	Sharpe	Max. DD
EW	0.0794	0.1153	0.6880	−0.2611
MV	0.0242	0.0535	0.4523	−0.1791
ERMV	0.0791	0.1150	0.6879	−0.2608

The ERMV solution is qualitatively different. At $\lambda = 0.01$, the weights range only from approximately 0.0828 to 0.0841; although not mathematically identical to the equal-weight portfolio, it is extremely close to it. The regularization does not merely make the MV solution moderately less concentrated—at the selected value of λ it effectively moves the solution almost all the way toward the maximum-entropy point of the long-only simplex. The empirical result should therefore be described as strong diversification regularization rather than as a subtle covariance–entropy compromise.

5.5. Full-sample portfolio performance

Table 4 compares the three portfolios using annualized mean return, annualized volatility, the reward-to-volatility ratio (Sharpe ratio at a 0% risk-free rate), and maximum drawdown, all computed from simple returns as in Sections 4.10–4.11. Sharpe ratios at the primary 4.5% risk-free rate are discussed in Section 6 and reported in Table 6.

The performance comparison produces a striking result. The MV portfolio achieves the lowest volatility, 5.35% annually, and the smallest maximum drawdown, approximately −17.91%, but also produces the lowest annualized mean return, 2.42%, and a reward-to-volatility ratio of 0.4523. The EW portfolio has an annualized mean return of 7.94%, volatility of 11.53%, and a reward-to-volatility ratio of 0.6880, with a maximum drawdown of approximately −26.11%. The ERMV portfolio produces an annualized mean return of 7.91%, annualized volatility of 11.50%, a ratio of 0.6879, and a maximum drawdown of approximately −26.08%.

Thus, ERMV does not materially outperform EW on the principal performance measures. The two portfolios exhibit almost identical realized performance while having slightly different weights. The difference in reward-to-volatility ratios is approximately $0.6879 - 0.6880 = -0.0001$, which is economically negligible and should not be characterized as meaningful outperformance without statistical evidence. The comparison with MV is different: MV substantially reduces volatility and drawdown relative to EW and ERMV, but this reduction is accompanied by a large reduction in realized return and a lower reward-to-volatility ratio. These results illustrate why diversification and minimum variance should not be treated as synonymous.




Figure 3: Cumulative wealth of the EW, MV, and ERMV portfolios over the full sample period, computed by compounding daily simple portfolio returns as $W_t = \prod_{s \leq t} (1 + r_{p,s})$, together with the SPY buy-and-hold benchmark (shown for reference; SPY is not one of the three constructed portfolios).

5.6. Cumulative portfolio performance

Figure 3 presents the cumulative wealth paths of the EW, MV, and ERMV portfolios together with the SPY buy-and-hold benchmark. The cumulative paths of EW and ERMV are visually almost indistinguishable, consistent with their very similar annualized mean returns and volatilities. The MV portfolio follows a substantially different trajectory, reflecting its lower realized return over the sample. The SPY series, shown as a reference equity benchmark, outperforms all three constructed portfolios over the full sample; it is included for context only and is not part of the three-strategy comparison.

5.7. The regularization effect and interpretation of λ

The near-equality of the ERMV and EW allocations at $\lambda = 0.01$ makes the sensitivity of the solution to λ particularly important. At $\lambda = 0$, the ERMV problem collapses to the constrained MV problem. As λ increases, the entropy component increasingly penalizes concentrated allocations. The observed solution at $\lambda = 0.01$ indicates that, for the covariance matrix estimated in this sample, this regularization strength is sufficient to push the solution very close to equal weighting. This result should not be generalized to other asset universes or datasets: the economically relevant magnitude of λ is inherently data- and frequency-dependent, and the appropriate interpretation is based on the regularization path rather than on the absolute numerical value of λ .

5.8. Portfolio entropy and effective diversification

The portfolio-weight results provide a direct illustration of the information-theoretic mechanism. The equal-weight portfolio has maximum Shannon entropy for the twelve-asset simplex, $H_{\max} = \ln 12$. The MV allocation, concentrating all economically meaningful capital in IEF and HYG, has considerably lower entropy and a much smaller effective number of assets, while the ERMV allocation is very close to the maximum-entropy point. A count of available securities would imply identical breadth for all three strategies; entropy instead captures the distribution of capital across those securities and reveals the substantial difference between the concentrated MV

solution and the approximately uniform ERMV and EW solutions. For EW, $N_{\text{eff}} = 12$; the ERMV value is very close to this upper bound, whereas the concentrated MV portfolio has an effective breadth substantially closer to two assets. Portfolio variance answers the question of how volatile the aggregate return is under the estimated covariance structure, whereas entropy answers how dispersed the allocation of capital is; these are distinct dimensions of portfolio quality.

5.9. Network entropy and portfolio diversification

The network analysis provides an additional qualification to the portfolio results. The rolling spectral entropy series demonstrates that the dependence structure of the ETF universe changes materially through time, with its sharpest contraction occurring during the 2020 stress episode. A diversified allocation in terms of portfolio weights does not therefore necessarily imply constant effective diversification in terms of asset dependence. The ERMV framework addresses the first form of concentration—the distribution of capital—but does not explicitly optimize against time-varying network dependence. The ERMV weights are determined from the estimated covariance matrix and the entropy of portfolio weights; the rolling network entropy measure is calculated independently as a diagnostic. In particular, the static EW and ERMV portfolios do not change their weights in response to the observed contraction in spectral network entropy, and the results should not be interpreted as demonstrating that entropy regularization provides dynamic protection against systemic correlation shocks. The evidence instead suggests a useful separation of roles: portfolio entropy controls allocation concentration, while network spectral entropy describes the effective dimensionality of the underlying dependence structure.

5.10. Synthesis of empirical findings

Three principal empirical findings emerge. First, constrained minimum-variance optimization generates a highly concentrated portfolio in this twelve-ETF universe, allocating 61.87% to IEF and 38.13% to HYG with effectively zero weights elsewhere. Second, entropy regularization almost completely eliminates this concentration at $\lambda = 0.01$, assigning approximately 8.3% to each ETF; the diversification effect is strong, but it does not translate into an economically meaningful improvement in realized performance relative to EW. Third, the network analysis indicates that diversification opportunities are not static: rolling spectral entropy falls sharply during the COVID-19 stress episode, reaching its sample minimum on 17 March 2020. Taken together, the results support a precise interpretation of entropy regularization: in this application, entropy acts primarily as an effective concentration-control mechanism, transforming an extremely concentrated minimum-variance allocation into an almost uniformly weighted portfolio while realized performance remains very close to the equal-weight benchmark.

Table 5: Selected points from the ERMV regularization path. Returns are annualized arithmetic means of daily simple returns; the Sharpe ratio uses the 4.5% simple-basis risk-free rate of Section 4.11. The $\lambda = 0.01$ row is the principal specification of Table 4.

λ	Vol.	Ret.	Sharpe	Entropy	N_{eff}
10^{-5}	0.0633	0.0388	-0.0979	1.7576	5.80
10^{-4}	0.0931	0.0625	0.1883	2.4119	11.16
0.01	0.1150	0.0791	0.2966	2.4842	11.99
0.10	0.1153	0.0793	0.2979	2.4849	12.00
1.00	0.1153	0.0793	0.2980	2.4849	12.00

6. Robustness and sensitivity analysis

The full-sample results indicate that entropy regularization produces a substantial change in portfolio concentration while leaving the realized performance of the ERMV portfolio very close to that of the equal-weight benchmark. Four additional analyses examine sensitivity to (i) the entropy regularization parameter λ ; (ii) the assumed risk-free rate; (iii) sampling variation in Sharpe-ratio differences; and (iv) rolling-window out-of-sample performance.

6.1. Sensitivity to the entropy regularization parameter

Since portfolio variance is measured in squared simple-return units while Shannon entropy is dimensionless, the numerical magnitude of λ is scale-dependent and is treated as a regularization hyperparameter rather than a universal structural constant. The ERMV problem is solved over a logarithmically spaced grid of 21 values covering $\lambda \in [10^{-5}, 1]$. Table 5 reports selected points and Figure 4 presents the complete path.

At $\lambda = 10^{-5}$, annualized volatility is approximately 6.33%, Shannon entropy is 1.758, the effective number of assets is approximately 5.80, and the maximum portfolio weight is approximately 42.02%: the solution remains relatively close to the concentrated minimum-variance solution. At $\lambda = 10^{-4}$, annualized volatility rises to approximately 9.31%, entropy rises to 2.412, and the effective number of assets increases to approximately 11.16, with the maximum weight falling to approximately 14.76%. The transition continues until the ERMV solution becomes effectively equal-weighted: at $\lambda = 0.1$, entropy is 2.4849, $N_{\text{eff}} \approx 12$, the maximum weight is 0.0834, and the L_2 distance from equal weighting is approximately 0.000166; at $\lambda = 1$ the distance is approximately 0.000017. These results establish a clear variance–diversification trade-off and clarify that $\lambda = 0.01$ is an application-specific point on a broader variance–entropy path, not an empirically optimal value; no formal cross-validation is employed. Beyond approximately $\lambda = 0.1$, further increases have virtually no effect on the allocation because the solution has already converged to the *maximum-entropy equal-weight point in the interior of the simplex* (not a boundary point).




Figure 4: ERMV regularization path across $\lambda \in [10^{-5}, 1]$. The five panels show (A) annualized volatility, (B) Shannon portfolio entropy, (C) the Sharpe ratio at the 4.5% simple-basis risk-free rate, (D) the effective number of assets $N_{\text{eff}} = \exp H(w)$, and (E) the L_2 distance of the weight vector from equal weighting, each as a function of λ (log scale).

Table 6: Sensitivity of Sharpe ratios to the assumed risk-free rate (simple basis; Section 4.11).

Risk-free rate	EW	MV	ERMV
0%	0.6880	0.4523	0.6879
2%	0.5146	0.0785	0.5140
4.5%	0.2979	-0.3887	0.2966

6.2. Risk-free-rate robustness

The Sharpe ratio depends directly on the assumed risk-free rate. The principal analysis uses 4.5% annually; Table 6 recalculates the Sharpe ratios under risk-free rates of 0%, 2%, and 4.5%, with the risk-free rate expressed on the same simple-return basis as the portfolio returns throughout (daily simple return $r_f^{\text{ann}}/252$).

The absolute level of the Sharpe ratio declines as the assumed risk-free rate increases, as expected, and the precise numerical ranking of EW and ERMV is marginally sensitive to the benchmark rate, but their difference remains economically very small at every rate. More importantly, the relative inferiority of MV in risk-adjusted terms is not eliminated by reducing the risk-free rate to 2%. The evidence does not support describing ERMV as uniformly superior to EW; instead, ERMV and EW form a closely matched performance pair whose ordering

depends marginally on the assumed risk-free benchmark.

6.3. Bootstrap inference for Sharpe-ratio differences

In the principal full-sample analysis (simple basis, $r_f = 4.5\%$), $\text{SR}_{\text{EW}} = 0.2979$ and $\text{SR}_{\text{ERMV}} = 0.2966$, so

$$\Delta \text{SR}_{\text{ERMV-EW}} = 0.2966 - 0.2979 = -0.0013. \quad (19)$$

The revised analysis employs a circular block bootstrap with 1,000 replications and a block length of 21 trading days. The common resampling scheme preserves the contemporaneous relationship among the portfolio return series while allowing the sampling distribution of Sharpe-ratio differences to be examined on the corrected return basis.

The confidence interval for ERMV relative to EW contains zero (its upper endpoint equals zero to the reported four-decimal precision), so the very small observed difference in Sharpe ratios is not statistically distinguishable from zero under the implemented block-bootstrap procedure. By contrast, the bootstrap confidence intervals for ERMV–MV and EW–MV are strictly positive. As before, the inference should be interpreted with caution: the bootstrap quantifies the uncertainty of the observed differences, and the statistical distinction between ERMV and MV does not imply that the entropy term itself is responsible for all of the performance difference, because the ERMV allocation is itself close to equal weighting.

6.4. Rolling-window out-of-sample performance

To examine whether the principal conclusions survive sequential portfolio formation, a rolling-window out-of-sample procedure is implemented with a 252-trading-day estimation window and a 21-trading-day evaluation period. At each rebalancing date, the covariance matrix is estimated exclusively from the preceding 252 trading observations; the EW, MV, and ERMV portfolios are then constructed, with $\lambda = 0.01$ for ERMV, and the resulting weights are applied to the subsequent 21 trading days. All out-of-sample performance statistics are computed from daily simple portfolio returns on the basis of Sections 4.10–4.11.

The rolling results closely reproduce the qualitative full-sample pattern (Table 8). EW generates an annualized mean return of 9.20% and annualized volatility of 11.75%, for a Sharpe ratio of 0.4000; ERMV produces an annualized mean return of 9.15%, volatility of 11.70%, and a Sharpe ratio of 0.3978, with an almost identical maximum drawdown. The out-of-sample Sharpe-ratio difference is

$$\text{SR}_{\text{ERMV}}^{\text{OOS}} - \text{SR}_{\text{EW}}^{\text{OOS}} = 0.3978 - 0.4000 = -0.0022. \quad (20)$$

The MV portfolio again produces the lowest volatility, 5.63%, and the smallest maximum drawdown, approximately -19.46%, but its annualized mean return is only 3.22%, implying a Sharpe ratio of -0.2276 at the 4.5% risk-free rate.

Table 7: Circular block-bootstrap inference for Sharpe-ratio differences (simple basis; 4.5% risk-free rate; 1,000 replications; block length 21 trading days).

Comparison	Observed Δ SR	Bootstrap mean	95% CI
ERMV – EW	–0.0013	–0.0014	[–0.0029, 0.0000]
ERMV – MV	0.6853	0.7042	[0.2636, 1.1792]
EW – MV	0.6866	0.7054	[0.2634, 1.1809]

Table 8: Rolling-window out-of-sample portfolio performance (252-day estimation window, 21-day evaluation step; simple basis; Sharpe ratio at the 4.5% annual risk-free rate).

Portfolio	Ann. Ret.	Ann. Vol.	Sharpe	Max. DD
EW	0.0920	0.1175	0.4000	–0.2611
MV	0.0322	0.0563	–0.2276	–0.1946
ERMV	0.0915	0.1170	0.3978	–0.2608

6.5. Full-sample versus out-of-sample comparison

The comparison between full-sample and rolling out-of-sample results reveals a stable qualitative ordering. In both settings, MV has substantially lower volatility than EW and ERMV but also considerably lower realized returns, and EW and ERMV generate similar returns, similar volatility, similar drawdowns, and similar Sharpe ratios. The persistence of this relationship is notable because the rolling procedure changes the estimated covariance matrix at every portfolio formation date. The evidence indicates that entropy regularization can reproduce equal-weight-like diversification and performance while providing an explicit optimization-based rationale for dispersing capital across the available assets.

6.6. Overall robustness assessment

The robustness analyses provide four complementary findings. First, the λ -sensitivity analysis confirms a systematic concentration-control effect along the regularization path. Second, the risk-free-rate analysis shows that the precise ranking of EW and ERMV is marginally sensitive to the benchmark rate, but their difference remains small at every rate. Third, circular block-bootstrap confidence intervals indicate that the Sharpe-ratio difference between ERMV and EW is not statistically distinguishable from zero, whereas both diversified strategies exhibit positive Sharpe-ratio differences relative to MV. Fourth, the rolling out-of-sample experiment reproduces the central full-sample pattern. These findings strengthen the interpretation of entropy regularization as a portfolio concentration-control mechanism rather than as a guaranteed source of superior investment performance.

7. Discussion

The empirical evidence provides a nuanced interpretation of the relationship between minimum-variance optimization, entropy regularization, and financial-network structure. The

principal result is not that entropy regularization produces uniformly superior investment performance. Rather, an entropy penalty can substantially alter the concentration properties of a variance-minimizing portfolio while retaining performance very close to a simple equal-weight benchmark, and the network analysis demonstrates that the diversification environment itself is time-varying, particularly during periods of systemic stress.

7.1. Minimum-variance concentration and the role of regularization

The classical constrained minimum-variance portfolio provides the clearest illustration of the concentration problem. Although the optimization is subject to full-investment and long-only constraints, the solution places approximately 61.87% of the portfolio in IEF and 38.13% in HYG, with the remaining ten ETFs receiving effectively zero weights. The minimum-variance objective penalizes aggregate portfolio variance but does not directly penalize concentration in the allocation vector. The low volatility of the MV portfolio should therefore not be interpreted automatically as evidence of superior diversification: in the present sample, its volatility reduction is achieved primarily through concentration in fixed-income instruments rather than through broad exposure across the six represented asset classes.

Entropy regularization changes the optimization criterion by assigning an explicit value to dispersion in portfolio weights. At the principal specification $\lambda = 0.01$, the ERMV weights range only from approximately 0.0828 to 0.0841, so the solution is extremely close to the maximum-entropy equal-weight allocation. For sufficiently large values of λ , the entropy penalty fundamentally changes the geometry of the optimization problem by penalizing concentration strongly enough to suppress the corner solution generated by the variance objective.

7.2. Entropy regularization versus equal weighting

The similarity between ERMV and EW is one of the most important findings of the study and requires careful interpretation. The full-sample annualized mean returns of EW and ERMV are approximately 7.94% and 7.91%, respectively, while their annualized volatilities are approximately 11.53% and 11.50%. Under the 4.5% risk-free-rate specification, their Sharpe ratios are approximately 0.2979 and 0.2966. ERMV therefore does not provide economically meaningful outperformance relative to equal weighting; the difference is approximately –0.0013, and the block-bootstrap confidence interval contains zero. This finding does not diminish the relevance of entropy regularization; it clarifies its nature. Equal

weighting implicitly maximizes Shannon entropy subject to the simplex constraints, whereas ERMV allows diversification to emerge from an explicit trade-off between estimated variance and information-theoretic concentration. The two approaches converge when the entropy penalty becomes sufficiently strong. The empirical question is therefore not whether ERMV always dominates EW, but whether an explicit variance–entropy objective provides a useful and interpretable route toward diversified allocations without materially worsening portfolio performance; the evidence in this sample is consistent with that more modest proposition.

7.3. The variance–diversification trade-off

The regularization path demonstrates that diversification is not costless in variance terms. At weak regularization, the portfolio retains the concentration properties of minimum-variance optimization and achieves relatively low volatility. Increasing λ progressively sacrifices some of this variance efficiency in exchange for greater allocation dispersion. The resulting path can be interpreted as a continuum between two limiting allocation principles: $\lambda \rightarrow 0$ implies a variance-dominated allocation, whereas $\lambda \rightarrow \infty$ implies a maximum-entropy allocation. For the present ETF universe, this transition occurs relatively rapidly.

7.4. Network entropy and the fragility of diversification

The correlation-based MST reveals a sparse dependence structure in which equity and real-estate exposures occupy a comparatively connected core, while fixed-income and precious-metal instruments tend to occupy more peripheral positions. The associated network entropy measure provides information that is not contained in the portfolio weight vector alone. Portfolio entropy measures how capital is distributed across assets; network spectral entropy measures how the dependence structure distributes variance across the eigenmodes of the correlation matrix. These quantities are complementary rather than interchangeable measures of diversification.

The rolling network entropy results are particularly informative. The sample minimum occurs on 17 March 2020, during the acute phase of the COVID-19-related market sell-off, while the 2020 average is approximately 1.38 against a full-sample average of approximately 1.50. During systemic stress, assets that normally provide diversification can become more strongly correlated, reducing the effective dimensionality of the opportunity set. An allocation can accordingly be highly diversified in terms of portfolio weights while remaining exposed to a highly concentrated dependence structure: high portfolio entropy does not imply high network diversification. ERMV controls the concentration of capital, but it does not explicitly minimize network concentration or dynamically respond to changes in correlation topology.

7.5. Static portfolio optimization and dynamic financial networks

The coexistence of static portfolio weights and time-varying network entropy suggests that the two components of the frame-

work operate at different temporal frequencies. During periods such as March 2020, the network indicates a substantial contraction in effective diversification, but a static ERMV portfolio does not automatically respond to this change. Evidence that the equity–bond dependence structure is itself dynamic and asymmetric [43], and that ETF diversification value varies across scales and turbulence regimes [44, 45], reinforces the case for such state-dependence. The implication is not that ERMV fails during stress; rather, the present implementation does not test a dynamic ERMV strategy whose regularization strength, covariance matrix, or admissible portfolio set responds to network conditions. A natural extension would be a state-dependent regularization parameter, $\lambda_t = f(H_t^{\text{network}})$, or the use of rolling covariance matrices and dynamically estimated network information in the portfolio optimization.

7.6. Why lower minimum variance does not imply superior quality

The contrast between MV and the diversified strategies illustrates a broader point about portfolio evaluation. MV achieves annualized volatility of approximately 5.35%, compared with approximately 11.53% for EW and 11.50% for ERMV, and its maximum drawdown is also lower. However, this lower risk is accompanied by substantially lower realized return, and under the 4.5% risk-free rate the MV portfolio has a negative Sharpe ratio whereas EW and ERMV retain positive risk-adjusted performance. Portfolio quality cannot be evaluated using volatility alone, and an optimizer estimated from historical observations may exploit sample-specific relationships that do not persist. The rolling out-of-sample analysis is therefore particularly relevant: the same qualitative pattern survives sequential estimation.

7.7. Implications for portfolio construction

The findings have several implications for practical portfolio construction. First, variance-only optimization should be accompanied by explicit monitoring of portfolio concentration. Second, entropy provides a transparent concentration-control mechanism that is computationally simple and differentiable over the interior of the portfolio simplex, providing a convenient alternative to arbitrary position-count or maximum-weight constraints. Third, entropy regularization can generate equal-weight-like portfolios without imposing equal weighting *ex ante*. Fourth, allocation diversification should be complemented by monitoring of dependence concentration. Finally, the appropriate regularization strength is inherently context-dependent; there is no basis in the present evidence for treating $\lambda = 0.01$ as a universally optimal parameter.

7.8. Contribution to the entropy and network literature

The principal contribution of the study is the empirical integration of three perspectives that are frequently examined separately: covariance-based portfolio optimization, information-theoretic diversification, and financial network structure. The portfolio analysis demonstrates that entropy regularization directly addresses a weakness of minimum-variance optimization

by penalizing concentration, while the network analysis simultaneously shows that the dependence structure underlying diversification is itself dynamic. The study thereby contributes an empirical case in which information entropy is used both as an allocation regularizer and as a network diagnostic, on a fully consistent simple-return measurement basis.

7.9. Limitations and directions for future research

Several limitations should be acknowledged. First, the study does not implement a fully dynamic network-informed ERMV strategy: the covariance matrix and regularization parameter are not jointly adapted to the rolling network entropy measure. Second, the λ analysis is a sensitivity exercise rather than a formal model-selection procedure; $\lambda = 0.01$ should be interpreted as a benchmark specification. Third, the asset universe consists of twelve ETFs, and the results may not generalize to individual securities, alternative funds, currencies, commodities, or higher-dimensional portfolios. Fourth, the network analysis is based on correlations and spectral properties of the correlation matrix; correlation is informative about linear dependence but does not fully characterize nonlinear or tail-specific dependence. Future research could integrate copula-based, tail-dependence, or nonlinear network measures into the portfolio construction process. Fifth, the analysis does not incorporate transaction costs or other implementation frictions; even modest rebalancing costs could offset the marginal realized difference between ERMV and EW, which is statistically indistinguishable from zero. Future research should therefore conduct a transaction-cost-aware backtest that explicitly accounts for turnover, bid-ask spreads, and rebalancing frequency. Finally, the analysis establishes an empirical relationship between network entropy and stress-associated changes in market dependence but does not establish predictive causality.

8. Conclusion

This study examined whether Shannon entropy regularization can mitigate the concentration associated with classical minimum-variance portfolio optimization while preserving economically meaningful portfolio performance. Using a twelve-ETF global multi-asset universe covering equities, fixed income, precious metals, and real estate, and computing all portfolio quantities from simple returns on a consistent measurement basis, the study combined entropy-regularized portfolio optimization with correlation-based network analysis and rolling network entropy.

Three principal conclusions emerge. First, the classical minimum-variance portfolio exhibits substantial allocation concentration: approximately 61.87% to IEF and 38.13% to HYG, with effectively zero weight on the remaining ten ETFs. This illustrates an important distinction between minimizing portfolio variance and achieving economically meaningful diversification. Second, entropy regularization provides an effective concentration-control mechanism: at $\lambda = 0.01$ the ERMV solution assigns approximately 8.3% to each of the twelve ETFs, and the λ -sensitivity analysis shows that this outcome is part

of a systematic regularization path rather than an isolated numerical result. The entropy term should not be viewed as a mechanism that necessarily generates superior investment returns; its principal function is to introduce an explicit penalty against concentration that is absent from the variance objective. Third, the evidence does not support a claim that ERMV materially outperforms equal weighting. EW and ERMV produce almost identical annualized mean returns and volatilities, and their Sharpe ratios differ by only approximately 0.0013 under the 4.5% risk-free-rate assumption, with a block-bootstrap confidence interval that contains zero. The rolling out-of-sample analysis strengthens this interpretation: EW produces an annualized mean return of approximately 9.20% and a Sharpe ratio of 0.4000, while ERMV produces approximately 9.15% and 0.3978; MV, despite substantially lower volatility, has a negative out-of-sample Sharpe ratio.

The network analysis provides a complementary perspective. The correlation-based MST reveals a structured dependency backbone across the twelve ETFs, and the rolling network entropy measure shows that the dependence structure is not stationary, reaching its sample minimum on 17 March 2020 during the acute COVID-19 market sell-off. This finding motivates an important conceptual distinction between portfolio entropy and network entropy: the former measures how capital is distributed across assets, the latter how the dependence structure distributes variance across the eigenmodes of the correlation matrix.

Looking ahead, a natural extension is a state-dependent ERMV model in which λ or the estimation window responds dynamically to changes in rolling network entropy, allowing the portfolio to increase its covariance-awareness during low-entropy stress episodes and relax toward equal weighting when the dependence structure is more evenly distributed. More broadly, the framework—combining entropy-regularized optimization with spectral network diagnostics on a consistent return basis—provides a reproducible foundation for investigating how information-theoretic measures of concentration operate at both the allocation and market-dependence levels.

Data availability

The ETF price data were obtained from Yahoo Finance for the twelve tickers listed in Section 3 and are publicly available. The complete replication script that generates every table and figure of the paper from the raw prices is available from the corresponding author upon reasonable request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this manuscript.

Funding

The authors received no specific funding from any public, commercial, or not-for-profit organisation for this research.

References

- [1] R. Aggarwal & L. Schofield, "The growth of global ETFs and regulatory challenges," *Advances in Financial Economics* **16** (2013) 77. [https://doi.org/10.1108/S1569-3732\(2013\)0000016003](https://doi.org/10.1108/S1569-3732(2013)0000016003).
- [2] S. Smodis & S. Smore, "Exchange traded fund risk management and resiliency," *Journal of Risk Management in Financial Institutions* **13** (2019) 59. <https://doi.org/10.69554/TDRB5524>.
- [3] R. Dias, R. Galvão, M. Irfan, P. Alexandre, S. Gonçalves & L. Almeida, "Delving into the exchange-traded funds (ETFs) market: understanding market efficiency," *Journal of Ecocumanism* **3** (2024) 2670. <https://doi.org/10.62754/joe.v3i4.3787>.
- [4] K. Saha, V. Madhavan & G. R. Chandrashekar, "Relative efficiency of equity ETFs: an adaptive market hypothesis perspective," *Applied Economics Letters* **28** (2021) 1202. <https://doi.org/10.1080/13504851.2020.1804045>.
- [5] K. Bae & D. Kim, "Liquidity risk and exchange-traded fund returns, variances, and tracking errors," *Journal of Financial Economics* **138** (2020) 222. <https://doi.org/10.1016/j.jfineco.2019.02.012>.
- [6] A. Valadkhani & A. Moradi-Motlagh, "An empirical analysis of exchange-traded funds in the US," *Economic Analysis and Policy* **78** (2023) 995. <https://doi.org/10.1016/j.eap.2023.05.002>.
- [7] Z. Li & Y. Jiang, "ETF connectedness and its applications: evidence from RCEP member countries," *Journal of Multinational Financial Management* **78** (2025) 100908. <https://doi.org/10.1016/j.mulfin.2025.100908>.
- [8] R. Geng, H.-K. Zhang, Y. Gao & G. Yuan, "Decoding global economic dynamic: a graph-based examination of contemporary ETF markets," *Chaos, Solitons and Fractals* **201** (2025) 117313. <https://doi.org/10.1016/j.chaos.2025.117313>.
- [9] A. K. Tiwari, C. T. Albulescu & S.-M. Yoon, "A multifractal detrended fluctuation analysis of financial market efficiency: comparison using Dow Jones sector ETF indices," *Physica A* **483** (2017) 182. <https://doi.org/10.1016/j.physa.2017.05.007>.
- [10] H. Markowitz, "Portfolio selection," *Journal of Finance* **7** (1952) 77. <https://doi.org/10.1111/j.1540-6261.1952.tb01525.x>.
- [11] H.-S. Choi & D. Min, "Efficiency of well-diversified portfolios: evidence from data development analysis," *Omega* **73** (2017) 104. <https://doi.org/10.1016/j.omega.2016.12.008>.
- [12] E. N. W. Aw, C. Carroll, D. Setola & H. Xie, "How many securities should an active manager hold?" *Journal of Investing* **32** (2023) 15. <https://doi.org/10.3905/joi.2023.1.282>.
- [13] Y. Hou & M. Holmes, "Do higher order moments of return distribution provide better decisions in minimum-variance hedging? Evidence from US stock index futures," *Australian Journal of Management* **45** (2020) 240. <https://doi.org/10.1177/0312896219879974>.
- [14] P. Ehling & S. B. Ramos, "Geographic versus industry diversification: constraints matter," *Journal of Empirical Finance* **13** (2006) 396. <https://doi.org/10.1016/j.jempfin.2006.04.001>.
- [15] G. Pola, "On entropy and portfolio diversification," *Journal of Asset Management* **17** (2016) 218. <https://doi.org/10.1057/jam.2016.10>.
- [16] J.-R. Yu, W.-Y. Lee & W.-J. P. Chiou, "Diversified portfolios with different entropy measures," *Applied Mathematics and Computation* **241** (2014) 47. <https://doi.org/10.1016/j.amc.2014.04.006>.
- [17] X. Deng, T. Lin & C. Chen, "Comparison and research on diversified portfolios with several entropy measures based on different psychological states," *Entropy* **22** (2020) 1125. <https://doi.org/10.3390/e22101125>.
- [18] L. Batra & H. C. Taneja, "Portfolio optimization based on generalized information theoretic measures," *Communications in Statistics – Theory and Methods* **51** (2022) 6367. <https://doi.org/10.1080/03610926.2020.1861294>.
- [19] M. Bajo Traver, "Enhancing diversification in fixed-income portfolios: an entropy-based optimization framework," *Journal of Asset Management* **26** (2025) 863. <https://doi.org/10.1057/s41260-025-00428-w>.
- [20] P. J. Mercurio, Y. Wu & H. Xie, "Portfolio optimization for binary options based on relative entropy," *Entropy* **22** (2020) 752. <https://doi.org/10.3390/e22070752>.
- [21] F. Pozzi, T. Di Matteo & T. Aste, "Spread of risk across financial markets: better to invest in the peripheries," *Scientific Reports* **3** (2013) 1665. <https://doi.org/10.1038/srep01665>.
- [22] Y. Tang, J. J. Xiong, Z.-Y. Jia & Y.-C. Zhang, "Complexities in financial network topological dynamics: modeling of emerging and developed stock markets," *Complexity* **2018** (2018) 4680140. <https://doi.org/10.1155/2018/4680140>.
- [23] A. Pallotta & V. Ciciretti, "Adaptive market anomaly detection (AMAD): enhancing minimum spanning tree stability in financial networks," *Finance Research Letters* **85** (2025) 107997. <https://doi.org/10.1016/j.frl.2025.107997>.
- [24] A. A. Dellow, M. Ismail, H. Bahaludin & F. A. Razak, "Comparing the impacts of past major events on the network topology structure of the Malaysian consumer products and services sector," *Journal of the Knowledge Economy* **16** (2025) 925. <https://doi.org/10.1007/s13132-024-02038-0>.
- [25] Y. Berouaga, M. Bouchekourte, C. El Msiyah & J. Madkour, "Dynamic portfolios via triangulated maximally filtered graph networks: a centrality-driven multi-strategy approach," *Cogent Business and Management* **13** (2026) 2638724. <https://doi.org/10.1080/23311975.2026.2638724>.
- [26] M. O. Jackson & A. Pernoud, "Optimal regulation and investment incentives in financial networks," *Journal of Public Economics* **252** (2025) 105516. <https://doi.org/10.1016/j.jpubeco.2025.105516>.
- [27] C. Yan, Y. Ding, W. Liu, X. Liu & J. Liu, "Multilayer interbank networks and systemic risk propagation: evidence from China," *Physica A* **628** (2023) 129144. <https://doi.org/10.1016/j.physa.2023.129144>.
- [28] D. Lee & H. Park, "Measuring global financial linkages: a network entropy approach," *Sustainability* **11** (2019) 4691. <https://doi.org/10.3390/su11174691>.
- [29] Ç. Sözen, "Stress-aware multiscale spillover networks: cross-market transmission via coherence-entropy centrality," *AIMS Mathematics* **10** (2025) 26313. <https://doi.org/10.3934/math.20251157>.
- [30] A. Sioofy Khoojine, L. Xiao, M. Shadabfar, G. Zhao & Z. Wang, "A mutual information-based Graph-VARX network analysis of global risk spillovers in the Chinese equity market," *Open Physics* **24** (2026) 20250274. <https://doi.org/10.1515/phys-2025-0274>.
- [31] I. Choi & W. C. Kim, "A multifaceted graph-wise network analysis of sector-based financial instruments' price-based discrepancies with diverse statistical interdependencies," *North American Journal of Economics and Finance* **75** (2025) 102316. <https://doi.org/10.1016/j.najef.2024.102316>.
- [32] J. Xu & Y. Li, "Mapping cross-market tail spillovers: a multilayer LASSO-quantile network approach," *Systems* **14** (2026) 394. <https://doi.org/10.3390/systems14040394>.
- [33] Y. Kitamura & M. Stutzer, "Connections between entropic and linear projections in asset pricing estimation," *Journal of Econometrics* **107** (2002) 159. [https://doi.org/10.1016/S0304-4076\(01\)00118-X](https://doi.org/10.1016/S0304-4076(01)00118-X).
- [34] F. Şerban & S. Dedu, "Maximizing portfolio diversification via weighted Shannon entropy: application to the cryptocurrency market," *Risks* **13** (2025) 253. <https://doi.org/10.3390/risks13120253>.
- [35] S. Dedu & F. Şerban, "Entropy-based portfolio optimization in cryptocurrency markets: a unified maximum entropy framework," *Entropy* **28** (2026) 285. <https://doi.org/10.3390/e28030285>.
- [36] A. Khosravi & S. J. Sadjadi, "A multi-period robust portfolio optimization framework using Yager's entropy," *PLOS ONE* **21** (2026) e0332725. <https://doi.org/10.1371/journal.pone.0332725>.
- [37] H. Zhao, H. Ahmadzade & M. GhasemiGol, "Tsallis entropy of uncertain sets and its application to portfolio allocation," *Journal of Industrial and Management Optimization* **20** (2024) 2885. <https://doi.org/10.3934/jimo.2024032>.
- [38] S. Gaied Chortane & K. Naoui, "The end of mean-variance? Tsallis entropy revolutionises portfolio optimisation in cryptocurrencies," *Journal of Risk and Financial Management* **18** (2025) 77. <https://doi.org/10.3390/jrfm18020077>.
- [39] O. M. Ardakani, "Portfolio optimization with transfer entropy constraints," *International Review of Financial Analysis* **96** (2024) 103644. <https://doi.org/10.1016/j.irfa.2024.103644>.
- [40] M. N. Sheikh & R. M. Rahman, *Entropy regularization and tradeoff strategies for policy gradient agent recommendation in Dhaka Stock Exchange market*, 2025 IEEE/ACIS 29th International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD), Busan, South Korea, 2025, pp. 848–855. <https://doi.org/10.1109/SNPD65828.2025.11252555>.
- [41] X. Yan, H. Yang, C. Hou, S. Zhang & P. Zhu, "Portfolio optimization by price-to-earnings ratio network analysis," *International Journal of Modern Physics B* **36** (2022) 2250107. <https://doi.org/10.1142/>

- [S0217979222501077](https://doi.org/10.3390/math10193598).
- [42] N. Magner, J. F. Lavín & M. A. Valle, "Modeling synchronization risk among sustainable exchange trade funds: a statistical and network analysis approach," *Mathematics* **10** (2022) 3598. <https://doi.org/10.3390/math10193598>.
 - [43] L. Cappiello, R. F. Engle & K. Sheppard, "Asymmetric dynamics in the correlations of global equity and bond returns," *Journal of Financial Econometrics* **4** (2006) 537. <https://doi.org/10.1093/jfinec/nbl005>.
 - [44] M. Antar, R. Hassan & R. Amami, "Analyzing diversification opportunities among top-performing ETFs: a multi-wavelet and MGARCH-DCC approach," *Journal of Risk Finance* **26** (2025) 653. <https://doi.org/10.1108/JRF-11-2024-0376>.
 - [45] R. Katoch & H. Gupta, "Can traditional ETFs enhance portfolio diversification alongside thematic funds? Insights from market turbulences," *Global Business Review* (2026). <https://doi.org/10.1177/09721509261422077>.
 - [46] H. R. Bhapkar, R. U. Shaban, S. A. Mir & J. Nisar, "A novel approach for constructing a minimum spanning tree," *Journal of the Nigerian Society of Physical Sciences* **8** (2026) 3739. <https://doi.org/10.46481/jnsps.2026.3739>.