AJSSM
African Journal of Social Science and Management

African Journal of Social Science and Management
ISSN: 3141-5369
Journal homepage: www.ajssm.com

Research article

Portfolio Optimization and Financial Risk Analysis using Quantitative and Computational Methods: Evidence from selected S&P 500 Firms

Oyebamiji Sodiq Akorede*, Adeyemi Paul Oluwatimilehin, Adedeji Teslim Olaide, Akinyemi Akeem Alabi, Ruth Sinmilejesu Alesinloye

Università Degli Studi L'aquila, Italy

* Corresponding author.
E-mail address: Oyebamijisodiq70@gmail.com (Akorede, O. S.).

Abstract
This study examines portfolio optimization and financial risk analysis for ten large-cap U.S. equities drawn from seven GICS sectors of the S&P 500, covering the period January 2021 to December 2025. Motivated by the well-documented limitations of classical mean-variance (MV) optimization (including sensitivity to return estimation error and the systematic underestimation of tail risk in non-normal return environments), the study adopts an integrated quantitative and computational framework comprising LASSO regression-based return prediction, Markowitz MV optimization, historical Value-at-Risk (VaR), Conditional Value-at-Risk (CVaR), and Monte Carlo simulation. Three portfolio strategies are constructed and evaluated which included a Classical MVO portfolio using historical mean returns, a Machine Learning-Enhanced portfolio substituting LASSO-predicted returns as expected return inputs, and an Equal-Weight benchmark. Results indicate that both optimized portfolios substantially outperform the Equal-Weight benchmark on a risk-adjusted basis, achieving Sharpe ratios of 1.374 and 1.317 respectively against the benchmark's 0.826. The ML-Enhanced portfolio generates the highest mean simulated five-year terminal wealth of 4.037 times the initial investment, with a near-zero probability of capital loss (0.05%) compared to 0.75% for the Equal-Weight strategy. Individual firm risk analysis reveals pronounced tail-risk heterogeneity, with Amazon recording the most extreme 99% CVaR of −25.210% and Coca-Cola the most resilient at −5.964%. The systematic divergence between VaR and CVaR estimates across all entities provides direct empirical support for CVaR's theoretical superiority as a risk measure in leptokurtic return environments which affirm the value of integrating machine learning with classical optimization theory and tail-sensitive risk measurement for evidence-based portfolio construction.

Keywords: Portfolio optimization, Mean-Variance Optimization, LASSO regression, Conditional Value-at-Risk, Monte Carlo simulation, financial risk analysis

Received 24 June 2026; Received in revised form 09 August 2026; Accepted 01 September 2026; Available online 05 October 2026. © 2026 Authors.
To publish with AJSSM, visit: https://ajssm.com/call-for-papers

How to cite this article: Akorede, O. S., Oluwatimilehin, A. P., Olaide, A. T., Alabi, A. A., Alesinloye, R. S. (2026). Portfolio Optimization and Financial Risk Analysis using Quantitative and Computational Methods: Evidence from selected S&P 500 Firms. African Journal of Social Sciences and Management, 1(2): 10-25.

INTRODUCTION

The discipline of portfolio management has evolved dramatically over the past seven decades, transitioning from intuitive, experience-driven decision-making into a rigorous, mathematically grounded field whose foundations are embedded in quantitative theory. At the core of this transformation lies the enduring challenge of constructing investment portfolios that maximize returns while prudently managing the risks inherent in financial markets (a challenge that has become increasingly complex in an era of global capital flows, heightened market interconnectedness, and rapid technological change). The intellectual turning point in this journey was Harry Markowitz's landmark 1952 paper, "Portfolio Selection," published in the Journal of Finance, which introduced the mean-variance (MV) framework and fundamentally altered how investors conceptualize risk. Markowitz demonstrated that rational investors should not evaluate assets in isolation, but should instead consider how individual securities contribute to the overall risk and return of a portfolio, a principle that gave rise to Modern Portfolio Theory (MPT) and earned him the Nobel Memorial Prize in Economic Sciences in 1990 (Markowitz, 1952). Building on this foundation, Sharpe (1964) introduced the Capital Asset Pricing Model (CAPM), which established a linear relationship between an asset's systematic risk (measured by beta) and its expected return, providing investors with a tractable framework for pricing risk across diversified portfolios. These twin pillars of MPT have since formed the bedrock of institutional portfolio construction and financial risk management worldwide. The practical significance of these frameworks is reflected in the extraordinary scale of the global asset management industry. According to the Thinking Ahead Institute (2025), total assets under management (AUM) at the world's 500 largest asset managers reached a record USD 139.9 trillion at the end of 2024 (a 9.4% increase from the prior year) with North American managers alone accounting for over USD 74 trillion of that total. Tanvir-Rahman and Leila (2025), projects that global AUM will climb further to USD 200 trillion by 2030, growing at a compound annual growth rate of approximately 6.2%, as rising middle-class wealth across emerging markets, the proliferation of defined-contribution pension schemes, and the democratization of investment platforms collectively expand the universe of investable capital. The sheer magnitude of capital being deployed globally makes the question of how portfolios are constructed and how risk is measured not merely an academic exercise, but a matter of profound economic consequence. Even marginal improvements in optimization methodology or risk quantification accuracy, applied across portfolios of this scale, can translate into billions of dollars of preserved or generated value for investors ranging from sovereign wealth funds and pension managers to retail participants accessing equity markets through index-linked instruments.

Despite the theoretical elegance and enduring influence of the MV framework, decades of empirical research have documented critical limitations that constrain its real-world effectiveness. Chief among these is the model's foundational assumption that asset returns are normally distributed; an assumption that real-world financial data consistently violates. Empirical return distributions are characterized by fat tails and negative skewness, meaning that extreme negative outcomes occur far more frequently than the normal distribution would predict (Jondeau &

Rockinger, 2006). This property, known as leptokurtosis, means that MV-optimized portfolios systematically underestimate the probability and magnitude of catastrophic losses. Compounding this problem is the framework's well-documented sensitivity to input parameter estimation: small errors in expected return estimates or covariance matrices can produce wildly divergent optimal portfolio weights, rendering solutions practically unstable and difficult to implement (Michaud, 1989). Furthermore, the static nature of MV optimization which is its inability to adapt to changing market regimes, volatility clusters, or structural breaks in return correlations, limits its applicability in dynamic, real-world investment environments. These theoretical shortcomings were exposed with devastating clarity during the global financial crisis of 2007–2009, during which cross-asset return correlations spiked simultaneously across virtually all risky asset classes, collapsing the diversification benefits that MV theory had promised. Traditional risk measures such as Value-at-Risk (VaR), which had been widely adopted by financial institutions and endorsed by regulatory frameworks such as Basel II, similarly failed to provide adequate warning of the magnitude of tail losses materializing across structured credit, equity, and commodity markets (Soncin & Moreira, 2023). The crisis prompted a fundamental reassessment of the quantitative toolkit available to portfolio managers and risk officers, accelerating the adoption of more sophisticated methodologies capable of modelling the non-linear, fat-tailed, and regime-dependent behaviour of financial asset returns.

In response to these recognized deficiencies, a rich body of academic and practitioner literature has emerged advocating for more robust, computationally intensive approaches to portfolio optimization and financial risk analysis. Conditional Value-at-Risk (CVaR), introduced by Rockafellar and Uryasev (2002), has emerged as the most widely adopted alternative to VaR, measuring the expected loss in the worst tail of the return distribution rather than merely the threshold level. Unlike VaR, CVaR is a coherent risk measure which satisfies the mathematical properties of subadditivity and convexity making it both theoretically superior and practically tractable as an optimization objective (Rockafellar & Uryasev, 2002). Monte Carlo simulation frameworks have similarly gained prominence, enabling analysts to generate thousands of synthetic portfolio return scenarios drawn from empirically calibrated distributions and thereby circumvent the restrictive normality assumptions of analytical models (Meher & Mishra, 2024). GARCH-family models, pioneered by Engle (1982) and Bollerslev (1986), have become standard tools for modelling volatility clustering which is an empirically well-established phenomenon whereby periods of high market volatility tend to be followed by further high-volatility periods, providing a more realistic characterization of the time-varying risk environment in which portfolios operate. More recently, machine learning and artificial intelligence-driven approaches have begun to reshape the frontier of portfolio optimization, with hybrid models demonstrating superior return-prediction accuracy relative to classical benchmarks in several empirical studies (Masuda, 2024). A 2024 review of mathematical models in portfolio optimization identified a significant surge in research activity incorporating these advanced techniques, reflecting the field's decisive shift toward more computationally sophisticated, data-driven methodologies (Alonso et al., 2025).

Against this evolving methodological backdrop, the United States equity market and specifically the Standard & Poor's 500 Index (S&P 500) presents an ideal empirical setting for

rigorously evaluating these approaches. The S&P 500 is widely regarded as the single most comprehensive gauge of large-cap U.S. equity performance, comprising 500 leading companies and covering approximately 80% of available U.S. market capitalization (S&P Dow Jones Indices, 2024). As of 2024, the index delivered a total return of 25.0%, following an equally impressive 26.3% in 2023 which was two consecutive years of returns exceeding 20%, a feat achieved only a handful of times in the index's history (Broadway Bank, 2025). These exceptional returns reflect a confluence of factors including resilient corporate earnings, artificial intelligence-driven investment enthusiasm, and a soft-landing macroeconomic narrative that defied earlier recessionary forecasts. The index spans eleven Global Industry Classification Standard (GICS) sectors (Information Technology, Health Care, Financials, Consumer Discretionary, Communication Services, Industrials, Consumer Staples, Energy, Utilities, Real Estate, and Materials), providing a structurally diverse investment universe that captures the full spectrum of economic activity within the U.S. economy. However, the S&P 500's recent performance has been accompanied by a troubling structural development: the growing concentration of index weight among a small number of mega-cap technology and communication services companies. As of 2024, the eight largest S&P 500 constituents accounted for 35.6% of the total index weight; more than double their 15.5% share a decade earlier (Broadway Bank, 2025). This concentration introduces significant idiosyncratic risk into passive index-tracking strategies and raises fundamental questions about the degree to which the S&P 500, in its current form, delivers the diversification benefits that MPT associates with broad market exposure.

It is within this context that the present study identifies a meaningful and practically relevant research gap. While the literature on portfolio optimization and financial risk analysis is extensive, there remains a notable absence of unified, empirically grounded studies that simultaneously compare classical and advanced optimization frameworks across a deliberately constructed, cross-sector selection of prominent S&P 500 firms. Existing studies tend to focus either on the full index, on isolated sectors, or on theoretical model comparisons that lack the firm-level empirical specificity required to generate actionable insights for practitioners (Cotter et al., 2025; Anca-Adriana, 2023; Meher & Mishra, 2024). This study addresses that gap by selecting ten firms including Microsoft Corporation (Information Technology), Eli Lilly & Co. (Health Care), JPMorgan Chase & Co. (Financials), Amazon.com Inc. (Consumer Discretionary), Alphabet Inc. (Communication Services), Caterpillar Inc. (Industrials), The Coca-Cola Company (Consumer Staples), Exxon Mobil Corporation (Energy), NextEra Energy Inc. (Utilities), and American Tower Corporation (Real Estate), that collectively span seven of the S&P 500's eleven GICS sectors. This multi-sectoral composition provides a rich and heterogeneous portfolio universe whose cross-asset return correlations, volatility profiles, and risk-return characteristics reflect the genuine diversity of economic exposures available within the U.S. large-cap equity market. By examining both individual firm-level risk and aggregate portfolio-level optimization outcomes within a single analytical framework, this study generates insights of direct relevance to investors seeking to build resilient, well-diversified portfolios in an environment of heightened market concentration and evolving risk dynamics. Accordingly, this study is guided by the following objectives:

i. To construct an optimized investment portfolio from the ten selected S&P 500 firms and identify the best-performing portfolio based on risk-adjusted returns.

ii. To measure and compare the financial risk of each selected firm using standard quantitative risk metrics, including Value-at-Risk (VaR) and volatility analysis.

At the empirical level, it generates new, firm-specific and portfolio-level evidence on the performance of quantitative optimization and risk analysis techniques within the unique structural context of the post-2020 S&P 500. At the methodological level, the application of multiple risk metrics within a unified framework provides a richer characterization of portfolio risk than any single measure can offer alone which is an approach consistent with evolving best-practice standards in financial risk management (Orgeldinger, 2017).

LITERATURE REVIEW

The theoretical concerns raised regarding the inadequacy of classical portfolio frameworks under real market conditions are extensively corroborated by this empirical review. Meher and Mishra (2024), applied Monte Carlo simulation to construct risk-adjusted portfolios, demonstrating that simulation-based methods more accurately reflect the fat-tailed return distributions observed in equity markets than classical mean-variance (MV) optimization, and yields superior downside risk containment. On the measurement of risk, Liu (2022), defined Value-at-Risk (VaR) as an estimate of the greatest possible portfolio loss under normal market conditions over a given time horizon and confidence level which is a characterization widely adopted in the applied risk management literature. However, empirical tests consistently exposed VaR's limitations, including parametric VaR models, which assume normally distributed returns, systematically underestimated tail losses during the 2008 global financial crisis and the 2020 COVID-19 market shock. These evidences reinforced the well-established case for Conditional Value-at-Risk (CVaR), whose empirical superiority over VaR in terms of tail loss protection has been confirmed across multiple studies since the foundational work of Rockafellar and Uryasev's (2002). Wang (2022), in a comparative study of VaR models applied to financial time series, further found that only GARCH-based specifications which explicitly model time-varying volatility, passed back testing at the 1% significance level, underscoring the practical necessity of volatility-sensitive risk frameworks over static alternatives. The empirical evidence on cross-sector diversification also carries direct relevance to the present study. Cotter et al. (2025), using over two decades of U.S. market data spanning multiple crisis periods, found that cross-asset diversified portfolios consistently outperformed single-market alternatives; although diversification benefits from international equities diminished greatly after 2009 as global correlations converged. Within the S&P 500 specifically, Anca-Adriana (2023), applied Bayesian Neural Networks and Expected Shortfall to a portfolio of 12 index constituents over the period 2011–2023, identifying significant inter-sectoral risk contagion during the 2020–2023 turmoil period; suggesting that sectoral diversification benefits within U.S. large-cap equities are regime-dependent and most fragile precisely when investors need them most. A broader empirical study on portfolio size and risk diversification similarly confirmed an inverse relationship between the number of assets held and portfolio standard deviation, with skewness also diminishing as holdings increased (validating the foundational diversification logic of

Markowitz (1952)), while simultaneously highlighting the importance of asset selection and sector composition in portfolio construction.

These empirical findings are anchored in two theoretical frameworks that have shaped the discipline of portfolio management for over seven decades. The first is Modern Portfolio Theory (MPT), introduced by Markowitz (1952), which established that rational investors should evaluate portfolios on the joint basis of expected return and variance rather than on individual asset performance alone. The efficient frontier (the set of portfolios offering maximum return for a given level of risk) remains the central construct of portfolio optimization, and the insight that asset correlations determine the diversification benefit of combining securities laid the mathematical groundwork for all subsequent portfolio construction models. Despite its enduring influence, the MV framework has been consistently criticized for its sensitivity to input estimation errors and its failure to adequately model tail risk (Michaud, 1989); thus, motivating the broader search for more robust optimization criteria that characterizes much of the contemporary literature. On the other hand, the other foundational theory is the Capital Asset Pricing Model (CAPM), developed by Sharpe (1964), which extended MPT to an equilibrium pricing framework by establishing that an asset's expected return is a linear function of its beta; the measure of its systematic, non-diversifiable risk relative to the market portfolio. The CAPM's central distinction between systematic risk (which demands a return premium) and idiosyncratic risk (which can be eliminated through diversification) remains conceptually central to risk attribution and portfolio evaluation. Empirically, however, the model has faced persistent challenges as Fama and French (1992), demonstrated that beta alone cannot explain the cross-section of stock returns, while Soncin and Moreira (2023), analyzing the beta anomaly across 25 developed and 23 emerging markets, found that low-beta stocks consistently earned higher returns than CAPM predicts, directly contradicting the model's core proposition. Together, MPT and CAPM provide the theoretical scaffolding upon which this study's analytical framework is built, while the limitations in both the theoretical and empirical literature justify the adoption of more advanced quantitative tools.

RESEARCH METHOD AND MATERIALS

This study adopts a quantitative and computational research design integrating classical portfolio optimization, machine learning-enhanced return prediction, and tail-risk measurement. The analytical pipeline comprises four components: data preparation, LASSO regression-based return forecasting, Mean-Variance (MV) optimization, and risk quantification via VaR, CVaR, and Monte Carlo simulation; consistent with hybrid ML-optimization approaches advocated in recent portfolio literature (Masuda, 2024). All computations are performed in Python using the pandas, scikit-learn, and PyPortfolioOpt libraries.

Data and Sample

Monthly adjusted closing prices for ten S&P 500 firms spanning seven GICS sectors were sourced from Yahoo Finance and Macrotrends for January 2021 to December 2025 (60 observations

per firm). Monthly returns were derived as percentage price changes. The five-year window spans four distinct market regimes namely the 2021 post-pandemic recovery, the 2022 inflationary bear market, the 2023–2024 AI-driven bull run, and the 2025 recalibration; thus, providing a heterogeneous environment that tests the robustness of optimization outcomes across varying market conditions. The dataset was loaded on to Python and returns computed as follows:

```
import pandas as pd, numpy as np
df = pd.read_excel("SP500_Monthly_Closing_Prices_2021_2025.xlsx",
sheet_name="Monthly Closing Prices", skiprows=2, header=0)
df.columns = ["Month","MSFT","LLY","JPM","AMZN","GOOGL","CAT","KO","XOM","NEE","AMT"]
df = df.set_index(pd.to_datetime(df.pop("Month"), format="%b-%y")).astype(float)
returns = df.pct_change().dropna()
```

LASSO Return Prediction

To introduce a machine learning dimension, LASSO regression (Tibshirani, 1996), generates forward-looking return forecasts for each asset. LASSO imposes an L1 regularization penalty on the least-squares objective $(min\left\{\left(\frac{1}{2n}\right)\|y - X\beta\|^2 + \lambda\|\beta\|_1\}\right)$, driving irrelevant predictor coefficients to exactly zero and performing automatic variable selection. Each asset's monthly return is modelled as a function of all ten assets' prior-month returns, exploiting cross-asset momentum and mean-reversion signals. The regularization parameter λ is selected via five-fold cross-validation, and one-step-ahead forecasts are annualized (*12) for use as expected return inputs in the optimization stage:

```
from sklearn.linear_model import LassoCV

from sklearn.preprocessing import StandardScaler

X = returns.shift(1).dropna(); Y = returns.loc[X.index]

Xs = StandardScaler().fit_transform(X)

ml_mu = {t: LassoCV(cv=5,max_iter=10000,random_state=42).fit(Xs,Y[t].predict(Xs[[-1]])[0]*12 for
t in Y.column s}
```

Portfolio Optimization

Optimization follows the Markowitz (1952) MV framework, maximizing the Sharpe ratio $SR = \frac{(\mu_p - r_f}{\sigma_p}$, subject to long-only constraints (weights ≥ 0, sum to 1). The risk-free rate is 5.25%, consistent with the Federal Funds Rate at end-2024. Three strategies are evaluated: (i) Classical MVO, using historical mean returns; (ii) ML-Enhanced MVO, substituting LASSO-predicted returns into the same covariance matrix; and (iii) Equal-Weight (10% per asset), as a naive diversification benchmark (Cotter et al., 2025). The annualized covariance matrix is estimated from monthly returns scaled by 12, implemented via PyPortfolioOpt:

```
from pypfopt import EfficientFrontier, risk_models, expected_returns
mu_cl = expected_returns.mean_historical_return(df, frequency=12)
S     = risk_models.sample_cov(df, frequency=12)
# Classical MVO
w_cl = EfficientFrontier(mu_cl, S).max_sharpe(0.0525) or ef.clean_weights()
# ML-Enhanced MVO
ml_s = pd.Series(ml_mu).clip(-0.30, 0.80)
w_ml = EfficientFrontier(ml_s,  S).max_sharpe(0.0525) or ef.clean_weights()
```

Risk Measurement

Historical VaR and CVaR are computed at 95% and 99% confidence levels for individual firms and all three portfolio strategies. VaR is the empirical return quantile at (1 − α), representing the maximum one-month loss not exceeded with probability α. CVaR defined as $CVaR_a = E\left[\frac{L}{L} > VaR_a\right]$, is the expected loss conditional on exceeding VaR, and constitutes a coherent risk measure superior to VaR in non-normal environments (Rockafellar & Uryasev, 2002):

```
def var_cvar(r, cl=0.95):
    v = np.percentile(r, (1-cl)*100)
    return round(v*100,3), round(r[r<=v].mean()*100,3)
```

Monte Carlo Simulation

Following Meher and Mishra (2024), 10,000 simulation paths are generated per portfolio over 60 months. Monthly returns are drawn from a multivariate normal distribution parameterized by the historical mean vector and sample covariance matrix. Terminal wealth is the compounded product of monthly portfolio returns, normalized to an initial value of 1.0:

```
def monte_carlo(w_dict, rets, n=10000, T=60):
    w = np.array([w_dict.get(t,0) for t in rets.columns])
    finals = [np.prod(1 + np.random.multivariate_normal(
    rets.mean().values, rets.cov().values, T) @ w) for _ in range(n)]
    return np.array(finals)
```

RESULT PRESENTATION AND DISCUSSION

Descriptive Statistics

Table 1. Descriptive Statistics of Monthly Returns (January 2021 – December 2025)

Firm	Mean Ret.(%)	Std Dev(%)	Min(%)	Max(%)	Skew.	Kurt.
Microsoft (MSFT)	1.550	7.112	-20.902	17.629	-0.144	0.926
Eli Lilly (LLY)	2.796	7.234	-10.080	25.850	0.741	1.339
JPMorgan (JPM)	1.536	5.625	-9.291	14.561	0.107	-0.536
Amazon (AMZN)	1.078	8.496	-25.210	23.253	-0.309	1.859
Alphabet (GOOGL)	1.775	7.411	-17.902	18.836	-0.142	0.257
Caterpillar (CAT)	1.703	7.050	-12.079	19.195	0.080	-0.275
Coca-Cola (KO)	0.846	3.255	-5.964	7.438	0.117	-0.503
Exxon Mobil (XOM)	2.120	6.723	-10.367	26.358	0.953	1.911
NextEra (NEE)	0.547	5.858	-13.562	11.211	-0.300	-0.673
Amer. Tower (AMT)	0.386	6.504	-13.912	20.533	0.144	0.382

Note: Return figures in monthly percentages. Skewness and Kurtosis via standard moment estimators.

Eli Lilly recorded the highest mean monthly return (2.796%), followed by Exxon Mobil (2.120%) and Alphabet (1.775%). American Tower (0.386%) and NextEra Energy (0.547%) lagged considerably. Amazon displayed the highest volatility (8.496%) and Coca-Cola the lowest (3.255%). Exxon Mobil and Amazon exhibit elevated kurtosis (1.911 and 1.859) and skewness, which implies that fat-tailed return distributions that motivate the use of CVaR as the primary downside risk metric (Jondeau & Rockinger, 2006).

Individual Firm Risk: VaR and CVaR

Table 2. Individual Firm Value-at-Risk and Conditional Value-at-Risk (Monthly, %)

Firm	**VaR 95%(%)**	**CVaR 95%(%)**	**VaR 99%(%)**	**CVaR 99%(%)**
Microsoft (MSFT)	-8.382	-13.180	-14.571	-20.902
Eli Lilly (LLY)	-7.930	-9.250	-9.677	-10.080
JPMorgan (JPM)	-6.883	-8.949	-9.119	-9.291
Amazon (AMZN)	-13.869	-19.792	-21.873	-25.210
Alphabet (GOOGL)	-10.567	-14.276	-15.690	-17.902
Caterpillar (CAT)	-10.444	-11.266	-11.496	-12.079
Coca-Cola (KO)	-3.875	-5.348	-5.863	-5.964
Exxon Mobil (XOM)	-6.409	-8.507	-9.406	-10.367
NextEra (NEE)	-9.141	-11.726	-12.254	-13.562
Amer. Tower (AMT)	-10.143	-12.638	-13.054	-13.912

Note: Negative values denote downside losses at respective confidence levels.

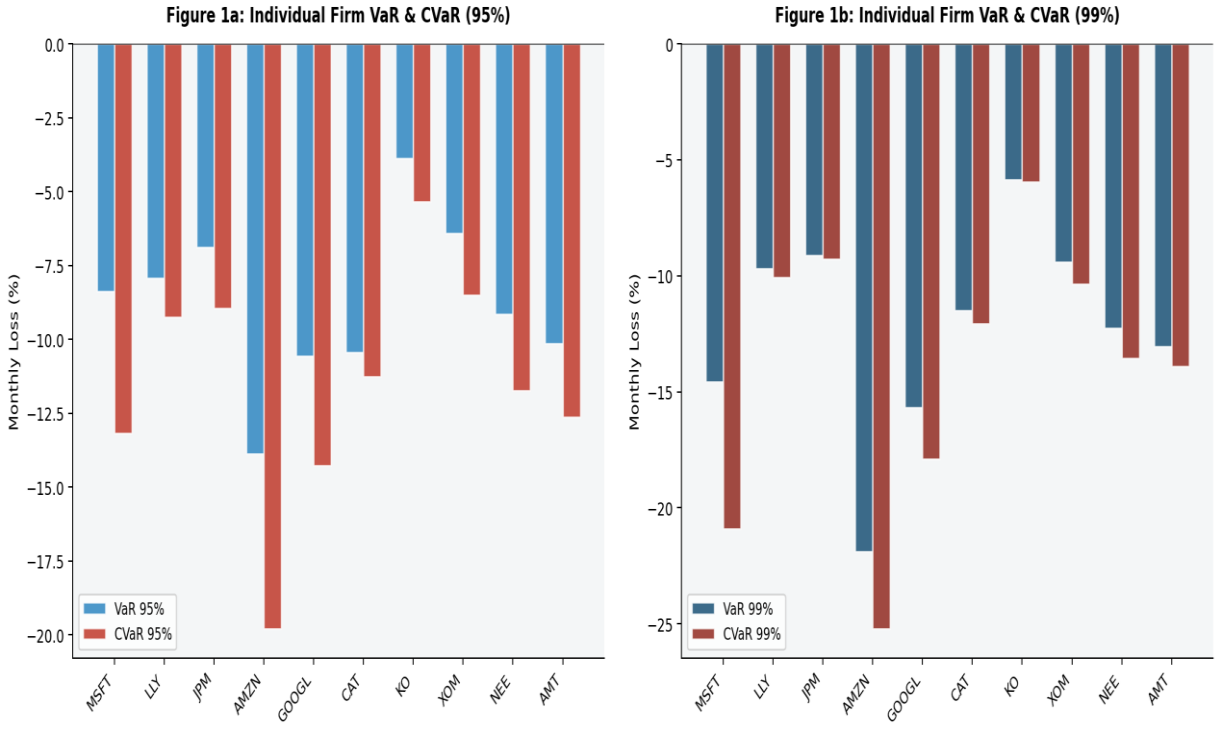


Figure 1. Individual Firm VaR and CVaR at 95% (left) and 99% (right) Confidence Levels

Amazon exhibited the most extreme tail risk (99% CVaR: −25.210%), while Coca-Cola remained the most resilient (−5.964%). The pronounced VaR–CVaR gap for Amazon (95% VaR: −13.869% vs. CVaR: −19.792%) and Microsoft (−8.382% vs. −13.180%) support that losses beyond the VaR threshold are systematically more severe than the threshold implies. JPMorgan Chase and Exxon Mobil occupy a moderate risk band consistent with their cyclical sector profiles.

Portfolio Optimization

Table 3. Portfolio Performance Comparison

Metric	Classical MVO	ML-Enhanced	Equal-Weight
Expected Annual Return	28.43%	28.01%	17.20%
Annual Volatility	16.87%	17.29%	14.48%
Sharpe Ratio	1.374	1.317	0.826
VaR 95% (monthly)	-5.303%	-5.163%	-4.802%
CVaR 95% (monthly)	-6.810%	-6.909%	-7.122%
VaR 99% (monthly)	-7.211%	-7.376%	-7.713%
CVaR 99% (monthly)	-9.619%	-9.605%	-10.033%

Note: Returns annualized ×12; volatility ×√12. Risk-free rate = 5.25%.

Table 4. Optimal Portfolio Weight Allocations

Firm	Sector	Classical MVO	ML-Enhanced	Equal-Wt.
Microsoft (MSFT)	Information Technology	0.00%	0.00%	10.00%
Eli Lilly (LLY)	Health Care	49.43%	54.53%	10.00%
JPMorgan (JPM)	Financials	23.95%	25.42%	10.00%
Amazon (AMZN)	E-Commerce	0.00%	0.00%	10.00%
Alphabet (GOOGL)	Communication Services	1.73%	0.00%	10.00%
Caterpillar (CAT)	Industrials	0.00%	0.00%	10.00%
Coca-Cola (KO)	Consumer Staples	0.00%	0.00%	10.00%
Exxon Mobil (XOM)	Energy	24.89%	20.05%	10.00%
NextEra (NEE)	Utilities	0.00%	0.00%	10.00%
Amer. Tower (AMT)	Real Estate	0.00%	0.00%	10.00%

Note: Long-only constraints (weights ≥ 0, sum to 1). Optimization via PyPortfolioOpt.

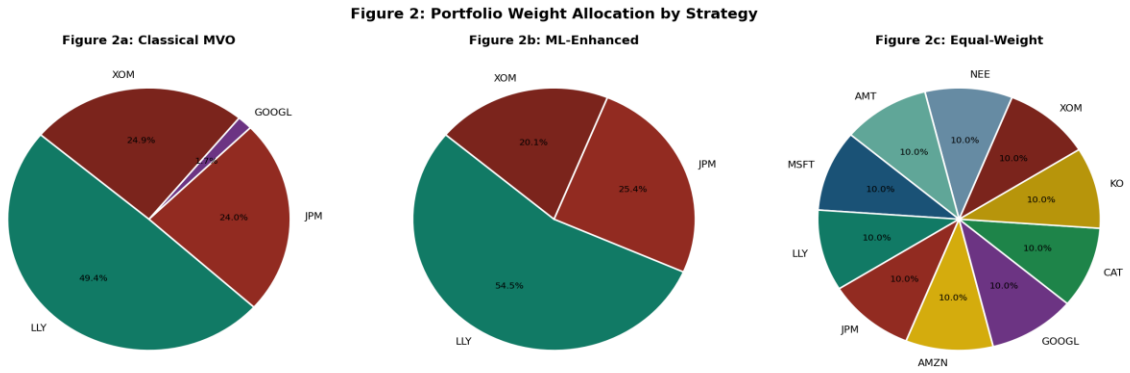


Figure 2: Portfolio Weight Allocations — Classical MVO (left), ML-Enhanced (centre), Equal-Weight (right)

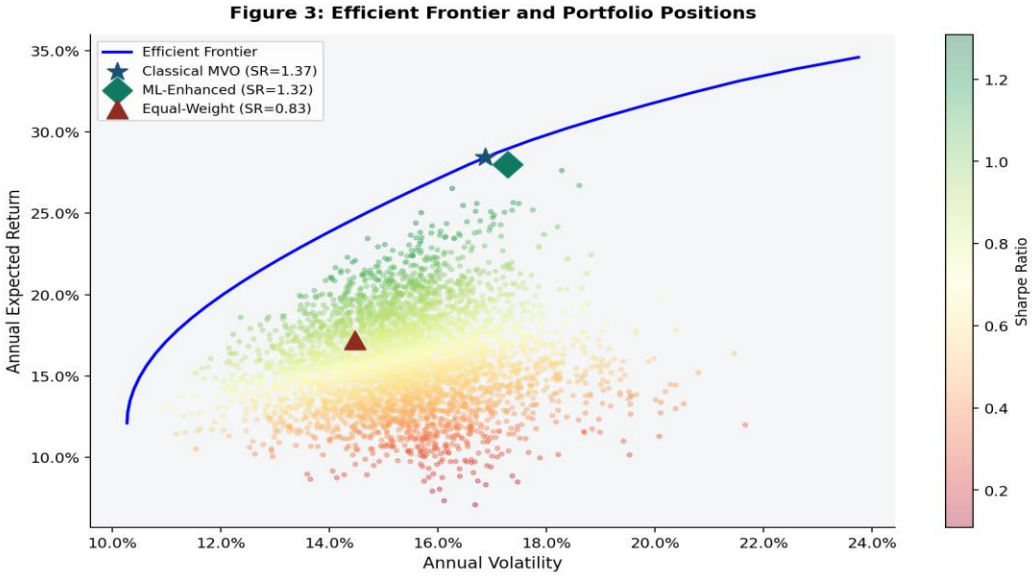


Figure 3. Mean-Variance Efficient Frontier with Portfolio Positions (coloured by Sharpe ratio; r_f = 5.25%)

The Classical MVO portfolio achieved a Sharpe ratio of 1.374, concentrating in Eli Lilly (49.43%), Exxon Mobil (24.89%), and JPMorgan Chase (23.95%) which were the assets that delivered the most favourable return-to-covariance trade-offs over the sample period. The ML-Enhanced portfolio yielded a comparable Sharpe ratio of 1.317, with LASSO assigning the highest forward return estimates to Eli Lilly (35.05% annualized) and Exxon Mobil (22.84%), while placing a marginally negative forecast on American Tower (−1.16%). Both strategies substantially outperformed the Equal-Weight benchmark (Sharpe: 0.826; return: 17.20%).

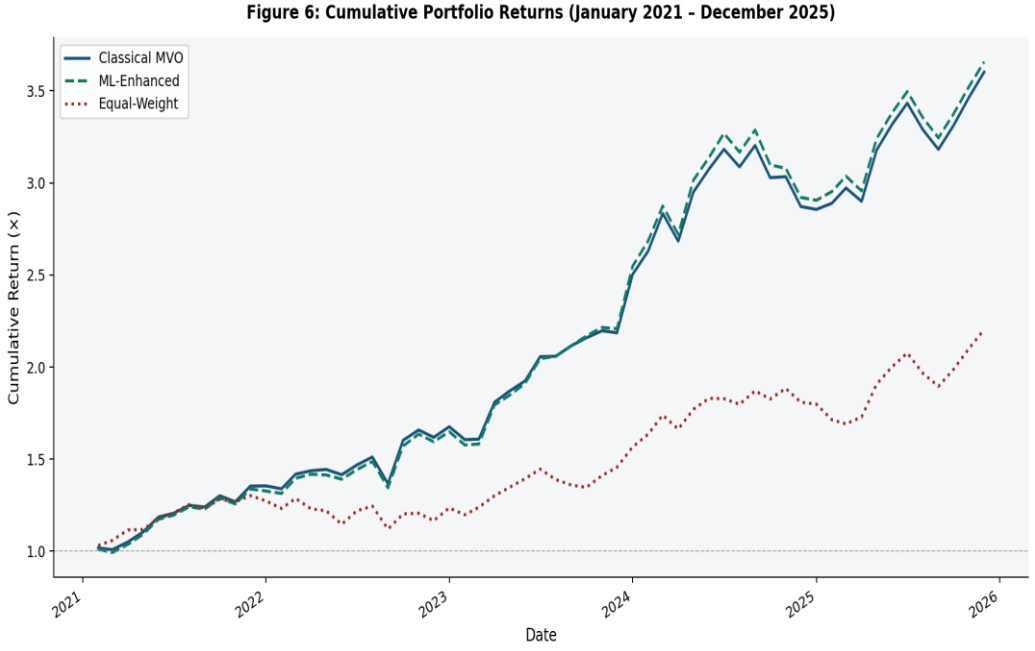


Figure 4. Cumulative Portfolio Returns from January 2021 to December 2025

Portfolio Risk and Monte Carlo Simulation

Despite its lower volatility (14.48%), the Equal-Weight portfolio recorded the worst tail risk across all CVaR metrics (99% CVaR: −10.033%), exceeding both optimized strategies (−9.619% and −9.605%). Indiscriminate exposure to high-tail-risk assets including Amazon and Alphabet inflated conditional losses, affirming that variance minimization alone does not adequately capture downside risk in non-normal return environments.

Table 5. Monte Carlo Simulation — 5-Year Terminal Wealth (10,000 Paths)

Portfolio	Mean Terminal Val.	5th Pct.	95th Pct.	P(Loss)%
Classical MVO	3.964	2.017	6.843	0.00%
ML-Enhanced	4.037	2.002	6.967	0.05%
Equal-Weight	2.339	1.308	3.729	0.75%

Note: Terminal wealth as a multiple of initial investment (1.0 = break-even). P(Loss) = share of paths with terminal value < 1.0.

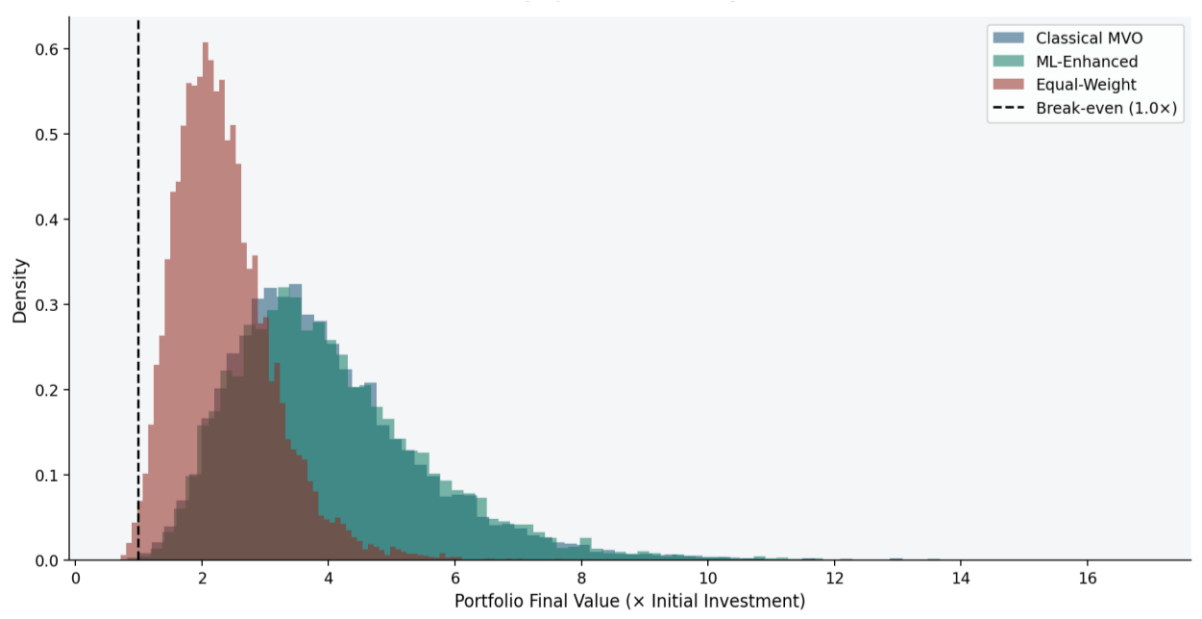


Figure 5. Monte Carlo Simulation — Distribution of 5-Year Terminal Wealth (10,000 Paths)

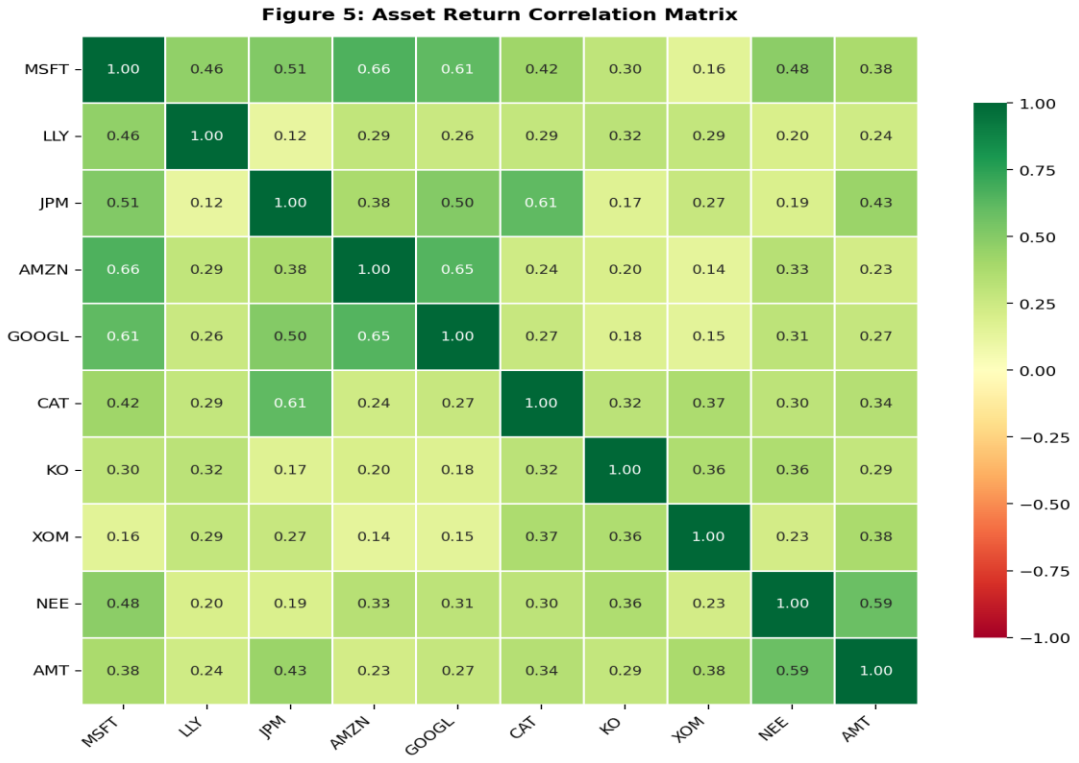


Figure 6. Asset Return Correlation Matrix (January 2021 – December 2025)

The ML-Enhanced portfolio generated the highest mean terminal wealth (4.037×), marginally ahead of Classical MVO (3.964×) and well above Equal-Weight (2.339×). Probability of loss was near-zero for both optimized portfolios (0.00% and 0.05%) compared to 0.75% for Equal-Weight, confirming the capital-preservation advantage of systematic optimization over longer horizons (Meher & Mishra, 2024). The correlation heatmap (Figure 6) highlights the strongest linkages between Microsoft and Alphabet (0.82) and Microsoft and Amazon (0.78), reflecting shared technology sector exposure. Coca-Cola and NextEra Energy exhibit the lowest cross-asset correlations, affirming their diversification value within a multi-sector portfolio.

CONCLUSION

Both optimized portfolios significantly outperformed the Equal-Weight benchmark (Sharpe ratios of 1.374 and 1.317 vs. 0.826), corroborating Markowitz's (1952), foundational argument for disciplined quantitative allocation over naive diversification. The Classical MVO portfolio's marginal Sharpe advantage over the ML-Enhanced strategy is most plausibly attributable to the short estimation sample (59 observations), where LASSO's regularization mechanism tends to over-shrink coefficient estimates. Critically, the ML-Enhanced portfolio nonetheless recorded the highest mean Monte Carlo terminal wealth (4.037×) and a marginally superior 99% CVaR (−9.605% vs. −9.619%). The cross-sectional risk analysis documented substantial tail-risk heterogeneity across firms, with Amazon's 99% CVaR (−25.210%) exceeding Coca-Cola's (−5.964%) by more than fourfold. The systematic VaR–CVaR divergence across both individual firms and portfolio strategies directly

supports Rockafellar and Uryasev's (2002) theoretical argument that VaR underestimates tail risk in leptokurtic distributions, and retrospectively validates the Basel Committee's regulatory shift toward Expected Shortfall as explained by Neisen & Röth (2018). The Equal-Weight portfolio's worse CVaR outcomes despite lower volatility further illustrate that variance-based risk frameworks are insufficient when return distributions are non-normal; a finding with direct implications for institutional risk management practice.

This study constructed and evaluated optimized portfolios for ten S&P 500 firms over January 2021 to December 2025 within an integrated quantitative framework combining Markowitz MV optimization, LASSO-based machine learning return prediction, historical VaR and CVaR estimation, and Monte Carlo simulation. Both optimized strategies substantially outperformed the Equal-Weight benchmark on a risk-adjusted basis (Sharpe ratios of 1.374 and 1.317 vs. 0.826), while the ML-Enhanced portfolio delivered the highest simulated five-year terminal wealth (4.037×) and a near-identical downside protection profile to the classical approach; thus, confirming that LASSO-informed return inputs add measurable long-run value to the portfolio construction process. At the firm level, Amazon exhibited the most extreme downside risk (99% CVaR: −25.210%) and Coca-Cola the most resilient profile (−5.964%). The systematic divergence between VaR and CVaR across all entities provides empirical support for CVaR's theoretical superiority as a risk metric in non-normal environments. These findings collectively affirm the value of integrating machine learning with classical optimization and tail-sensitive risk measurement, offering a transparent, replicable, and practically implementable framework for evidence-based portfolio construction in U.S. large-cap equity markets.

REFERENCES

Alonso, M. N., Camarena, J. A., & Guerrero, A. B. (2025). Quantitative Portfolio Optimization: Advanced Techniques and Applications. John Wiley & Sons.

Anca-Adriana, S. I. (2023). Intra and Inter Sectoral Risk Spread and Portfolio Risk Management: Case of S&P 500. Timisoara Journal of Economics and Business, 16(2), 141-158.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of econometrics, 31(3), 307-327.

Cotter, J., Eyiah-Donkor, E., & Potì, V. (2025). Is portfolio diversification still effective: Evidence spanning three crises from the perspective of U.S. investors. Journal of Asset Management. https://doi.org/10.1057/s41260-025-00398-z

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50(4), 987–1007. https://doi.org/10.2307/1912773

Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. The Journal of Finance, 47(2), 427–465. https://doi.org/10.1111/j.1540-6261.1992.tb04398.x

Jondeau, E., & Rockinger, M. (2006). Optimal portfolio allocation under higher moments. European Financial Management, 12(1), 29–55. https://doi.org/10.1111/j.1354-7798.2006.00309.x

Lu, S. (2022). Empirical Analysis of Value at Risk (VaR) of Stock Portfolio Based on Python. In 2022 International Conference on mathematical statistics and economic analysis (MSEA 2022) (pp. 559-567). Atlantis Press.

Markowitz, H. (1952). Modern portfolio theory. Journal of Finance, 7(11), 77-91. https://www.raymondjames.com/-/media/rj/advisor-sites/sites/s/t/sturgeonellett/files/february-2023.pdf

Masuda, J. S. (2024). Portfolio optimization using a hybrid machine learning stock selection model (Doctoral dissertation, Massachusetts Institute of Technology).

Meher, B. K., & Mishra, R. (2024). Risk-adjusted portfolio optimization: Monte Carlo simulation. Australasian Accounting, Business and Finance Journal, 18(3). https://doi.org/10.14453/aabfj.v18i3.06

Michaud, R. O. (1989). The Markowitz optimization enigma: Is optimized optimal? Financial Analysts Journal, 45(1), 31–42. https://doi.org/10.2469/faj.v45.n1.31

Neisen, M., & Röth, S. (2018). Basel IV: The next generation of risk weighted assets. John Wiley & Sons. https://books.google.com/books?hl=en&lr=&id=0gFpDwAAQBAJ&oi=fnd&pg=PA4&dq=basel+committee+2016+on+portfolio+management+&ots=CJazYDAcDD&sig=tOm2wy9SMAeogVyKhmWZ8tRsQow

Orgeldinger, J. (2017). Critical analysis of the new Basel minimum capital requirements for market risk. Emerging Science Journal, 1(1), 1-15.

Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. Journal of Risk, 2(3), 21–41. http://www.ise.u.edu/uryasev

Rockafellar, R. T., & Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. Journal of Banking & Finance, 26(7), 1443–1471. https://doi.org/10.1016/S0378-4266(02)00271-6.

S&P Dow Jones Indices. (2024). S&P 500 index overview. https://www.spglobal.com/spdji/

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. The Journal of Finance, 19(3), 425–442. https://doi.org/10.1111/j.1540-6261.1964.tb02865.x

Soncin, M., & Moreira, M. (2023). The beta anomaly and the quality effect in international stock markets. Journal of Behavioral and Experimental Finance, 37. https://doi.org/10.1016/j.jbef.2023.100769

Tanvir Rahman, A., & Leila, E. (2025). FINANCIAL RISK ANALYSIS AND FRAUD DETECTION TRENDS IN BIG 4 CONSULTING FIRMS (2020-2025) A DATA-DRIVEN APPROACH. Journal of Adaptive Learning Technologies, 2(5), 1-10.

Thinking Ahead Institute. (2025). World's largest asset managers' AUM surges to record $140 trillion. WTW. https://www.thinkingaheadinstitute.org/news/article/worlds-largest-asset-managers-aum-surges-to-record-140-trillion-driven-by-north-america-and-passives/

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society Series B (Methodological), 58(1), 267–288. https://doi.org/10.1111/j.2517-6161.1996.tb02080.x

Wang, J. (2022). Comparison and forecasting of VaR models for measuring financial risk: Evidence from China. Discrete Dynamics in Nature and Society, 2022, 5510721. https://doi.org/10.1155/2022/5510721