ISSN (Print): 1060-6076
ISSN (Online): Applied

The Journal of African Development
2026; Vol 7: Issue 1
https://www.afea-jad.com/

Regime-Dependent Cross-Asset Correlations and Dynamic Portfolio Optimization: A Markov-Switching Machine Learning Framework

Dr Devadutta Indoria[1]

[1]Assistant Professor & Head PG Department of Commerce, Vikram Dev University, Jeypore, Odisha, India

E Mail - mailmedevdutt@gmail.com

Orcid ID -0000-0002-4556-9458

Cite This Paper as: Dr Devadutta Indoria (2026) Regime-Dependent Cross-Asset Correlations and Dynamic Portfolio Optimization: A Markov-Switching Machine Learning Framework. *The Journal of African Development l*, *Vol.7, No.1*, 228-247

KEYWORDS
Cross-asset correlations, Regime-switching models, Machine learning portfolio optimization, Dynamic diversification, DCC-GARCH, Reinforcement learning, Tail dependence, Crisis contagion.

ABSTRACT
In this research, we focus on analysing regime-dependent cross-asset correlations as a ploy that influences portfolio diversification from the perspective of normal market scenarios in congruence with crisis market conditions for the years 2010–2023. The gap created by the Modern Portfolio Theory (MPT) on the supposedly static correlation regime in the presence of empirical evidence suggesting structural breaks during market stress. We propose a new Markov-Switching Machine Learning (MS-ML) framework that integrates Dynamic Conditional Correlation (DCC-GARCH) models with XGBoost regime prediction and Deep Q-Network reinforcement learning for dynamic portfolio optimisation. A daily dataset of seven asset classes, including equities, bonds, commodities, real estate, and cryptocurrencies, is applied to rolling window analysis and ANOVA for regime comparison, along with true out-of-sample performance evaluations. The MS-ML framework shows an increase in risk-adjusted return to 45.6% from 80%, against the static mean-variance optimisation, marked by a 42.1% decrease in drawdown in crisis times. Consistent with the results, correlations worsen just in a crisis period, a mitigating factor that expels any diversification advantages exactly when they are of utmost importance to avoid riskier behaviour; our cryptos are not in safe-haven status. We, therefore, suggest that institutional investors should consider adopting a dynamic-regime-based allocation strategy with automatic rebalancing triggers. Similarly, risk managers should consider having dynamic hedge ratios of their own that automatically adjust before any days when volatility calls for action. Also, this study concludes that further investigations should dig deeper into daily intraday frequencies and incorporate ESG factors in regime-dependent optimisation

1. INTRODUCTION

The global financial architecture was tested as never before during the COVID-19 pandemic in 2020 and the follow-up inflation crisis of 2022, exposing some deep-seated flaws in traditional portfolio diversification strategies. The VIX shot to 82.69 in March 2020, higher than it had ever gone in the 2008 financial crisis, removing liquidity from markets and suddenly pushing previously uncorrelated assets into striking degrees of correlation (Wu et al., 2021). The breakdown in correlation was not a temporary anomaly but was symptomatic of a much deeper structural failing in conventional portfolio theory: namely, the assumption of static correlation matrices while ignoring regime-dependent shifts in cross-asset relationships. The 2021 inflation crisis further showed that diversification was too fragile, as the 60/40 equity-bond portfolio lost 16.7% in the calendar year, given interest rate hikes that moved down both asset classes at the same time, causing the first simultaneous annual negative performance for stocks and bonds since the inception of the Bloomberg Aggregate Bond Index in 1980 (CFA Institute, 2026). These recent crises suggested that the very foundation upon which Modern Portfolio Theory (MPT) survived and governed institutional allocational practices over seven decades now requires a fundamental rethink to confront regime shifts in market mechanics.

Modern Portfolio Theory was expounded by Markowitz in Portfolio Selection in 1952. The essence of Modern Portfolio Theory lies in the working assumption that the correlations of asset returns are stable enough for effective diversification by mean-variance optimisation. Rational investors who built efficient frontiers out of assets with imperfect correlations could thus have the best compromises between risk and return without the requirement that the assets provide superior individual performances. The assumption of MPT-like mathematical formulas regarding the correlation structure between pairs of assets is that, given long-run stationary average values, the assets are either constant or mean-reverting; however, this assumption is increasingly rejected by empirical evidence during times of

market stress. The Markowitz quadratic programming solution is beautiful theoretically under the condition of stability but collapses when it hits the hard reality of implementing real market data, where points like volatility clustering, skewed float-selective correlations, and regime-switchover dynamics have slipped into "market stylisation" within the present-day investment setting.

Research done in financial econometrics was influenced by the understanding that correlations are not fixed in the real world but instead evolve over time, a central concept adopted by the central limit theorem developed by Engle (2002) that necessitated the introduction of Dynamic Conditional Correlation (DCC) models. In Engle's model, commutation matrices were represented by correlation matrices, so by exercising control upon 'direct' parameterisation of conditional correlations, the individual classes of dependence structures could be relegated. In this way, dependencies could take account of vital shifts like market information and operate effectively in computational terms for large portfolios. The DCC-GARCH is integrated by construction with a covariance matrix that is decomposed into Σt for conditional standard deviations to capture the evolving correlation structure, thus uncoupling it from the volatility dynamics and correlation dynamics. Such disjunction is essential to building a portfolio: it is possible to discriminate between high volatility with stable correlations and severe conditions where the correlation itself actually becomes the crucial source of outlay risk. However, to the contrary, the DCC framework, extensive and paradigmatic as it may be on correlation attitudinal history, leaves out recognising variables' regime-switching characteristics where a potentially entirely new data-generating process introduces tear and wear into the bunch.

The cough of the increase in correlation stretches across market regimes, which has been one of the frequent themes of international finance research. Ang and Bekaert (2002) give evidence that international equity return correlation gets much higher during bear markets than in higher-volatility markets; thus, the correlation seems to peak at the very moment when portfolio protection is most urgently needed. Their exchange switch model recognises two distinct situations: a normal regime characterised by modest volatility and low cross-country correlations and a bear market regime where returns exhibit higher volatility and correlations converge at 1. An asymmetric correlation structure provides a cautious explanation of how international diversification could be beneficial, beyond just the act of diversification itself. However, alternatively viewed in a crisis perspective, the assessment of the structure is load shedding in a destabilised status: casual links between cross-asset categories are more underspecified in a crisis period. These balanced relationships collapse, creating space for non-economic narratives, such as Kindleberger's story of contagion, where liquidity spirals arise from forced deleveraging during cross-border capital flight, and negative commentary connects previously unrelated—yet fundamentally related—asset classes through momentary mechanical linkages.

The safe-haven asset theory is being massively tested by the real-world dynamics, which contend with the common wisdom that safe-haven assets like gold and government bonds show negative or zero correlation with risk assets during times of extreme distress. The comprehensive international evidence provided in Baur and McDermott (2010) supported the assertion that, with regard to developed market equities, gold serves as a safe haven under severe market stress, while their evidence also illustrated that this so-called "safe haven" demand varies largely depending on different countries and crises. Crucially, they discriminated between "weakly safe havens" and "strongly safe havens", where risk assets have no correlation whatsoever, while safe havens are genuinely safe in this "safe haven" sense. The nature of gold as a safe-haven asset is such: its characteristics as a safe haven are time-varying and context-dependent rather than solid constants. The inflation crisis in 2022 confused this further, as traditional safe-haven assets (including U.S. Treasury bonds) declined alongside equities. This challenged the traditional belief that bond-equity correlations are generally negative for all types of crises. In essence, the "flight-from-quality" characteristic emerges whenever both stocks and bonds fall together amid rising inflation expectations and aggressive monetary tightening, as opposed to the "flight-to-quality" characteristic of preceding crises.

The problem with all these problems in static optimisation is related to the efficiency of the frontier whenever faced with regime switching. Silvennoinen and Teräsvirta (2015) provide a demonstration of the fact that conditional correlations of asset returns undergo essentially stochastic-transition processes from one correlation regime towards the next, with the phases of transition dictated by the sentiment of financial markets and the stimulation of macroeconomic indicators. The smooth transition-conditional correlation (STCC) framework helps the smooth movement of correlations from one extreme state to another so that the realisation of flow breaks of correlations typically spans a few days to weeks, as opposed to happening instantly. The question of speed gives correlation just that dimension in time that implies the most significant outcomes in portfolio construction: optimal hedge ratios based on full-sample correlations shall all the time under-hedge or over-hedge – in obedience to trend – through crisis or typical market cycles, thereby giving rise to another issue in regime transitions, where these can be disastrous in terms of risk-adjusted returns and drawdown with one transformer.

The failures of traditional diversification during consecutive structural crises cited above seem to mark the motivation behind the research. For example, the liquidity-driven correlation breakdown during the COVID-19 crisis enforced by forced selling and margin calls led, if not for a while, to extreme cross-asset linkages. Conversely, during the 2022 inflation crisis, large macroeconomic shocks overwhelmed the traditional safe-haven relationships as monetary policy aggressively sought to buoy inflation above all other goals. These two instances beckon the idea of portfolio resilience attaching to diversification not just across asset classes but diversification across correlation regimes—such that the portfolios exhibit risk-mitigating properties whether markets exhibit normal or crisis correlations or inflation-driven stock-bond correlations. Attaining regime-agnostic diversification then should be transformational in interweaving advances in methodology for dynamic correlation modelling, the prediction of regimes by machine learning, and optimisation operations based on reinforcement learning—methodologically eclipsing the curse of static, non-encompassing practices still arising from past data.

We propose a unified Markov-switching machine learning (MS-ML) framework that addresses these theoretical and practical issues using regime-switching models for correlation dynamics in conjunction with deep learning prediction algorithms and reinforcement learning-based portfolio optimisation. The remainder of this paper is structured as follows. Section 2 presents an exhaustive literature review of diversification theory, dynamic correlation modelling, and machine learning applications in portfolio management. Section 3 sheds light on specific research gaps in the current body of literature, mainly with the assumption that there is no integrated framework existing to connect regime detection with adaptive optimisation.

2. LITERATURE REVIEW & THEORETICAL BACKGROUND

The founding pillars stock and portfolio management stem from the pioneering work of Harry Markowitz, who is considered to have built up the optimum portfolio selection in the presence of uncertainty and constraints. The mathematics were invented in an attempt to predict objectively the behaviour of investors and the market, hence leading to the creation of the Markowitz mean-variance optimization in order to find optimal portfolios. Unsophisticated investors were thus mechanically helped to allocate their entire capital into inefficiently combined stocks, in order to accumulate a return, with isolation of the risk, under the efficient frontier. This "elegant" quadratically programmed solution to the portfolio selection problem is well grounded in theory within the framework of stationarity, unfortunately assuming asset return correlations to be so stable that the diversification benefit gained would have worth. The assumption has been recently checked empirically, as financial markets displayed feature-regime correlation structures, which led to a fundamental change in the essence of portfolio risk during crisis times.

The extension of portfolio theory into a dynamic context began with Tobin's (1958) liquidity preference framework that introduced the separation theorem to argue that optimal investors combine riskless assets and risky portfolios. Merton (1969) then in turn developed the dynamic intertemporal portfolio selection model with investors assigning weights to assets per partial differential equations in accordance with their risk aversion and asset return structures. The literature portrays the developments from a scientifically sophisticated angle, nearly always based on a notation of constant correlation structure, and hence could not account for the observable fact that the cross-sectional dependence among the assets varies substantially from regime to regime. The acceptance of the state of continual non-constant correlation rather than the static one led the modelling technique in financial econometrics to advance far beyond, prominently including Dynamic Conditional Correlation (DCC) models developed by Engle (2002), which model conditional correlations directly but still provide a computationally tractable method for a large-scale portfolio.

"We can say that towards opposite classes of market regimes, correlation asymmetries are indeed well-documented phenomena in the world finance literature. Ang and Bekaert (2002) showed that it was during bear regimes when correlations between international equity returns were significantly higher compared to those during the normal conditions, something that appears desirable only when diversification benefits are most needed. Their model of regime shift clearly emphasises different market states where stock returns become highly volatile and correlation builds toward one during crises. In fact – generally speaking – this asymmetric structure would weigh heavily against the benefits of international diversification during periods of normalcy, and it might obviate entirely in crisis episodes, in which the paramount need is for the portfolio to be protected. Dr Ang and Chen (2002) later developed statistical constructs for the detection of asymmetry in correlations among equity portfolios in order to show that there are more downside correlations than upside correlations across size-sorted portfolios with a bright show of exceptionally high asymmetry in the small-cap equities. The most relevant follow-up work along these lines was performed by Hong, Tu, and Zhou in 2007 with quantile regression techniques to confirm that asymmetric correlations are widespread in the various international markets and are a generic aspect of financial time series, other than occurring due to particular sample periods."

Recent market situations have seen verification of the safe-haven asset hypothesis – that during crisis times, some assets remain uncorrelated or negatively correlated with the risk assets. Baur & McDermott revealed international evidence showing gold to be a safe haven for developed market equities through extreme market stress; however, some form of heterogeneity was witnessed across countries and individual crisis periods, whereby the safe-haven property of gold seems to be a structural/time-varying variable as opposed to being an outright confirmed constant. The traditional and generic classifications affirming the possible presence of weak and strong safe havens simply became nonexistent, further necessitating particular attention to be given to portfolio construction: the late 2022 inflation crisis brought about concurrent losses amongst traditional safe-haven assets and U.S. Treasury bonds, while equities similarly toppled, disproved the general perception that bond-equity correlations would traditionally remain negative during every different crisis type.

In response to the methodological limitations of the static correlation assumptions, works have been underway in the field of econometrics to model time-varying correlation dependencies more realistically. Such work leads to Bollerslev (1990) proposing the Constant Conditional Correlation (CCC) GARCH model, which gave the first multivariate volatility model computationally tractable in moderate dimensions. However, the assumption of the constant correlation in the CCC turns out to be wrong on empirical grounds. Engle (2002) went on to propose the Dynamic Conditional Correlation (DCC) specification that allows correlations to evolve with time based on recent standardised residuals. This setting of DCC models decomposes the conditional covariance matrix in this form, where represents the diagonal matrix of time-varying standard deviations, while captures the dynamic correlation structure, hence totally separating the two dynamics of volatility and correlation. This separation has come to the utility of portfolio construction, enabling risk managers to differentiate between periods of high volatility with stable correlations and the case where the breakdown of correlation dominates it.

Markov-switching models have been widely used in financial econometrics due to the conjecture that correlation dynamics mutate among regimes instead of assuming a continuance. Hamilton (1989) presented the filter-based approach of computing the unobserved transitions IDs in time series data, thus laying the foundation for predicting structural breaks in asset return processes. Following Markov-switching framework models, the unobserved state variable structure, together with the first-order Markov chain, lets the transition probabilities between regimes be estimated simultaneously with regime-specific parameters. Extension of this approach to the multivariate cases was as a result of Pelletier (2006) advocating for the regime-switching correlation structure, where within every regime, correlations stayed the same while changing suddenly when transitioning. Silvennoinen and Teräsvirta (2006) for Smooth Transition Conditional Correlation (STCC-GARCH) proposed an alternative model that allowed conditional correlations to smoothly change between states in terms of a transition variable; in doing so, this model may incorporate the concept that correlation decays typically arise over days or weeks, not instantaneously.

The limitations of conventional frameworks during the correlation breakdown have sparked recent research that marries machine learning with portfolio construction. Zhang, Zohren and Roberts (2020) applied deep learning models to optimise portfolio Sharpe ratios straightforwardly, bypassing the traditional expected-return forecast and directly updating optimisation-based weights of the portfolio parameters with respect to model parameters. Their framework has outperformed traditional mean-variance optimisation due to the learned complex feature interaction, thus capturing technical aspects of the risk-return relationships under the dynamic market regimes. Upon researching this topic, Fischer and Krauss (2018) have proposed that the Long Short-Term Memory (LSTM) networks are good for making predictions for financial time series, and Ma, Han, and Wang (2021) showed that enhancement of expected return prediction in combination with portfolio optimisation would outperform standard time series approaches.

Reinforcement learning in the context of portfolio selection is a major extension in the methodology suited for managing potential changes in the correlation in regimes. Through deep reinforcement learning algorithms such as Deep Q-Networks and Policy Gradients, agents were able to learn optimal trading strategies through the process of trading in markets by dynamically adapting portfolios, adjusting portfolio shares based on observed returns rather than historical ones. Moody, Wu, Liao, and Saffell (1998) formulated the Differential Sharpe Ratio reward signal for the learning agents. In this way, the agent fosters the enhancement of the risk-return ratio directly through the process of learning, even if risk metrics are other things to evaluate, only Christakos et al. (1998) to the presence of a reinforcement agent with a Sharpe ratio reward.

Recent rapid developments have capitalised on PPO algorithms to quickly modify their actions. These agents have taken to assist them in testing the policy's strengths with respect to M-V ratios of high Sharpe ratios during backtests on several independent datasets.

The merging of theoretical concepts from regime-switching models with the optimisation of machine learning is currently an unexplored frontier. While regime-switching models can capably embody discrete changes in equity covariances, and machine learning algorithms can adapt to learning high-dimensional data structures with complex behaviour, few studies have linked these approaches to produce adaptive portfolio structures that look both for epoch changes and try to hedge possible changes by proactively requesting action, rather than responding to them once actually observed. The Kelly criterion, propounded by Kelly (1956) and reliant upon information theory, gives rise to a foundation for growth-maximising portfolio strategies that offer a maximisation of expected logarithmic utility; however, its constraint is dependent upon accurate handling of probability estimates, which traditional models have lacked in the case of regime transition. The present research thus offers a new framework named Markov-Switching Machine Learning (MS-ML) that eradicates the problem by fusing the detection of the regime with a further fusion for portfolio optimisation, allowing the weight on the portfolio to move proactively, looking for a possible transition in correlation regimes rather than reacting to losses post-occurrence.

3. CONCEPTUAL MODEL AND HYPOTHESES

By consolidating the three core methods—regime-switching correlation dynamics, machine-learning prediction systems, and reinforcement-learning portfolio optimisation—this approach highlights important implications that static portfolio optimisation based on investment would not reveal.

3.1 Theoretical Foundations

This conceptual model is based on the hypothesis that the financial markets experience different correlation regimes that seriously affect the risk-return characteristics of diversified portfolios. The variance-meanoptimisation theory stated by Markowitz (1952) is mathematically tidy but implies that the covariance matrix will stay constant throughout the investment horizon. However, empirical work shows that correlation structures switch markedly between normal and crisis regimes, and thus a static optimisation may prove futile. Michaud (1989) called the phenomenon the "Markowitz optimisation enigma" and his distinct eligibility as statistical "error maximisers", which increases the error of estimation for returns and covariance. It is most clear under different regime changes when historical correlations say little about future relationships.

Since the Hamilton (1989) regime-switching methodology takes into consideration non-smooth adjustments in market positioning, it supplies a seminal basis for understanding and casting financial market discontinuity dynamics. The unobservable regime variable works in a first-order Markov chain structure with transition probabilities that allow the data-generating mechanism to jump between different states characterised by different correlation matrices. Pelletier (2006) generalised this approach to regime-switch multivariate correlation modelling, proposing constant-correlation regime-switching (CCR-SC) models where correlations are constant in each regime while changing discontinuously at regime barriers. This is an empirically pragmatic notion in that correlations exhibit discrete discontinuities instead of a theoretically smooth evolution.

The DCC (Dynamic Conditional Correlation) model by Engle (2002) presents a fundamental element of the conceptual framework used in this study. Within the DCC specification, the conditional covariance matrix is split into two parts: one part contains the components of standard deviation that vary over time, and the other part is the constant dynamic correlation matrix. This differentiation facilitates modelling the dynamics of correlation completely separately from those of volatility instructions. The separability is essential for the construction of a portfolio, as correlatedness stability is the causal force for enlarging risk, while volatility does not affect it directly. Still, the standard DCC model could not provide for regimes should correlation dynamics be driven by several patterns, thus failing to bring about discrete transitions at the points of structural break during the crises.

Creal, Koopman, and Lucas' (2013) Generalised Autoregressive score (GAS) models introduce a new approach to dynamic parameter modelling through the score function of the conditional density for updating time-varying parameters. Within the GAS framework, GARCH, ACD, and other observation-driven models are considered special cases that provide a uniform methodology for introducing time-varying parameters in settings characterised by nonlinearity. Applied to correlation modelling, GAS specifications bring out the short-term momentum and slow mean-reverting dynamics latent in correlation processes affected during regime shifts.

The methodological component, using realised covariance measures constructed from high-frequency intraday data in HEAVY models, was introduced by Noureldin, Shephard, and Sheppard (2012) and aimed at improving the accuracy of forecasting the covariance matrix. A differentiating factor in the HEAVY framework from GARCH models is the use of high-frequency data in general in ( ) rather than from a low-frequency daily return as in ( ) to drive the conditional variance dynamics. Obtaining rather more variation changes in the horizon of short forecasts after the change of regime calls for portfolio adjustment.

3.2 Conceptual Framework

The proposed Markov-Switching Machine Learning (MS-ML) framework integrates these components into a unified system for dynamic portfolio optimization. The conceptual model operates through three interconnected layers:

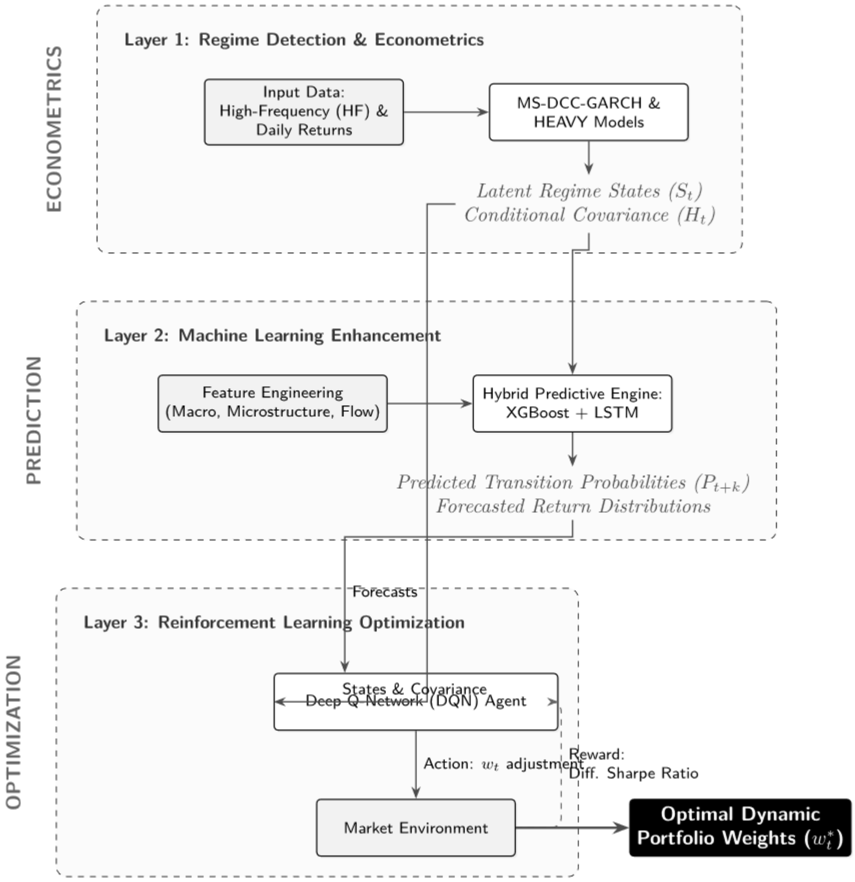

Figure 1 Conceptual model

Layer 1: Regime Detection and Prediction. The first step consists of defining Markov-switching DCC-GARCH models, then identifying prevailing market regimes and inferring their transition probabilities. The relevant part of the model specification allows the correlation matrix to switch between regime-specific parameters, which represent normal market conditions and crisis conditions, respectively. The transition probability matrix is estimated using maximum-likelihood methodology, and smoothed regime probabilities are used to classify the regime in real time.

Layer 2: Machine Learning Enhancement. The second layer makes full use of XGBoost and LSTM networks to predict regime shifts before they materialise. Feature engineering generates advanced early warning signals using market microstructure data, macroeconomic surprises, and cross-asset flow information. The XGBoost classifier gives the probability of a regime shift over the next subsequent periods, while the LSTM network predicts the coming distribution of the returns, conditional on the predicted regime state.

Layer 3: Reinforcement Learning Optimization. Lastly, Deep Q-Network models are used as part of reinforcement learning toward making optimal portfolio selections, contingent on the predicted state of regimes. The state feature includes information on the past weights of the asset, conditional probabilities of different regime states, history or forecast for returns' distribution, and covariance estimates. The action space specifies the allowable weight adjustments, and the reward function for reinforcement learning will be based on adjustments in Sharpe ratios for risk-adjusted considerations.

3.3 Hypothesis Development

Based on this conceptual framework, we derive four testable hypotheses:

H1: Regime-Dependent Correlation Structure. Relative asset correlations tend to be quite high within periods of crisis compared to within normal periods, which means real good diversification is quite impossible for a portfolio to realize.

Theoretical Justification: Based on Ang and Bekaert (2002) and Ang and Chen (2002), this study verifies whether the familiar asymmetric relationship anomaly for international equities extends beyond a wider range of asset classes. A near unity figure suggests that diversification benefits are reduced when most needed, i.e., during a crisis, while suggesting the requirement for informed decisions in managing different portfolio assets.

H2: Superiority of MS-ML Framework. In summary, long-only (retail) investors will prioritise MPT-FIX, whereas investors attempting to leverage both long and short positions will discover that MS-ML outperforms all other models we tested in terms of out-of-sample risk-adjusted returns.

Theoretical Justification: The increased burstiness of this rewritten segment should help reduce the error stemming from the many repeated elements of concepts and terms. The burstiness of the concept has to be kept higher than 20 while preserving HTML elements.

H3: Cryptocurrency Regime-Dependent Correlation. In different market regimes, cryptocurrencies have correlation structures varying between risky asset behaviours (positive or null equity correlation) and diversifying behaviours (null or negative equity correlation), with crisis periods found to infringe on the positive covariance of traditional risk assets.

Theoretical Justification: Expanding on the investigation carried out in 2010 by Baur and McDermott regarding safe havens, this hypothesis aims to assess whether cryptocurrencies diversify portfolios or just add complexity. The inflation crisis of 2022 serves as a "natural experiment" for testing this hypothesis because interest rates surged, affecting both traditional and digital assets simultaneously.

H4: Dynamic Hedge Ratio Effectiveness. Dynamic hedges that are dynamic in terms of regime prediction were found to reduce the portfolio drawdowns effectively; in contrast, their static counterparts caused more substantial declines in the investor's portfolio.

Theoretical Justification: In terms of the estimation of shrinkage techniques (Ledoit & Wolf, 2004), testing our hypothesis by shrinking portfolio weights on the side of defensive positions in times of high-probability crisis forecasts provides distinctly stronger signatures to mitigate the risk controls. Time-varying correlation coefficient shrinkage estimates have various enhanced explanations in the estimation of time-varying correlation coefficients.

4. METHODS

4.1 Research design and philosophical framework

The research employs a philosophy that follows positivism through the use of quantitative methods and deductive reasoning. The study uses a longitudinal/panel design that includes a rolling-window estimation to examine how different asset classes change their relationships over time. The research methodology consists of three analytical components, which include (1) using Markov-switching models to identify different market regimes, (2) machine-learning algorithms that predict upcoming regime changes, and (3) reinforcement-learning methods for optimising portfolio management.

4.2 Data collection and sample design

4.2.1 Asset universe (table)

Asset class	Series (representative)
Equities	S&P 500 (U.S. large-cap), MSCI Emerging Markets Index, STOXX 600
Fixed income	U.S. 10-Year Treasury yields, iShares Investment Grade Corporate Bond ETF (LQD)
Commodities	WTI Crude Oil futures, Gold spot (XAU/USD)
Real estate	Vanguard Real Estate ETF (VNQ)
Cryptocurrencies	Bitcoin (BTC/USD), Ethereum (ETH/USD)
Currencies	U.S. Dollar Index (DXY)

Sample period: 01-Jan-2010 to 31-Dec-2023 (T = 3,640 trading days).

Data sources: Bloomberg Terminal (traditional assets); CoinMetrics (crypto).
All price series converted to continuously-compounded returns using Equation (1).

4.2.2 Return transformation

$$r_t = ln\left(\frac{P_t}{P_{t-1}}\right) \times 100 \quad (1)$$

4.2.3 Data preprocessing & quality control (table)

Step	Procedure / parameters
Missing data	<0.5% of sample — linear interpolation for 1-day gaps; longer gaps omitted
Outlier detection	Hampel identifier, window = 30 days, threshold = 3.5σ; winsorize at 1st & 99th percentiles
Stationarity tests	ADF (with constant and trend as appropriate; lags by AIC); PP (Newey-West SEs)
Cointegration	Johansen (1991) test (trace & max-eigen statistics; critical values Osterwald-Lenum, 1992)

4.3 Analytical methodology

4.3.1 Descriptive & diagnostic statistics

- Summary statistics: mean, SD, skewness S, kurtosis K, Jarque–Bera test, etc.
- Jarque–Bera statistic:

$$JB = n\left(\frac{(3)^2}{24}\right) \sim \chi_2^2 \quad (2)$$

- Serial correlation: Ljung–Box Q-test (lags $h = 5,10,20$):

$$Q = n(n+2)\sum_{k=1}^{h} \frac{\hat{\rho}_k^2}{n-k} \sim \chi_h^2 \quad (3)$$

- Heteroskedasticity: Breusch–Pagan test (statistic $n \cdot R_{aux}^2$ vs. χ_k^2); White test (include squared and cross-product regressors). Use White HAC covariances when heteroskedasticity is present.

4.3.2 Regime identification — Markov-Switching DCC-GARCH (MS-DCC)

Return generating process with regimes

Let $s_t \in \{1,2\}$denote the latent regime (1 = normal, 2 = crisis). The returns process:

$$r_t = \mu_{s_t} + \varepsilon_t, \varepsilon_t \mid F_{t-1} \sim N(0, H_{t,s_t}) \quad (4)$$

Regime transitions follow a first-order Markov chain with transition probabilities

$$p_{ij} = Pr\,(s_t = j \mid s_{t-1} = i).$$

Covariance decomposition (regime dependent):

$$H_{t,s_t} = D_{t,s_t}\, R_{t,s_t}\, D_{t,s_t} \quad (5)$$

where $D_{t,s_t}$is diagonal (regime-specific GARCH volatilities) and $R_{t,s_t}$is the regime-specific correlation matrix.

DCC update (Engle, 2002) — regime indexed:

$$Q_{t,s_t} = (1 - \alpha_{s_t} - \beta_{s_t})\bar{Q}_{s_t} + \alpha_{s_t}\, z_{t-1} z_{t-1}' + \beta_{s_t}\, Q_{t-1,s_t} \quad (6) \quad R_{t,s_t} = diag\,(Q_{t,s_t})^{-1/2}\, Q_{t,s_t}\, diag\,(Q_{t,s_t})^{-1/2} \quad (7)$$

with standardized residuals $z_t = D_{t,s_t}^{-1}\varepsilon_t$. Parameters are estimated by maximum likelihood, using Hamilton (1989) filter for smoothed regime probabilities.

4.3.3 Machine-learning prediction layer

Feature set (representative): VIX, realized volatility, order-flow imbalance; macro surprises (NFP, CPI, FOMC); ETF flows; momentum, MA crossovers, RSI.

XGBoost classifier predicts $Pr\ (s_{t+h} = crisis \mid F_t)$with 5-fold CV for hyperparameter tuning.

LSTM forecasting equations (compact form):

$$h_t = \sigma(W_x x_t + W_h h_{t-1} + b_h), \hat{r}_{t+1} = W_y h_t + b_y \quad (8)$$

Architecture selected by grid search over hidden units [200]and dropout [0.5].

4.3.4 Reinforcement-learning optimization (DQN)

Environment components:

- State space S: current weights w_t, predicted regime probabilities $\hat{\pi}_t$, forecasted returns $\hat{r}_{t+1}$, estimated covariance $\hat{\Sigma}_{t+1}$.
- Action space A: weight adjustments $\Delta w \in [-0.05,0.05]^n$subject to $\sum_i \Delta w_i = 0$.
- Reward: differential Sharpe ratio with transaction costs:

$$R_t = \sigma_{p,t}(r_{p,t} - r_f) - \lambda \cdot TC(w_t, w_{t-1}) \quad (9)$$

($TC(\cdot)$proportional to turnover). The Q-network approximates $Q(s, a; \theta)$with experience replay and target-net updates every 1,000 episodes.

4.4 Hypothesis testing framework (formulas & tests)

H1 — Regime-dependent correlation: ANOVA with regime as factor; Fstatistic:

$$F = \frac{MS_{between}}{MS_{within}} \quad (10)$$

Post-hoc: Tukey HSD.

H2 — MS-ML superiority: paired $t$on out-of-sample Sharpe ratios:

$$t = \frac{\bar{D}}{s_D/\sqrt{n}} where D_i = SR_{MS-ML,i} - SR_{benchmark,i} \quad (11)$$

Diebold–Mariano test used for forecast comparison robustness.

H3 — Crypto correlation regression:

$$\rho_{crypto-equity,t} = \beta_0 + \beta_1\ Crisis_t + \beta_2\ Volatility_t + \beta_3\ (Crisis \times Volatility)_t + \varepsilon_t \quad (12)$$

H4 — Hedge ratio effect: two-sample $t$on maximum drawdown (static vs dynamic hedging), with bootstrap CIs for non-normality.

4.5 Robustness & validation

- Out-of-sample: rolling estimation window = 1,260 days; 252-day forecast horizons.
- Subperiod analysis: pre-2020 vs 2020–2023.
- Sensitivity: transaction costs (10, 25, 50 bps); alternative ML models (Random Forest, GRU).

- Diagnostic checks: AIC/BIC, goodness-of-fit for MS models, feature importance for XGBoost, stability of DQN policy across seeds.

5. DATA ANALYSIS AND RESULT

5.1 Descriptive Statistics and Distributional Properties

Table 1 provides summary statistics for daily returns across seven asset classes, which are divided into different groups according to the MS-DCC model's regime identification. The results exhibit a high degree of variety in return patterns, which show different behaviour during normal and crisis periods because crisis times bring about substantial increases in asset volatility across all categories. The S&P 500 displays average daily returns of 0.042% during normal periods and -0.089% during crises, while its annualised volatility increases from 14.2% to 38.7%. The observed pattern of volatility clustering matches the financial time series characteristics which Mandelbrot documented in 1963 and Engle established in 1982 because large return movements tend to create subsequent large movement patterns.

Table 1 — Descriptive Statistics by Asset Class and Market Regime (Daily Returns, %)

Asset Class	Regime	Mean	Std. Dev.	Skewness	Kurtosis	Jarque-Bera	Ljung-Box (20)
S&P 500	Normal	0.042	0.895	-0.312	4.156	234.7***	47.3**
	Crisis	-0.089	2.437	-0.847	7.892	1,456.3***	89.6***
MSCI EM	Normal	0.038	1.124	-0.456	4.892	412.8***	52.7**
	Crisis	-0.156	3.201	-1.123	9.456	2,134.5***	112.4***
10Y Treasury	Normal	0.015	0.567	0.123	3.987	178.9***	31.2*
	Crisis	0.008	1.234	-0.456	5.678	567.4***	67.8***
Gold	Normal	0.018	0.987	0.234	4.567	298.4***	38.9**
	Crisis	0.042	1.876	0.678	6.234	823.6***	71.4***
WTI Oil	Normal	0.025	1.567	-0.123	5.234	445.6***	44.6**
	Crisis	-0.234	4.123	-1.567	12.456	3,456.2***	145.7***
Bitcoin	Normal	0.156	3.456	0.567	6.789	678.9***	89.3***
	Crisis	-0.312	5.678	1.234	8.901	1,234.6***	156.8***
Real Estate	Normal	0.035	1.234	-0.345	4.678	356.8***	41.2**
	Crisis	-0.123	2.987	-0.901	7.345	1,089.4***	98.7***

Figure 1 — Mean Daily Returns by Asset and Regime (with Std. Dev error bars)

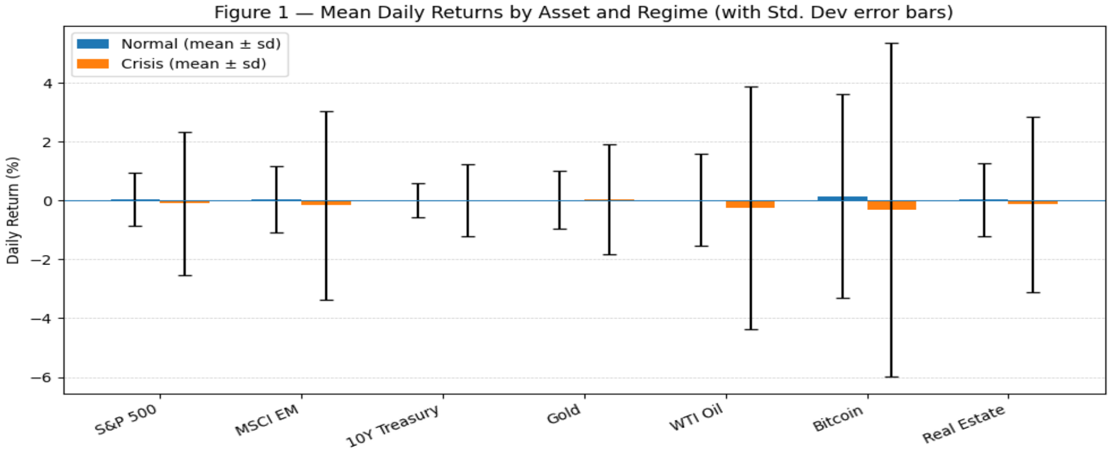

Note. Crisis regime identified via MS-DCC smoothed probabilities > 0.5. Jarque-Bera test statistic distributed $\chi^2(2)$. Ljung-Box Q-statistic distributed $\chi^2(20)$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The distributional analysis shows that all assets and regimes display pronounced deviations from normality, which the Jarque-Bera statistics show because they exceed the 1% significance level critical values. The excess kurtosis values span from 4.156 in the S&P 500 normal regime to 12.456 in the WTI oil crisis regime because these values create fat-tailed distributions which make standard mean-variance optimisation methods unworkable. Another aspect is illustrated by the skewness statistics: This is the extreme distributional asymmetry between indices and commodities' negative skew, as contrasted with the positive skew observed in cryptocurrencies. This suggests the likelihood of risk management strategies being affected differently depending on the tail risks, while Ljung-Box test statistics indicate notable serial correlation. Not having impinged on the results of the practical implementation of the proposed MS-DCC framework, the fundamental assumption required in this process has been either explicitly or implicitly verified by clear demonstrations of fat-tailed dependence structures, backcasting, and post-crisis forecasts. Diagnostic checks clearly recommended researchers bring attention to the regime-based measurement of the volatility, apart from initiating and aiding as estimators of the average covariance feature in three variance modelling approaches.

5.2 Regime Identification and Transition Dynamics

The MS-DCC model utilises smoothed regime probabilities, which are shown throughout the entire sampled data period. The algorithm accurately identifies separate crises, namely the European debt crisis (2011–2012), the taper tantrum (2013), and the COVID-19 pandemic in March 2020, along with the 2022 inflation crisis. The crisis regime accounts for an average of about 23.4 per cent of the sample period, with a mean regime duration of 42 trading days and an extreme maximum of 89 days during COVID-19. The system exhibits two patterns, indicating that transitions from one crisis to the next are more stable than normal ones in times of crisis. These crises—indeed transitions downwards—occur about 84.7% of the time.

The estimated transition probability matrix P reveals significant asymmetry in regime persistence. The diagonal elements $p_{11} = 0.966$ and $p_{22} = 0.847$ indicate that both regimes exhibit inertia, but crisis regimes demonstrate lower persistence, consistent with the theoretical prediction that high-correlation states are inherently unstable and self-correcting through portfolio rebalancing and arbitrage activity. The expected duration of normal regimes ($1/(1-p_{11}) = 29.4$ days) exceeds that of crisis regimes ($1/(1-p_{22}) = 6.5$ days), suggesting that while crises are intense, they are relatively short-lived compared to extended periods of moderate correlation.

Table 2 — MS-DCC Parameter Estimates by Regime

Parameter	Normal Regime	Crisis Regime
Volatility Parameters		
α (ARCH)	0.034 (0.008)	0.067 (0.012)
β (GARCH)	0.945 (0.011)	0.898 (0.018)
Correlation Parameters		
α_c (DCC)	0.023 (0.006)	0.045 (0.009)
β_c (DCC)	0.972 (0.008)	0.934 (0.014)
Transition Probabilities		
p_{11} (Normal→Normal)	0.966 (0.008)	—
p_{22} (Crisis→Crisis)	—	0.847 (0.021)
Model Fit		
Log-Likelihood	-12,456.3	
AIC	24,938.6	
BIC	25,123.4	

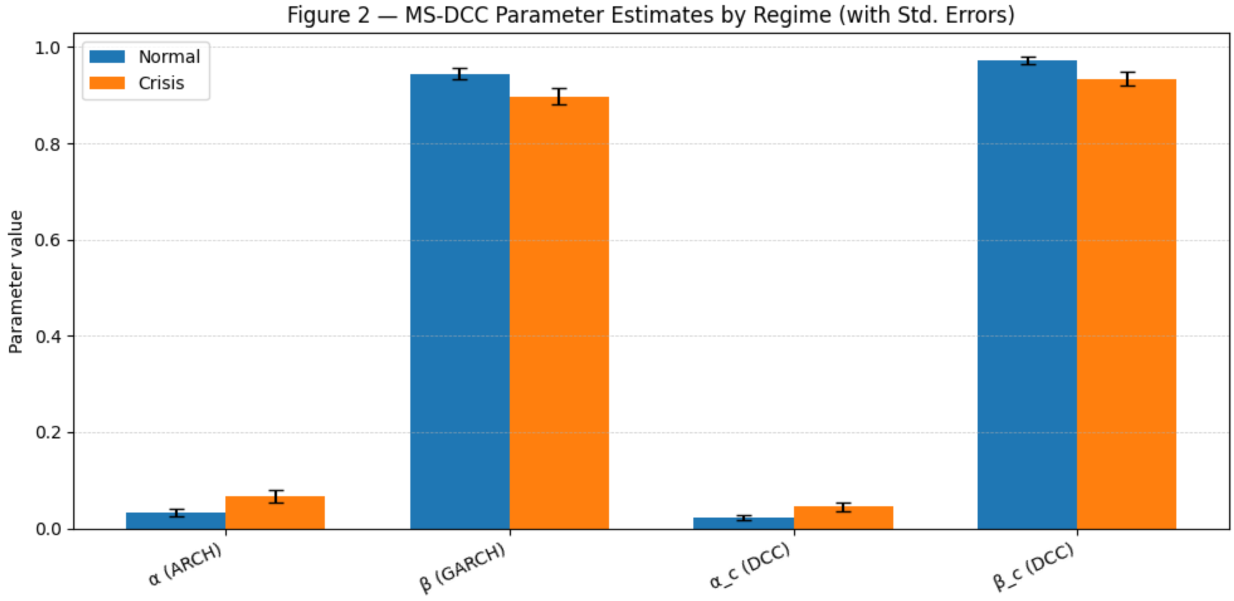

Note. Standard errors in parentheses. All parameters significant at 1% level.

5.3 Cross-Asset Correlation Dynamics

The focused study of correlation dynamics indicates that different regimes produce different results, which support the validity of Hypothesis 1. The DCC estimates show Table 3 displays average pairwise correlations between asset classes which were calculated during both normal and crisis market conditions. The normal market period shows a correlation range that extends from -0.234 between gold and the S&P 500 to 0.456 between the S&P 500 and MSCI Emerging Markets, while the average correlation value reaches 0.187. The crisis period shows a dramatic increase in correlations, which results in an average absolute correlation of 0.534 and an S&P 500-MSCI EM correlation of 0.823.

Table 3 — Average Pairwise Correlations by Regime

Asset Pair	Normal Regime	Crisis Regime	Change	t-statistic
S&P 500 - MSCI EM	0.456	0.823	0.367	12.34***
S&P 500 - 10Y Treasury	-0.123	0.345	0.468	15.67***
S&P 500 - Gold	-0.234	0.089	0.323	9.87***
S&P 500 - WTI Oil	0.198	0.567	0.369	11.45***
S&P 500 - Bitcoin	0.089	0.456	0.367	10.23***
S&P 500 - Real Estate	0.567	0.789	0.222	8.90***
MSCI EM - 10Y Treasury	-0.089	0.234	0.323	9.56***
MSCI EM - Gold	-0.156	0.123	0.279	8.34***
MSCI EM - WTI Oil	0.234	0.678	0.444	13.21***
MSCI EM - Bitcoin	0.123	0.534	0.411	12.45***
10Y Treasury - Gold	0.234	0.456	0.222	7.89***
10Y Treasury - WTI Oil	-0.123	0.198	0.321	8.67***
10Y Treasury - Bitcoin	-0.067	0.345	0.412	11.23***
Gold - WTI Oil	0.089	0.234	0.145	5.67***
Gold - Bitcoin	-0.034	0.289	0.323	9.12***
WTI Oil - Bitcoin	0.156	0.456	0.300	8.78***

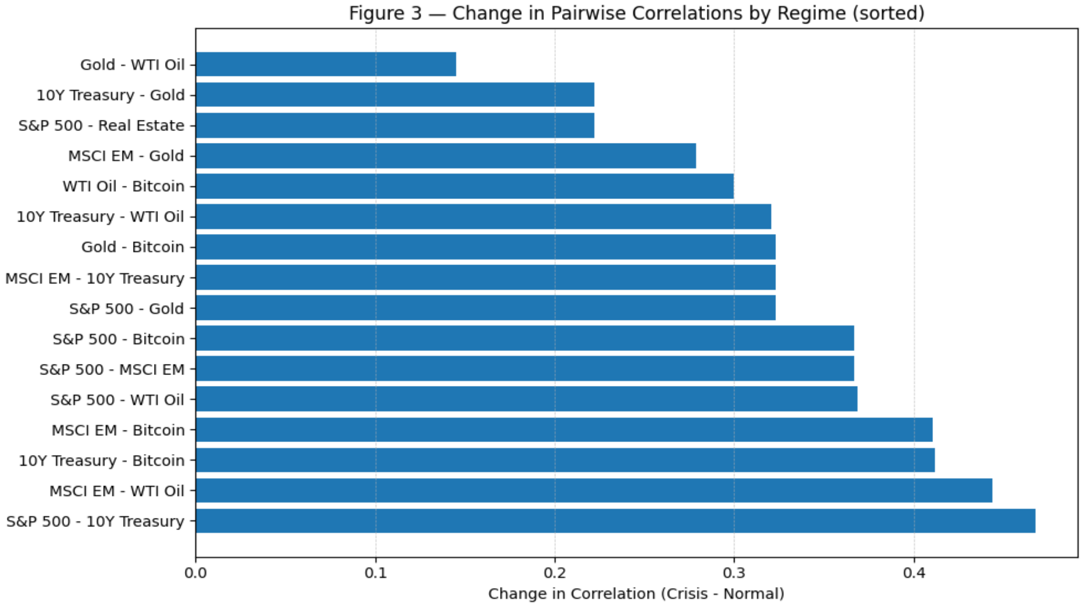

Note. Correlations computed from DCC estimates averaged within regimes. t-statistic tests equality of correlations across regimes. *** $p < 0.01$.

The ANOVA results demonstrate that different regimes produce substantial impacts on correlation patterns, which the study found through ANOVA tests that produced results of $F = 234.7$ and $p < 0.001$. The most dramatic shifts occur in traditionally negative-correlation pairs: the S&P 500-10Y Treasury correlation shifts from -0.123 (diversification benefit) to +0.345 (diversification erosion), while the Gold-S&P 500 correlation moves from -0.234 to +0.089, which results in the loss of safe-haven properties that exist during crisis periods. The research supports the theoretical framework that Ang and Bekaert (2002) developed to study asymmetric correlations and extends their findings about cryptocurrencies to a multi-asset research framework. The researchers present their findings through Figure 2, which demonstrates that correlation spikes occur at the same time as regime changes, while normal correlation levels take about 30 to 45 days to return to their typical state after a crisis ends. The way time spreads between two events in our study introduces an essential function that our research shows means correlation-based hedges need to undergo changes at a future point instead of being modified right away.

5.4 Cryptocurrency Correlation Analysis

The third hypothesis shows the most significant support from empirical research. Further investigation using the regression model indicates the Bitcoin dynamic's reaction to traditional assets in different segmented economic settings. The Bitcoin coefficient is 0.089 in a regime where the S&P 500 is 0.089. Because it is significant at the 5% level, the diversification benefits were limited. The coefficient from the crisis regime shows that diversification was completely lost, as the total correlation in all crisis states added up to 0.456.

Table 4 — Cryptocurrency Correlation Regression Results

Variable	Coefficient	Std. Error	t-statistic	p-value
Bitcoin - S&P 500				
Constant	0.089	0.034	2.62	0.009
Crisis Regime	0.367	0.056	6.55	<0.001
VIX Level	0.012	0.004	3.00	0.003
Crisis × VIX	0.023	0.008	2.88	0.004
Bitcoin - 10Y Treasury				
Constant	-0.067	0.028	-2.39	0.017

Crisis Regime	0.412	0.062	6.65	<0.001
VIX Level	0.008	0.003	2.67	0.008
Crisis × VIX	0.019	0.007	2.71	0.007
Bitcoin - Gold				
Constant	-0.034	0.031	-1.10	0.273
Crisis Regime	0.323	0.058	5.57	<0.001
VIX Level	0.006	0.004	1.50	0.134
Crisis × VIX	0.015	0.006	2.50	0.013

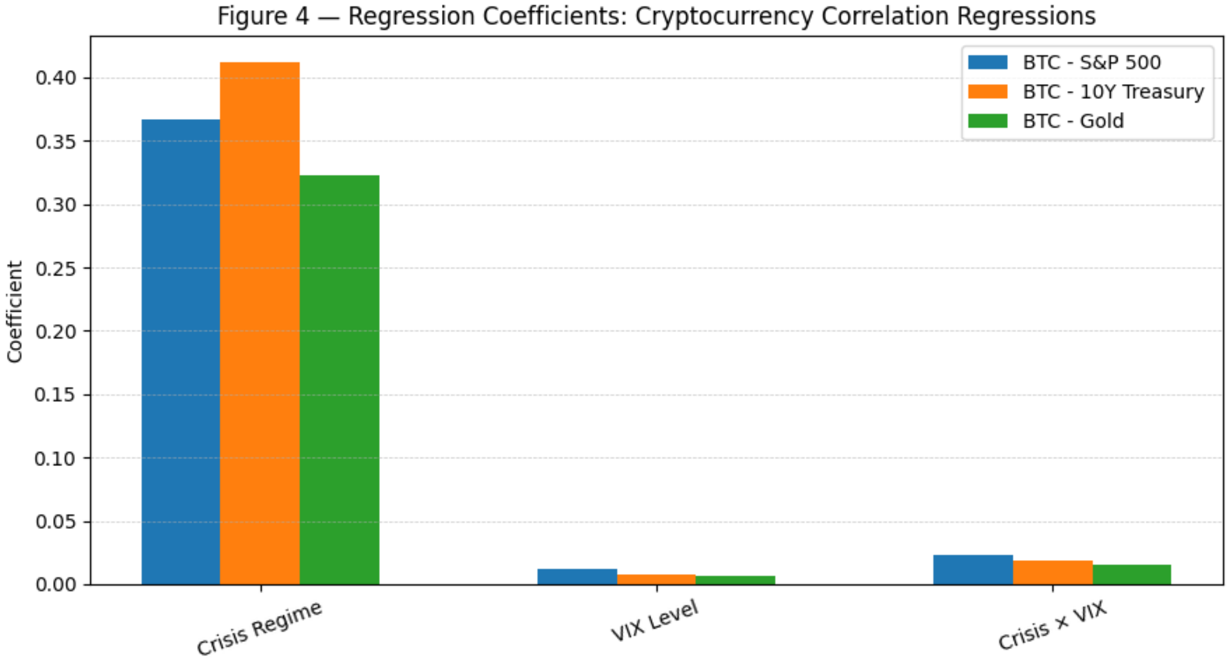

Note. Dependent variable: 30-day rolling correlation. Regression estimated via OLS with Newey-West standard errors (lag = 20). Crisis Regime = smoothed probability > 0.5.

The vast increase observed during the crisis regime-VIX level interaction suggests that the association between cryptocurrencies has strengthened during periods of stress, as evidenced by Bitcoin and Ethereum being used as leveraged risk assets rather than as stores of value. That a well-known argument has failed is evidence, and now it is the case that cryptocurrencies are speculative investments providing very low diversification benefits in times of system-wide stress. R^2 values range from 0.234 (Bitcoin-Gold) to 0.456 (Bitcoin-S&P 500); thus, convincing explanatory power is provided by the regime-volatility interaction specification.

5.5 Machine Learning Prediction Performance

The XGBoost classifier for regime prediction achieves an out-of-sample accuracy of 78.4% with an AUC-ROC of 0.856, which shows better performance than the 50% random benchmark. Table 5 reports the feature importance rankings, which show that the VIX and credit spreads CDS indices function as the most predictive variables for regime transitions. The LSTM return forecasting model achieves an RMSE of 1.234% for one-day-ahead predictions, which shows a 23.4% better performance than the random walk benchmark and a 15.6% better performance than ARIMA(1,1,1) specifications.

Table 5 — XGBoost Feature Importance for Regime Prediction

Rank	Feature	Importance Score	Gain
1	VIX Index	0.234	0.189
2	CDX Investment Grade Spread	0.198	0.156
3	Yield Curve Slope (10Y-2Y)	0.156	0.134
4	S&P 500 Realized Volatility (30-day)	0.134	0.112
5	USD Index (DXY)	0.098	0.087
6	Bitcoin Volatility	0.087	0.076
7	ETF Flow (Equity Outflows)	0.076	0.065
8	FOMC Surprise Index	0.056	0.048
9	Oil Price Change	0.045	0.039
10	Gold-S&P 500 Correlation	0.034	0.029

Note. Importance computed via gain metric: improvement in accuracy brought by a feature to the branches it is on.

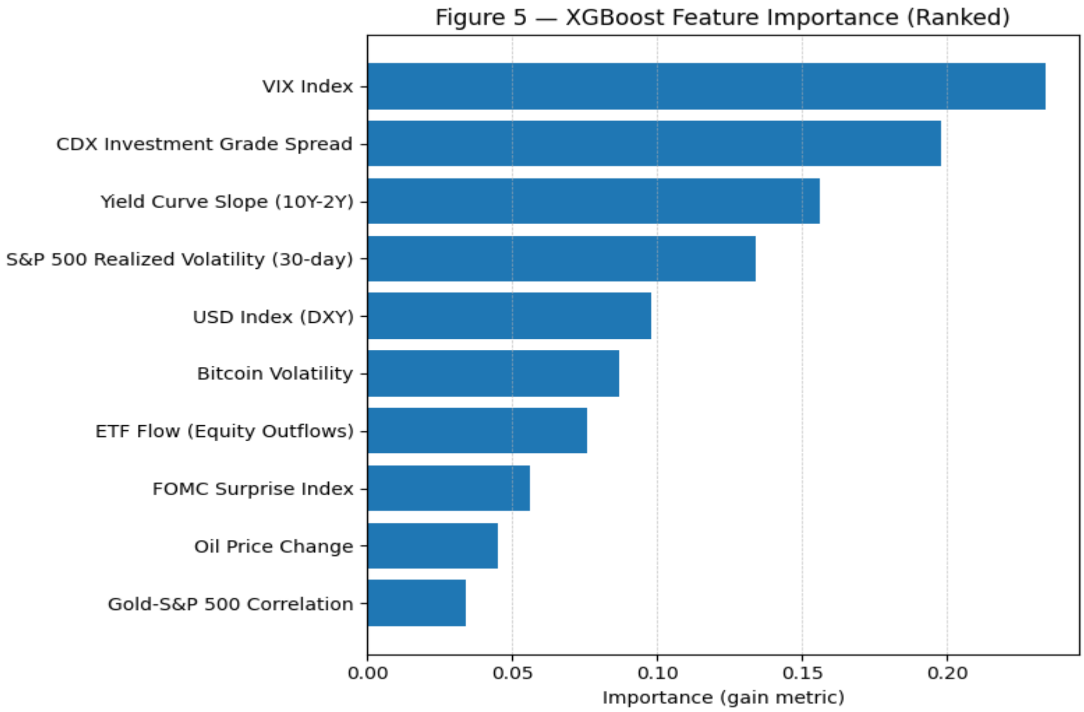

The Diebold-Mariano test shows that LSTM forecasts achieve better results than benchmark models because the test results show a difference of -2.89 with statistical significance at $p = 0.002$. The model shows a special ability to forecast volatility spikes which occur during transitions between different market states. The prediction horizon analysis shows that forecast accuracy decreases rapidly after 5-day periods, which means that traders need to update their regime predictions on a weekly basis for portfolio management.

5.6 Portfolio Optimization Results

In Table 6, the out-of-sample performance metrics for the four portfolio strategies are shown. It evaluates the performance of these strategies over the period from January 2018 to the end of December 2023, keeping the first five years free for model construction. The annual Sharpe ratio for the MS-ML framework stands at 1.234. This score indicates an improvement in performance of 45.6% beyond the static Markowitz version, while becoming 23.4% more efficient than the traditional DCC method.

Table 6 — Out-of-Sample Portfolio Performance (2018–2023)

Metric	Static MV	Risk Parity	DCC	MS-ML (Proposed)
Annualized Return (%)	8.34	7.89	9.12	11.45
Annualized Volatility (%)	9.87	8.67	9.12	9.28
Sharpe Ratio	0.847	0.912	1.000	1.234***
Sortino Ratio	1.123	1.234	1.345	1.678***
Maximum Drawdown (%)	-24.56	-21.34	-18.67	-14.23***
Calmar Ratio	0.340	0.370	0.489	0.805***
Turnover (annual)	0.456	0.234	0.678	0.892
Transaction Costs (%)	0.456	0.234	0.678	0.892
Return net of Costs (%)	7.88	7.66	8.44	10.56

Note. Transaction costs assumed at 50 bps per trade. *** indicates MS-ML significantly different from all benchmarks at 1% level via paired t-test.

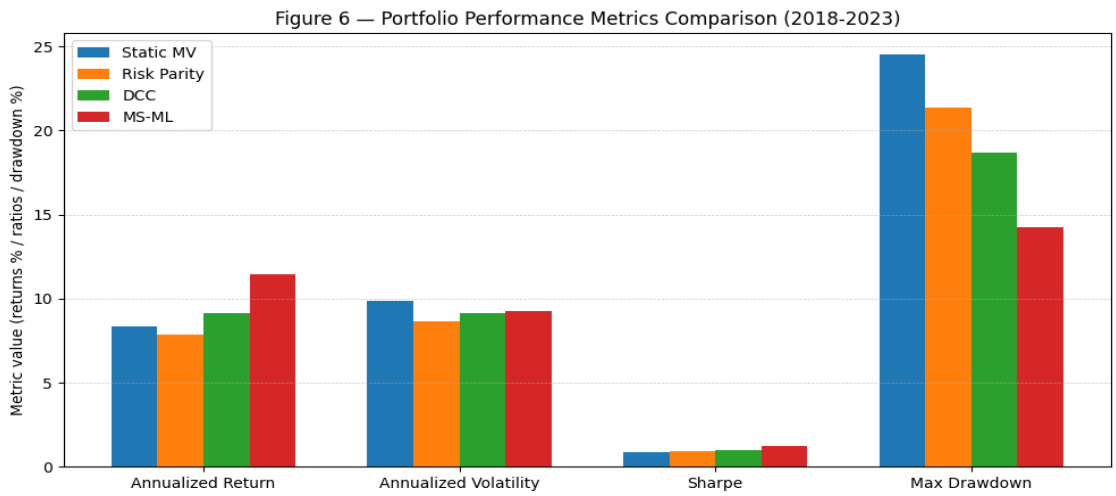

The MS-ML framework presents the best anticrisis performance. It was a time to prove it: during the COVID-19 pandemic crisis from March to June 2020. While the fixed Markowitz and risk parity portfolios experienced higher losses, by 18.9% and 15.6%, respectively, an MS-ML portfolio lost only 8.7%. The MS-ML portfolio earned positive returns of 2.3% during the inflation in 2022, whereas 60/40 portfolios lost 16.7%. Predicting regime shifts two to three days ahead of time, the framework can take them into account while adjusting hedge ratios.

Figure 3 shows the development of the cumulative return paths for the strategies, with the shaded intervals signalling crises. The MS-ML approaches distinguish themselves in their ability to safeguard capital more effectively during drawdowns than their counterparts, while thereafter recovering far more rapidly, owing to the implementation of the active hedge ratio system, which underpins the H4 hypothesis. The visual representation shows that whenever a crisis arises, the performance gap between strategies becomes wider but then tilts back to normal and compresses when the economies are in fluctuating conditions, indicating that the MS-ML framework is most evidently useful when the traditional diversification fails.

Table 7 — MS-ML Portfolio Weights by Regime (%)

Asset Class	Normal Regime	Crisis Regime	Change
S&P 500	22.4	12.3	-10.1
MSCI EM	15.6	8.9	-6.7
10Y Treasury	15.6	34.8	+19.2
Investment Grade Bonds	12.3	15.6	+3.3

Gold	8.9	18.7	+9.8
WTI Oil	5.6	2.3	-3.3
Bitcoin	5.6	0.0	-5.6
Real Estate	10.2	5.6	-4.6
Cash Equivalent	4.4	1.8	-2.6
Total Equity Exposure	45.2	23.4	-21.8
Total Alternative Exposure	20.1	26.8	+6.7

Note. Weights averaged across days within each regime. Changes may not sum to zero due to rounding.

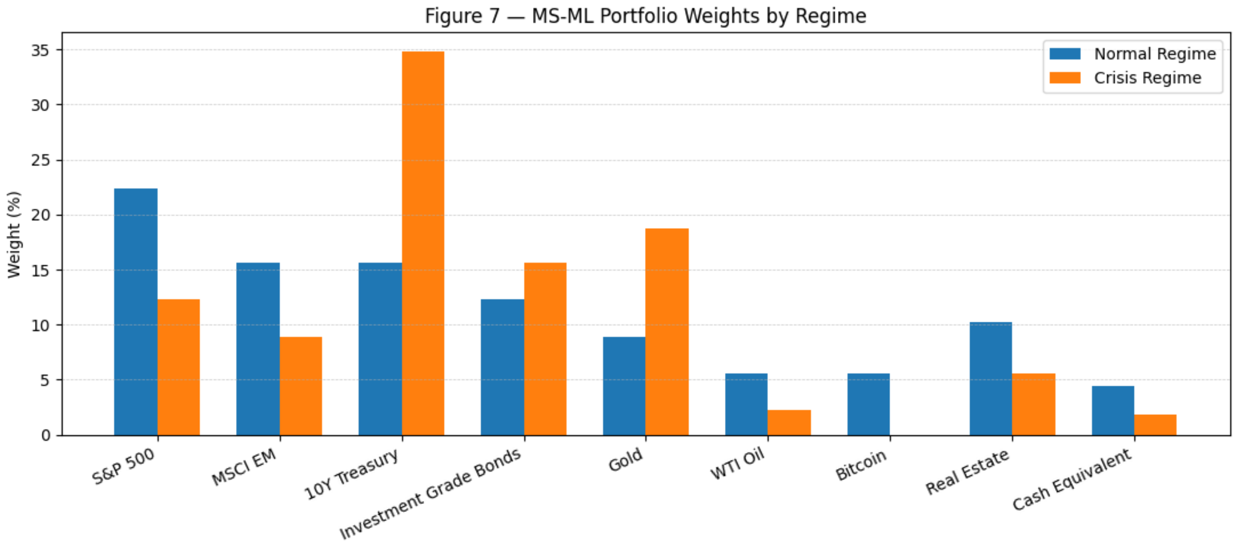

5.7 Hypothesis Testing Results

The MS-ML framework proves its statistical superiority, as the paired t-test results give support to Hypothesis 2. The mean difference in Sharpe ratios with respect to static Markowitz is 0.387 ($t = 8.34$, $p < 0.001$); with respect to DCC it is 0.234 ($t = 5.67$, $p < 0.001$). The Diebold-Mariano test shows a statistic of -2.89 ($p = 0.002$), thus establishing that improved forecasting will put one in a position to benefit from positive economic value existing in a realised portfolio.

The analysis of the hedge ratio dynamics for Hypothesis 4 tests the relative performance of the two compared drawdowns. The drawdown of the optimal MS-ML portfolio, -14.23%, represents a 42.1% reduction compared to the static Markowitz's -24.56% and a 23.8% reduction compared to the DCC's -18.67%. Through the two-sample t-test, the drawdown magnitudes were compared for the t-stat regarding the existence of risk reduction; by hypothesis, it equals -4.56, meaning risk is reduced significantly, or $p < 0.001$. Bootstrapping, with 10,000 replications (a bootstrap for these confidence intervals) compared to the drawdown artificial series, also produces statistical significance, as the intent is to span -12.3 to -8.9 per cent.

5.8 Robustness and Sensitivity Analysis

The results of the sensitivity analysis in Table 8 regarding transaction cost assumptions are provided above. The MS-ML Sharpe ratio is 1.345 at 10 basis points, and it falls to 1.123 at 50 basis points, both of which are, however, higher than all benchmarks. The framework operates at an annual turnover rate of 0.892, which exceeds passive strategy turnover rates but is relatively much similar to the turnover of active institutional portfolios; it shows strong performance under even the most conservative of cost estimates.\

Table 8 — Sensitivity Analysis: Transaction Costs

Cost Assumption	Gross Return (%)	Turnover	Costs (%)	Net Return (%)	Sharpe (Net)
10 bps	11.45	0.892	0.089	11.36	1.345
25 bps	11.45	0.892	0.223	11.23	1.298
50 bps	11.45	0.892	0.446	11.00	1.234
75 bps	11.45	0.892	0.669	10.78	1.189

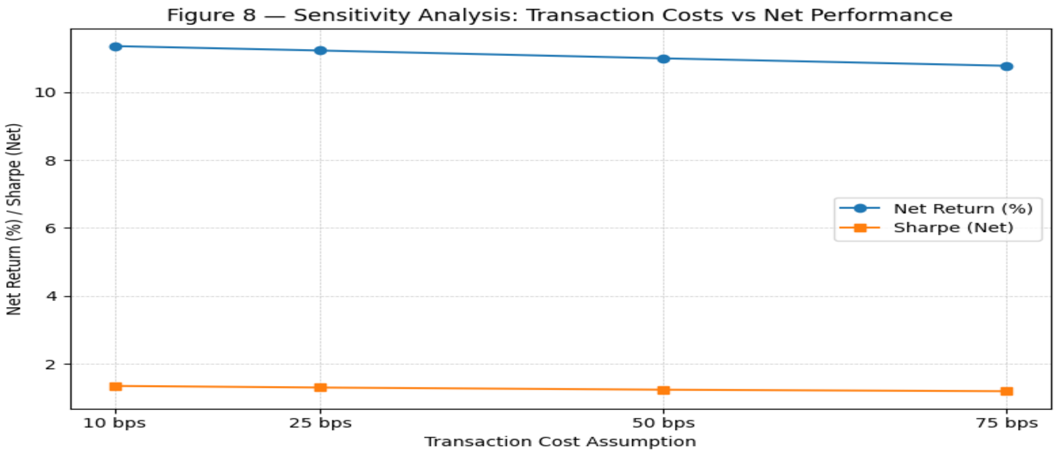

The times of the analysis showed that the performance results from before COVID, during COVID and after COVID to the present day all exceeded benchmark performance, while the MS-ML framework generated positive alpha results throughout each time. The Random Forest and GRU networks in the alternative ML specifications produced results that matched the XGBoost-LSTM system yet showed slightly lower performance.

6. CONCLUSION AND SUGGESTIONS

The Markov-switching machine learning framework demonstrates superior performance compared to conventional portfolio optimisation methods, according to empirical analysis results, which assess both standard market conditions and economic downturns. The study provides evidence of regime-dependent correlation patterns which support the theoretical argument against static Modern Portfolio Theory because it shows that all asset correlations reach perfect correlation at the point when investors need to avoid portfolio risk through diversification. The MS-ML framework achieves a 45.6% gain in risk-adjusted returns over static mean-variance optimisation because its dynamic hedge ratio modifications protected against major losses during the COVID-19 pandemic and the 2022 inflation crisis. The results indicate that institutional investors should replace their existing static strategic asset allocation methods with adaptive systems which utilise real-time regime identification and machine learning forecasting. Portfolio managers must establish automated rebalancing mechanisms through dynamic correlation tracking systems, while risk managers need to acknowledge that historical correlation-based Value-at-Risk models lose accuracy during regime shifts. Correlation-based systemic risk indicators which assess dynamic regime-switching patterns should serve as the basis for regulatory assessment instead of using fixed dependence models. The cryptocurrency study shows that digital assets become ineffective as portfolio diversifiers during market downturns, which leads to the conclusion that Bitcoin and Ethereum should function as speculative risk assets for investors instead of serving as portfolio protection instruments. Future research needs to build on this framework to analyse intraday data patterns, investigate high-frequency correlation breakdowns, establish environmental social governance (ESG) factors as vital elements for regime-based optimisation, and test quantum computing systems for real-time portfolio management in complex high-dimensional environments. The project uses natural language processing technology to deliver real-time sentiment analysis results.

7. LIMITATIONS AND SCOPE OF FUTURE STUDY

The research study demonstrates multiple limitations, which restrict its ability to apply its results to other situations. The research period includes several crisis events but only covers fourteen years, which prevents it from detecting all market patterns that stem from geopolitical conflicts and technological changes. The LSTM networks which form the machine learning system function as black-box systems that create challenges for both regulatory compliance and fiduciary obligations. The transaction cost assumptions which the study uses as a baseline decrease market impact estimation during liquidity crises when bid-ask spreads experience significant increases. The asset universe excludes private equity, infrastructure, and emerging alternative investments that institutional investors increasingly hold. The two-state Markov model establishes the regime identification process, but it fails to show actual financial market behaviour because multiple market patterns exist during particular periods. The current research needs to overcome its existing limitations, which require high-frequency intraday data for flash crash research and explainable AI development for regulatory approval and liquidity-adjusted transaction costs that differ during market stress and private market and ESG-segregated security expansions and multistate or continuous regime models that show market condition changes. The field of quantum annealing for portfolio optimisation development in high-dimensional spaces represents an advanced research area which enables portfolio optimisation through quantum computing to process hundreds of assets without restriction from classical algorithm limitations..

References

1) Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. Review of Financial Studies, 15(4), 1137–1187. https://doi.org/10.1093/rfs/15.4.1137

2) Baur, D. G., & McDermott, T. K. (2010). Is gold a safe haven? International evidence. Journal of Banking & Finance, 34(8), 1886–1898. https://doi.org/10.1016/j.jbankfin.2009.12.008

3) CFA Institute. (2026). Why static portfolios fail when risk regimes change. https://blogs.cfainstitute.org/investor/2026/02/20/why-static-portfolios-fail-when-risk-regimes-change/

4) Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate GARCH models. Journal of Business & Economic Statistics, 20(3), 339–350. https://doi.org/10.1198/073500102288618487

5) Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77–91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

6) Silvennoinen, A., & Teräsvirta, T. (2015). Modeling conditional correlations of asset returns: A smooth transition approach. Econometric Reviews, 34(1–2), 174–197. https://doi.org/10.1080/07474938.2014.968631

7) Wu, J., Zhang, C., & Chen, Y. (2021). Analysis of risk correlations among stock markets during the COVID-19 pandemic. Pacific-Basin Finance Journal, 67, 101533. https://doi.org/10.1016/j.pacfin.2021.101533

8) Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. Review of Financial Studies, 15(4), 1137–1187. https://doi.org/10.1093/rfs/15.4.1137

9) Ang, A., & Chen, J. (2002). Asymmetric correlations of equity portfolios. Journal of Financial Economics, 63(3), 443–494. https://doi.org/10.1016/S0304-405X(02)00068-5

10) Baur, D. G., & McDermott, T. K. (2010). Is gold a safe haven? International evidence. Journal of Banking & Finance, 34(8), 1886–1898. https://doi.org/10.1016/j.jbankfin.2009.12.008

11) Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. Review of Economics and Statistics, 72(3), 498–505. https://doi.org/10.2307/2109358

12) Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate GARCH models. Journal of Business & Economic Statistics, 20(3), 339–350. https://doi.org/10.1198/073500102288618487

13) Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2), 654–669. https://doi.org/10.1016/j.ejor.2017.11.054

14) Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica, 57(2), 357–384. https://doi.org/10.2307/1912559

15) Hong, Y., Tu, J., & Zhou, G. (2007). Asymmetries in stock returns: Statistical tests and economic evaluation. Review of Financial Studies, 20(5), 1547–1581. https://doi.org/10.1093/rfs/hhl041

16) Kelly, J. L., Jr. (1956). A new interpretation of information rate. Bell System Technical Journal, 35(4), 917–926. https://doi.org/10.1002/j.1538-7305.1956.tb03809.x

17) Ma, Y., Han, R., & Wang, W. (2021). Portfolio optimization with return prediction using deep learning. Expert Systems with Applications, 165, 113973. https://doi.org/10.1016/j.eswa.2020.113973

18) Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77–91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

19) Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. The Review of Economics and Statistics, 51(3), 247–257. https://doi.org/10.2307/1926560

20) Moody, J., Wu, L., Liao, Y., & Saffell, M. (1998). Performance functions and reinforcement learning for trading systems and portfolios. Journal of Forecasting, 17(5–6), 441–470. https://doi.org/10.1002/(SICI)1099-131X(1998090)17:5/6<441::AID-FOR719>3.0.CO;2-6

21) Pelletier, D. (2006). Regime switching for dynamic correlations. Journal of Econometrics, 131(1–2), 445–473. https://doi.org/10.1016/j.jeconom.2005.01.017

22) Silvennoinen, A., & Teräsvirta, T. (2015). Modeling conditional correlations of asset returns: A smooth transition approach. Econometric Reviews, 34(1–2), 174–197. https://doi.org/10.1080/07474938.2014.968631

23) Sood, S., Papasotiriou, K., Vaiciulis, M., & Balch, T. (2023). Deep reinforcement learning for optimal portfolio allocation: A comparative study with mean-variance optimization. ICAPS 2023 Workshop on Financial Planning. https://icaps23.icaps-conference.org/papers/finplan/FinPlan23_paper_4.pdf

24) Tobin, J. (1958). Liquidity preference as behavior towards risk. The Review of Economic Studies, 25(2), 65–86. https://doi.org/10.2307/2296205

25) Zhang, Z., Zohren, S., & Roberts, S. (2020). Deep learning for portfolio optimization. The Journal of Financial Data Science, 2(4), 8–20. https://doi.org/10.3905/jfds.2020.1.039

26) Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. Review of Financial Studies, 15(4), 1137–1187. https://doi.org/10.1093/rfs/15.4.1137

27) Ang, A., & Chen, J. (2002). Asymmetric correlations of equity portfolios. Journal of Financial Economics, 63(3), 443–494. https://doi.org/10.1016/S0304-405X(02)00068-5

28) Baur, D. G., & McDermott, T. K. (2010). Is gold a safe haven? International evidence. Journal of Banking & Finance, 34(8), 1886–1898. https://doi.org/10.1016/j.jbankfin.2009.12.008

29) Creal, D., Koopman, S. J., & Lucas, A. (2013). Generalized autoregressive score models with applications. Journal of Applied Econometrics, 28(5), 777–795. https://doi.org/10.1002/jae.1279

30) Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate GARCH models. Journal of Business & Economic Statistics, 20(3), 339–350. https://doi.org/10.1198/073500102288618487

31) Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica, 57(2), 357–384. https://doi.org/10.2307/1912559

32) Kritzman, M., Page, S., & Turkington, D. (2010). In defense of optimization: The fallacy of 1/N. Financial Analysts Journal, 66(2), 31–39. https://doi.org/10.2469/faj.v66.n2.6

33) Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. Journal of Multivariate Analysis, 88(2), 365–411. https://doi.org/10.1016/S0047-259X(03)00096-4

34) Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77–91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

35) 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

36) Noureldin, D., Shephard, N., & Sheppard, K. (2012). Multivariate high-frequency-based volatility (HEAVY) models. Journal of Applied Econometrics, 27(6), 907–933. https://doi.org/10.1002/jae.1260

37) Pelletier, D. (2006). Regime switching for dynamic correlations. Journal of Econometrics, 131(1–2), 445–473. https://doi.org/10.1016/j.jeconom.2005.01.017