University of Southampton

FACILITY OF SOCIAL SCIENCES
SOUTHAMPTON BUSINESS SCHOOL

Essays on the Role of Memory in Financial Markets.

by

Dmitri Mustanen

ORCiD: 0009-0008-7918-7790

A thesis for the degree of
Doctor of Philosophy in Quantitative Finance

January 2026

University of Southampton

Abstract

FACILITY OF SOCIAL SCIENCES
SOUTHAMPTON BUSINESS SCHOOL

Doctor of Philosophy

Essays on the Role of Memory in Financial Markets.

by Dmitri Mustanen

This thesis investigates the phenomenon of persistence as a key and economically meaningful characteristic of financial markets, one that critically influences forecasting performance, price discovery, and long-term investment outcomes. Despite extensive research on long memory in asset prices, its economic significance and practical importance remain insufficiently understood. This study helps close this gap by demonstrating that persistence captures how information propagates through interconnected financial systems and can be systematically measured and exploited.

To operationalise these insights, the thesis develops an integrated framework linking system identification, memory dynamics, and economic application. A spillover-based variable selection approach is introduced to define financial systems through directional shock transmission, ensuring that persistence is estimated within coherent and internally connected market structures. Applied to energy markets, this framework improves forecast accuracy and yields more stable and interpretable memory estimates.

Building on this foundation, the thesis shows that persistence shapes price formation through the interaction of historical dependence and forward-looking anticipation. Using fractional cointegration and bidirectional memory measures, it demonstrates that fundamentally driven markets are dominated by backward-looking persistence, while speculative markets exhibit stronger anticipatory dynamics. These distinct memory profiles explain systematic variation in predictability and establish persistence as an information-rich determinant of asset price dynamics.

Finally, persistence is embedded into a multi-period asset-pricing framework, transforming memory from a descriptive time series feature into a structural intertemporal risk characteristic. By incorporating fractional memory into closed-form expressions for effective variance and skewness, the thesis develops a memory-augmented Capital Asset Pricing Model that links serial dependence to long-horizon wealth dynamics. Empirical evidence from US and Japanese equity markets shows that memory-aware strategies enhance risk-adjusted portfolio performance and long-horizon efficiency.

Contents

List of Figures **ix**

List of Tables **xi**

Declaration of Authorship **xiii**

Acknowledgements **xv**

1 Introduction **1**
1.1 Research Context . . . 2
1.2 Research Aims . . . 3
1.3 Research Objectives . . . 3
1.4 Original Contributions of the Thesis . . . 4
1.5 Thesis Structure . . . 5

2 Improving System Generalisation & Forecastability via Spillover-Based Variable Selection Approach. **7**
2.1 Introduction . . . 8
2.2 Literature Review . . . 10
2.3 Data and Model Specification . . . 12
2.3.1 Model Specification . . . 13
2.4 Empirical Methodology . . . 14
2.4.1 Preliminary Tests and VECM Specification . . . 14
2.4.2 Spillover Analysis . . . 16
Frequency Domain Spillover Analysis. . . . 17
Dynamic (Rolling Window) Spillover Analysis. . . . 18
2.4.3 Variable Forecastability Validation . . . 18
FCVAR Approach. . . . 18
Robustness Check - LSTM Benchmark. . . . 19
Rolling Window Forecasting and Evaluation. . . . 19
2.5 Empirical Results . . . 20
2.5.1 Spillover Analysis . . . 20
2.5.2 Network Analysis . . . 21
2.5.3 Dynamic Spillover Analysis . . . 24
2.6 Concept Validation . . . 25
2.6.1 Out-of-Sample FCVAR Forecasting Performance . . . 26
2.6.2 Forecast Comparison Tests and Model Dominance . . . 27
Conditional Predictive Ability (CPA). . . . 29

Model Confidence Set (MCS). 29
Forecast Encompassing. 29
2.6.3 Robustness Check - Evaluation of Model 6 Forecast Performance Against Analogous Research 30
2.6.4 Synthesis with Prior Research 31
2.7 Policy and Practical Implications 33
2.8 Conclusion 33

3 Forward and Backward Memory in Fractionally Cointegrated Systems. 35
3.1 Introduction 36
3.2 Literature Review 37
3.3 Methodology 40
3.3.1 FCVAR Model and Bidirectional Memory Estimation 40
3.3.2 Empirical Strategy 40
3.4 Data and Preprocessing checks 41
3.4.1 Pre-Estimation Memory Validation 43
3.5 Empirical Results 45
3.5.1 Relation of Memory Dynamics to Prior Research 48
3.6 Implications of Bidirectional Memory Dynamics 49
3.6.1 Forecast Performance: Brent Oil vs. Bitcoin 49
3.6.2 Interpretation and Broader Implications 50
3.6.3 VAR Granger Causality Tests 52
3.6.3.1 Brent Market Analysis 53
3.6.3.2 Bitcoin Market 54
3.7 Conclusion 55

4 A Memory-driven Multi-period Capital Asset Pricing Model 57
4.1 Introduction 58
4.2 Literature 60
4.3 Theoretical Framework - Compounding-Adjusted Risk 62
4.3.1 Log-Utility, Jensen's Inequality and Volatility Drag 62
4.3.2 Taylor Expansion and Single–Period Approximation 64
4.3.3 Multi-Period Drag and Effective Variance 64
4.3.4 Fractional Memory and the $ARFIMA(0,\ d,\ 0)$Model 66
4.3.5 Variance Decomposition in Presence of Autocorrelation in Returns 66
4.3.6 Portfolio-Level Variance and Drag 67
4.3.7 Memory-Augmented Mean-Variance Optimisation 69
4.4 Empirical Methodology 70
4.4.1 Economic Significance of Memory 70
Memory estimation. 70
Interpretation of the memory parameter. 70
Memory and mean–reversion link examination. 71
CAPM extension. 71
Out-of-sample return forecasts. 72
4.4.2 Theoretical Framework Validation 72
Portfolio Performance Analysis 72
Efficient–Frontier Diagnostics 72

4.5 Data . . . 74
4.6 Empirical Analysis . . . 74
4.6.1 Portfolio Level Statistics Overview . . . 77
4.6.2 Mean-Reversion Analysis . . . 78
4.6.3 Memory-Factor CAPM and GRS Test Results . . . 79
4.6.4 CAPM Forecasting Performance. . . . 80
4.6.5 Portfolio Performance Analysis . . . 81
4.6.5.1 Intermediate-Memory Portfolios Dominance . . . 83
4.6.6 Memory Portfolios vs Market Relative Performance. . . . 85
4.6.7 Efficient–Frontier Diagnostics . . . 89
4.7 Robustness Analysis: Evidence from the Japanese Market . . . 90
4.7.1 Link with Existing Studies . . . 93
4.8 Limitations . . . 93
4.9 Conclusion . . . 94

5 Conclusions **97**

Appendix A Supplement to Chapter 2 **101**
Appendix A.1 Static (Time-Domain) Spillover Analysis . . . 101
Appendix A.2 Fractional Cointegration Vector Autoregressive Model . . . 102
Appendix A.2.1 Model Derivation . . . 103
Appendix A.2.2 Forecasting Using FCVAR . . . 104
Appendix A.3 Six Model Spillover Analysis . . . 105
Appendix A.3.1 Statics Analysis of Intra-Market Dynamism - Isolated Systems. . . . 105
Appendix A.3.1.1 *Model 1 - BSP, BE, BL, BR.* . . . 105
Appendix A.3.1.2 *Model 2 - WSP, WE, WL, WR.* . . . 106
Appendix A.3.1.3 *Model 3 - BSP, BE, BL, BR, WSP, WE, WL, WR.* . . . 108
Appendix A.3.2 Statics Analysis of Intra-Market Dynamism - Global Systems. . . . 110
Appendix A.3.2.1 *Model 4 - BSP, E, L, R.* . . . 110
Appendix A.3.2.2 *Model 5 - WSP, E, L, R.* . . . 112
Appendix A.3.2.3 *Model 6 - BSP, WSP, E, L, R.* . . . 114
Appendix A.4 Spillover Analysis for Model 6 vs Guo et al. dataset. . . . 115

Appendix B Supplement to Chapter 3 **117**
Appendix B.1 Fractional Cointegrated Vector Auto Regression Model . . . 117
Appendix B.1.1 Forecasting Using FCVAR . . . 119
Appendix B.2 Rolling Window Size Estimation . . . 120
Appendix B.3 Variable Interdependence Evaluation Methods . . . 121
Appendix B.3.1 VAR Granger Causality & Wald Block Exogeneity Tests 121
Appendix B.3.2 Wald Block Exogeneity Tests . . . 122
Appendix B.4 Loss Functions . . . 123

Appendix C Supplement to Chapter 4 **125**
Appendix C.1 Japan's market analysis . . . 125

References **133**

List of Figures

2.1 : Network Analysis Summary of Intra-Market Dynamics in the Brent and WTI Oil Markets. 23
2.2 : Total Spillover Index vs Global Economic Policy Uncertainty Index. . . 25
2.3 : Brent and WTI Spot Out-of-Sample Forecasts - FCVAR Model 1, 2 & 6. 27

3.1 : Forward and backward system long-memory profiles of Brent and Bitcoin markets - obtained via FCVAR framework using expanding and rolling window estimation. 46
3.2 : Error profiles of FCVAR out-of-sample of Brent and Bitcoin spot price forecasts in forward and backward direction. Expanding vs rolling window analysis. 50

4.1 : Illustration of The Effect of Volatility Drag on Investor's Expected Log-Returns. 63
4.2 : Cumulative Returns of Frictionless, Equally Weighted Portfolios under Daily Rebalancing. 83
4.3 : Cumulative Returns of Equally Weighted Portfolios with Transaction Costs 30 Day Rebalancing Frequency. 85
4.4 : Average Estimated Memory vs. Average Daily Returns of US Market. . 86
4.5 : Illustration of The Effect of Volatility Drag of Investor's Expected Log-Returns in Presence of Positive Skewness. 87
4.6 : Efficient Frontier Graphs of Different Covariance Models. 91

Appendix A.1 : Model 1 - Network analysis of intra-market dynamics in Brent oil market. 106
Appendix A.2 : Model 2 - Network analysis of intra-market dynamics in WTI oil market. 107
Appendix A.3 : Model 3 - Network analysis of inter-market dynamics between Brent & WTI markets. 109
Appendix A.4 : Model 4 - Network analysis of dynamics between Brent spot and global oil market variables. 111
Appendix A.5 : Model 5 - Network analysis of dynamics between WTI spot and global oil market variables. 113
Appendix A.6 : Model 6 - Network analysis of dynamics between Brent and WTI spot and global oil market variables. 114

Appendix C.1 : Cumulative Returns of Frictionless, Equally Weighted Portfolios under Daily Rebalancing (Japan Market). 130
Appendix C.2 : Cumulative Returns of Equally Weighted Portfolios with Transaction Costs and 30-Day Rebalancing Frequency (Japan Market). . 131

Appendix C.3 : Cumulative Returns of Equally Weighted Portfolios with Transaction Costs and 30-Day Rebalancing Frequency (Japan Market) - Excluding 7012.T. 132

List of Tables

2.1 Bloomberg tickers and construction of variables used in the spillover analysis. . . . 12
2.2 : Summary statistics. . . . 13
2.3 : Model Selection Summary. . . . 14
2.4 : Unit Root Test Results: DF–GLS vs KPSS. . . . 15
2.5 : VECM Model Parameters and Complexity Summary. . . . 16
2.6 : Guo et al. (2023b) variables. . . . 19
2.7 : LSTM network architecture and hyperparameter settings across forecast horizons. . . . 20
2.8 : Spillover TSI Summary for Models M1 vs M4, M2 vs M5, and M3 vs M6. 22
2.9 : TSI Summary statistics. Dynamic analysis. . . . 26
2.10 : FCVAR Out-of-Sample Forecast Evaluation. . . . 28
2.11 : Unified Forecast Evaluation Summary (BSP and WTI Spot Models) . . 28
2.12 : LSTM Out-of-Sample Forecast Evaluation. . . . 30

3.1 : Summary of Analysis Variables. . . . 42
3.2 : Summary Statistic of Analysis Variables 43
3.3 : Summary Statistic of Analysis Variables 44
3.4 : DF–GLS and KPSS Unit Root Tests. . . . 45
3.5 : Summary Statistic of Memory Profiles 47
3.6 : Out-of-sample forecast error profiles of Brent and Bitcoin Sport price forecasts for 100 observations in forward and backward directions. 51
3.7 : VAR Granger Causality Tests analysis of causality structure within Brent & Bitcoin markets. . . . 53

4.1 : Asset Universe Returns - Descriptive Statistics. . . . 75
4.2 : Asset Universe Memory - Descriptive Statistics. . . . 76
4.3 : Portfolio Returns - Descriptive Statistics. . . . 77
4.4 : Portfolio Memory - Descriptive Statistics. . . . 77
4.5 : Mean Reversion Analysis Results. . . . 79
4.6 : HML CAPM Extension - GRS Test Results. . . . 80
4.7 : Conventional vs Extended CAPM Out-Of-Sample Forecast Performance. 81
4.8 : Portfolio Performance Metrics - Daily Rebalanced Frictionless Equally Weighted Portfolios. . . . 82
4.9 : Portfolio Performance Metrics - Equally Weighted Portfolio with Transaction Cost & 30 Day Rebalancing Frequency. . . . 84
4.10 : Tangency Portfolio Analysis of Different Covariance Models. . . . 89

Appendix A.1 : Model 1 - Static Directional Connectedness in Brent Market. 105

Appendix A.2 : Model 1 - Frequency Decomposition of Static Directional Connectedness in Brent Market. 106
Appendix A.3 : Model 2 - Static Directional Connectedness in WTI market. . 107
Appendix A.4 : Model 1 - Frequency Decomposition of Static Directional Connectedness in WTI Market. 108
Appendix A.5 : Model 2 - Static Directional Connectedness between Brent & WTI markets. 108
Appendix A.6 : Model 3 - Frequency Decomposition of Static Directional Connectedness between Brent & WTI Markets. 110
Appendix A.7 : Model 4 - Static Directional Connectedness in Brent Market represented through global oil market variables. 110
Appendix A.8 : Model 4 - Frequency Decomposition of Static Directional Connectedness Between Brent and Global Oil Market. 112
Appendix A.9 : Model 5 - Static Directional Connectedness in WTI oil market represented through global oil market variables. 112
Appendix A.10 : Model 5 - Frequency Decomposition of Static Directional Connectedness Between WTI and Global Oil Market. 113
Appendix A.11 : Model 6 - Static Directional Connectedness in Brent and WTI oil markets represented through global oil market variables. 114
Appendix A.12 : Model 6 - Frequency Decomposition of Static Directional Connectedness Between Brent, WTI and Global Oil Market. 115
Appendix A.13 : Static Spillover Analysis - Model 6 Variables. 116
Appendix A.14 : Static Spillover Analysis - Guo et al. (2023b) Variables. 116

Appendix C.1 : Asset Universe Returns – Descriptive Statistics (Japan Market). 126
Appendix C.2 : Asset Universe Memory - Descriptive Statistics (Japan Market). 127
Appendix C.3 : Quintile and Market Portfolio Returns – Descriptive Statistics (Japan Market). 128
Appendix C.4 : Quintile and Market Portfolio Memory – Descriptive Statistics (Japan Market). 128
Appendix C.5 : HML CAPM Extension – GRS Test Results (Japan Market). 129
Appendix C.6 : Conventional vs Extended CAPM Out-Of-Sample Forecast Performance (Japan Market). 129
Appendix C.7 : Portfolio Performance Metrics – Daily Rebalanced Frictionless Equally Weighted Portfolios (Japan Market). 131
Appendix C.8 : Portfolio Performance Metrics – Equally Weighted Portfolios with 30-Day Rebalancing Transaction Costs (Japan Market). 132

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. None of this work has been published before submission

Signed:..... Date:.....

Acknowledgements

I would like to take this opportunity to express my heartfelt gratitude to all those who have supported and guided me through my PhD journey. First and foremost, I extend my sincere appreciation to the Southampton Business School for generously providing the scholarship, and to Professor Tapas Mishra for nominating me for this prestigious opportunity. Without this nomination and the School's support, achieving this lifelong ambition would have remained merely a dream.

I would like to express my profound gratitude to my main supervisor, Professor Taufiq Choudhry, whose unwavering support, trust, and belief in my abilities have been fundamental to the successful completion of this thesis. The foundations laid by his teaching during my MSc studies have significantly shaped my academic growth and PhD research interests. My special thanks and heartfelt appreciation also extend to my second supervisor and dear friend, Dr Ahmad Maaitah, whose insightful advice, intellectual and emotional generosity greatly enriched and helped this journey.

Above all, I owe an immense debt of gratitude to my wife and my children for their unconditional love, endless patience, and steadfast support throughout these years. The debt of family moments lost can never fully be repaid, but my love and gratitude for their care, strength, and understanding will endure long beyond this achievement.

Thank you all for standing by my side through every challenge and milestone along this journey.

To my children ...

Personal growth is both a journey and a choice, one that rewards perseverance as much as it demands curiosity.

Chapter 1

Introduction

1.1 Research Context

Persistence, often characterised econometrically as long memory, is a well-documented feature of financial time series. Despite extensive statistical evidence, its deeper economic meaning, behavioural origins, and practical relevance for forecasting and investment decisions remain insufficiently articulated. Much of the existing literature treats persistence primarily as a statistical regularity, with limited attention to how it arises from investor behaviour, information frictions, and systemic interdependence. As a result, persistence is frequently modelled without a clear link to the structure of the financial system in which it is observed.

Throughout this thesis, a distinction is drawn between the concepts of *memory* and *persistence*. Memory refers to the economic and behavioural mechanism through which past information, beliefs, and expectations are retained and propagated within interconnected financial systems. Persistence denotes the observable statistical manifestation of this mechanism, typically characterised by slow decay of dependence and measured using fractional integration parameters. Within this framework, persistence is interpreted as the measurable expression of underlying market memory rather than as a separate or purely technical construct.

To operationalise this perspective, the thesis addresses three interdependent challenges. First, persistence is inherently system-dependent, such that its estimation is meaningful only if the underlying system is coherently defined. Conventional approaches often estimate persistence from isolated time series without establishing whether they reflect the internal dynamics of a connected market structure. This thesis therefore develops a spillover-based system definition, in which variables are selected based on directional shock transmission and their roles as material transmitters or receivers within a connected network. This approach ensures that persistence is measured within an empirically coherent system rather than an arbitrary collection of predictors.

Second, the thesis examines how persistence influences price formation through the interaction of backward-looking historical dependence and forward-looking anticipatory effects. Financial markets differ markedly in the extent to which prices are anchored in slow-moving fundamentals or shaped by expectations and sentiment. By comparing a fundamentally driven market (Brent crude oil) with a sentiment driven market (Bitcoin), and measuring persistence using bidirectional memory estimates within a fractional cointegration framework, the thesis shows that persistence captures distinct information about how prices adjust, how expectations are incorporated, and why forecastability varies across market environments.

Finally, the thesis integrates these insights into a practical, memory-aware asset pricing framework. Traditional models often abstract from serial dependence and higher-order distributional effects, limiting their relevance for long-horizon investors. By explicitly

incorporating persistence into a multi-period asset pricing setting, the thesis shows how memory shapes effective risk, return dynamics, and long-term portfolio performance.

Overall, the thesis shows both analytically and empirically that persistence is an economically meaningful feature of financial markets that shapes predictability, price formation, and long-horizon investment outcomes. First, it demonstrates that defining forecasting systems via spillover connectedness improves stability and forecast performance. Second, it shows that fundamentally driven markets exhibit stronger backward-looking dependence, whereas speculative markets are dominated by anticipatory dynamics, with forecast accuracy improving as the relative strength of anticipatory forces increases. Third, it establishes that embedding memory into asset pricing and portfolio frameworks improves risk-adjusted portfolio performance and long-term efficiency. These findings imply that persistence should be treated as a structural input in forecasting design and as an intertemporal risk characteristic in investment practice.

1.2 Research Aims

The primary aim of this research is to develop a unified empirical and theoretical framework that establishes persistence as a measurable, interpretable, and economically meaningful property of financial systems. Rather than treating persistence as a statistical artefact, the thesis demonstrates how it reflects the propagation of information, beliefs, and expectations within interconnected markets, and how this structure can be exploited to improve forecasting, explain price formation, and enhance long-horizon investment outcomes.

To achieve this aim, the thesis integrates system identification, memory dynamics, and economic application within a single analytical framework. It seeks to demonstrate that persistence must be understood jointly as a system-dependent phenomenon, a driver of price formation operating through both historical dependence and anticipatory effects, and a priced intertemporal characteristic with direct implications for asset pricing and portfolio performance.

1.3 Research Objectives

Consistent with these aims, the thesis pursues a sequence of interconnected objectives. First, it designs and implements a spillover-based variable selection framework that defines financial systems through internal interdependence, ensuring that persistence is measured within coherent and empirically connected market structures. Second, it empirically analyses persistence in both fundamentally driven and sentiment-driven markets using backward and forward-looking memory measures, demonstrating that these

markets exhibit systematically different persistence profiles with direct implications for predictability.

Third, the thesis evaluates the forecasting relevance of anticipatory effects by embedding forward-looking memory into established econometric frameworks, including fractional cointegrated vector autoregressions and Vector Autoregressive (VAR)-based causality testing. Fourth, it derives closed-form expressions for effective variance and skewness under fractional integration and operationalises these results within a memory-augmented, multi-period asset pricing framework. Finally, the framework is validated empirically using US and Japanese equity markets, where portfolios constructed by sorting assets on their memory characteristics are shown to improve risk-adjusted returns and enhance long-horizon compounding efficiency.

1.4 Original Contributions of the Thesis

This thesis makes the following original contributions to the literature on financial market persistence, forecasting, and asset pricing:

- **Methodological innovation:** Operationalised existing spillover-based connectedness measures as a variable-selection criterion, using directional shock transmission to define coherent and internally connected financial systems for forecasting and reliable memory estimation.
- **Empirical characterisation:** Introduced and applied a bidirectional memory analysis[1] to distinguish and measure the roles of informational inertia and anticipatory effects in both fundamentally driven (Brent oil) and sentiment-driven (Bitcoin) markets, demonstrating how these memory profiles influence price formation and forecastability.
- **Econometric integration and application:** Integrated forward-looking persistence measures into established econometric frameworks (Fractionally Cointegrated Vector Autoregressive Model (FCVAR) and VAR-based Granger causality testing implemented via Wald restrictions) to assess the forecasting relevance of anticipatory dynamics.
- **Theoretical development:** Derived closed-form solutions for effective variance and skewness under fractional integration, and operationalised these within a novel, memory-augmented multi-period Capital Asset Pricing Model (CAPM), explicitly

[1] By bidirectional memory analysis, we refer to the estimation of persistence using both historical observations (backward memory) and future-realised trajectories (forward memory) in an ex-post sense. Forward memory does not imply reverse causality but measures the extent to which subsequent price dynamics were already implicit in the system at time t_0.

linking persistence to distributional moment effects that shape asset portfolio performance.

- **Practical validation:** Demonstrated, through comprehensive empirical analysis on US and Japanese equity markets, that the explicit incorporation of persistence characteristics into asset pricing and portfolio optimisation delivers substantial improvements in forecasting accuracy, risk-adjusted returns, and compound growth for long-term investors.

1.5 Thesis Structure

This thesis is structured as a series of three self-contained but thematically connected chapters.

Chapter 2 addresses the foundational challenge of system definition and variable selection in financial forecasting by introducing a spillover-based approach grounded in network theory. This method enables the construction of empirically coherent and internally connected systems in which persistence can be credibly estimated, reducing ambiguity about system boundaries and ensuring that memory parameters reflect genuine internal dynamics rather than artefacts of arbitrary model specification. Using data from the Brent and West Texas Intermediate (WTI) oil markets, the chapter demonstrates the critical role of systemic interdependence in supporting robust and stable forecasting performance.

Chapter 3 extends this foundation by investigating persistence as a bidirectional mechanism encompassing both backward-looking historical dependence and forward-looking anticipatory effects. Through a comparative analysis of the fundamentally driven Brent oil market and the sentiment-driven Bitcoin market, the chapter shows that different market environments exhibit distinct memory profiles with direct implications for price formation and forecastability. In particular, it establishes the empirical relevance of anticipatory dynamics and highlights the importance of integrating both backward and forward-looking memory into models of asset price behaviour.

Chapter 4 synthesises these insights within a novel, memory-augmented multi-period asset pricing framework. By incorporating fractional persistence into closed-form expressions for effective volatility and skewness, the chapter develops a tractable extension of the classical Capital Asset Pricing Model that links serial dependence to long-horizon wealth dynamics. The proposed framework is validated using data from US and Japanese equity markets, demonstrating that memory-aware portfolio strategies deliver superior risk-adjusted performance and improved long-horizon efficiency.

Collectively, these chapters establish a comprehensive and actionable framework that reinterprets persistence as a structural and behavioural attribute of financial systems,

with direct implications for forecasting design, price formation, and investment decision-making. By uniting behavioural mechanisms, econometric modelling, and practical application, the thesis advances understanding of how memory operates in financial markets and equips researchers and practitioners with tools for navigating complex, memory-driven environments.

Chapter 2

Improving System Generalisation & Forecastability via Spillover-Based Variable Selection Approach.

Abstract

Accurately identifying the determinants of oil spot prices remains a persistent challenge. This paper shows that a spillover-based approach to variable selection provides a stronger foundation for forecasting oil prices. Drawing on systems theory and the established financial contagion literature, we conceptualise markets as adaptive information networks and employ the Total Spillover Index (TSI) to quantify their dynamic interdependencies. Within a cointegration framework, we construct various models designed to capture the behaviour of Brent and WTI markets from both isolated and globally integrated perspectives. A dynamic spillover analysis, combining static, frequency domain and rolling window methods, demonstrates that models incorporating global indicators exhibit higher TSI values and respond more sharply to macroeconomic shocks, aligning closely with Global Economic Policy Uncertainty. Out-of-sample forecasts from a Fractional Cointegration VAR (FCVAR) and Long Short-Term Memory (LSTM) networks show that the global model consistently outperforms isolated market specifications, especially at medium and long-term horizons. Our results provide consistent theoretical and empirical evidence that information coherence, capturing the structured interdependencies across markets, should guide variable selection in complex systems such as oil markets.

Keywords: variable selection, spillover analysis, FCVAR, LSTM, forecast accuracy.

JEL classification: C22; C53; G17; Q41

2.1 Introduction

Forecasting models in finance and energy economics ultimately depend on the quality of the variables that enter them. Selecting these variables determines whether a model captures the genuine structure of information transmission or merely fits transient correlations. In practice, variable selection in multivariate settings is usually a two-stage process, in which candidate variables are first selected for their economic relevance, and then retained or excluded according to their incremental forecasting contribution or parsimony criteria such as the Akaike or Bayesian information criteria (Stock and Watson, 2002; Inoue and Kilian, 2006; Giacomini and Rossi, 2017). While this approach is well established, it implicitly assumes that information enters the model exogenously. Such treatment neglects the endogenous propagation of shocks and expectations that characterises modern, highly interconnected markets (Billio et al., 2012; Andries et al., 2019).

A large literature documents that multivariate dependence in macro-financial and commodity systems is shaped by the joint dynamics of prices, inventories, liquidity, and uncertainty. Vector Autoregression (VAR) and Vector Error Correction (VECM) models have long been used to capture such interrelations (Sims, 1980; Johansen, 1995b), while the generalised forecast error variance decomposition (GFEVD) has served as the principal tool for quantifying how shocks in one variable transmit to others. Building on this foundation, the connectedness framework of Diebold and Yilmaz (2009, 2012) interprets the GFEVD as a system-wide measure of spillovers, with frequency domain extensions capturing horizon-specific transmission patterns (Baruník and Křehlík, 2018). In commodity markets, studies such as Kilian (2009), Antonakakis and Gabauer (2018) and Zhou and Duan (2021) show that oil prices are embedded in dense networks of financial and macroeconomic linkages. Yet, despite this extensive evidence, spillover measures are almost exclusively used to describe transmission patterns or systemic risk, rather than to determine which variables should constitute a forecasting model.

We propose reframing variable selection as the detection of endogenous attention within the system, the extent to which variables respond to and amplify mutual innovations. When such mutual responsiveness is present, the system behaves as a self-contained information network rather than as a collection of isolated predictors. This notion draws from systems theory (Simon, 1962), signal extraction models of information diffusion (Grossman and Stiglitz, 1980; Hong and Stein, 1999), and the financial contagion literature that studies how shocks propagate through network linkages (Allen and Gale, 2000; Diebold and Yilmaz, 2012). These strands collectively suggest that the degree of

internal information circulation governs the ability of a system to generalise, with more interconnected systems producing more stable and predictable aggregate outcomes.

To operationalise this principle, we introduce a spillover-based variable selection framework. Using the Total Spillover Index (TSI), we measure how much of each variable's forecast error variance is explained by shocks originating in others within a system. Variables exhibiting stronger spillovers constitute the system's informational core, whereas weakly connected variables are peripheral. Conceptually, higher spillovers imply greater internal coherence within the model, which enhances its ability to capture persistent cross-variable effects.

Our framework complements rather than replaces existing approaches to variable selection. Penalised regression and sparse model methods such as the LASSO and elastic net (Tibshirani, 1996; Zou and Hastie, 2005), Bayesian model averaging (Raftery, 1997), and dynamic factor or state-space frameworks (Bai and Ng, 2008; Forni and Gambetti, 2014) focus primarily on parsimony or dimensionality reduction. In contrast, existing connectedness studies use spillovers primarily to monitor systemic risk or to characterise the transmission structure of financial and commodity markets (Billio et al., 2012; Diebold and Yilmaz, 2009, 2012; Baruník and Křehlík, 2018; Antonakakis and Gabauer, 2018; Zhou and Duan, 2021). The methodological gap we address is the absence of a procedure that uses spillover intensity as an explicit, endogenous criterion for selecting variables that sustain dynamics within the system.

Empirically, we apply this framework to Brent and WTI oil benchmarks, analysing them both as isolated markets and as components of a broader global system encompassing variables that represent expansion, liquidity, and regime dynamics. Holding model size and estimation class fixed, a higher TSI is interpreted as greater internal coherence of the information set. Out-of-sample analyses using both an FCVAR model and LSTM networks confirm that models constructed from highly interdependent variables consistently outperform conventional specifications.

The contributions of this study are threefold. First, it formalises a systems-theoretic rationale for variable selection, linking endogenous information flow to forecast generalisability. Second, it introduces a methodological innovation by integrating spillover analysis into the variable selection stage. Third, it provides empirical evidence that systems characterised by stronger information circulation yield superior forecasts, highlighting the importance of recognising endogenous interdependence when modelling complex, adaptive markets.

The remainder of the paper is organised as follows. Section 2.2 reviews the literature on variable selection, oil price forecasting, and spillover analysis, positioning the proposed framework within existing econometric and network-based approaches. Section 2.3 describes the dataset, construction of global factors, and validation procedures. Section 2.4 details the empirical methodology, linking the spillover-based variable selection to the

forecasting frameworks used for validation. Section 2.5 presents the empirical analysis, including static, frequency, and dynamic spillover results, together with forecast evaluation. Section 2.6 provides a focused concept validation, combining FCVAR performance, formal forecast comparison tests, and the LSTM benchmark. Finally, Sections 2.7 & 2.8 conclude with the policy and theoretical implications, outlining how the framework advances the understanding of information propagation in adaptive market systems.

2.2 Literature Review

The literature on variable selection, multivariate modelling, and market interdependence is extensive, yet fragmented across econometric, network, and forecasting traditions. This section synthesises these strands with an explicit emphasis on the methodological lineage from VAR and VECM models to generalised variance decomposition and modern spillover frameworks.

Traditional variable selection typically relies on economic judgement, benchmark predictors, and parsimony criteria. Penalised regression approaches such as the LASSO and elastic net shrink coefficients to induce sparsity (Tibshirani, 1996; Zou and Hastie, 2005), while Bayesian model averaging accounts for specification uncertainty by weighting models by posterior probabilities (Raftery, 1997). Dimensionality reduction methods, including dynamic factor and state-space models (Bai and Ng, 2008; Forni and Gambetti, 2014), summarise common variation through latent components. Although these methods improve tractability in high-dimensional systems, they prioritise parsimony or compression rather than the strength of inter-variable feedback, treating interdependence as a statistical artefact rather than as an organising principle of model design.

A second strand of literature focuses on multivariate time series models rooted in the VAR and VECM. Since Sims (1980), VARs have been used to capture dynamic interactions among economic variables without strong identifying restrictions. VECMs extend this framework by incorporating long-run equilibrium relations among non-stationary variables, following the Johansen methodology (Johansen, 1995b). In both VAR and VECM systems, generalised forecast error variance decomposition is central to understanding how shocks propagate, providing a model-based decomposition of the share of each forecast error variance attributable to innovations in other variables. This decomposition underpins the modern connectedness literature and serves as the mathematical basis for spillover indices.

The transformation of FEVD into a system-wide measure of interdependence originated with the connectedness framework of Diebold and Yilmaz (2009, 2012). Their Total Spillover Index summarises the degree to which shocks in one variable transmit to others, shifting the interpretation of FEVD from a descriptive output of VAR models to a

structural diagnostic of information transmission. Subsequent contributions refined this framework for different forms of interdependence. Diebold and Yılmaz (2014) embedded spillovers within a network-theoretic representation, while Baruník and Křehlík (2018) introduced spectral decompositions that distinguish short and long-horizon spillovers using frequency domain methods. These developments show how VAR and VECM structures can be used not just for estimating interrelations but also for quantifying their horizon-specific effects.

The empirical relevance of spillover techniques is particularly evident in the oil market. Studies based on VAR and VECM methodologies such as Kilian (2009), Baumeister and Peersman (2013), and Zhang et al. (2019) document the presence of strong cross-market interactions and shared responses to macroeconomic shocks. Spillover applications, including Antonakakis and Gabauer (2018), Reboredo and Ugolini (2016), and Zhou and Duan (2021), show that oil markets exhibit substantial short and long-run connectedness with financial assets, exchange rates, and uncertainty indicators. These studies consistently demonstrate that oil price behaviour is governed by system-wide interactions rather than isolated causal channels, reinforcing the suitability of spillover measures for assessing informational structure in commodity markets.

A separate but related literature examines forecasting frameworks. Econometric models such as VECM and FCVAR incorporate long-run equilibrium adjustments and fractional persistence (Johansen and Nielsen, 2012a), while machine learning models such as LSTM networks capture nonlinear and long memory structures (Hochreiter and Schmidhuber, 1997b; Salisu and Vo, 2020a). Forecast comparison studies show that models incorporating economically meaningful variables or measures of interdependence tend to outperform purely algorithmic benchmarks (Malliaris and Malliaris, 2008; Guo et al., 2023a). However, this literature largely treats interdependence as an observable outcome of estimation rather than as a criterion for selecting variables ex ante.

Across these strands, a gap remains. Existing work acknowledges that oil markets and financial systems are interconnected, but no study systematically integrates the strength of dynamic interdependencies, quantified through spillover measures, into the variable selection stage of model construction. While spillovers are widely used to monitor systemic risk or describe interrelations, they have not been used to determine which variables enter a forecasting model. Similarly, although forecasting models benefit from including variables that capture market structure, no established methodology identifies such variables through a formal measure of multivariate coherence.

The spillover-based variable selection framework proposed in this chapter addresses this gap by using the Total Spillover Index as an endogenous mechanism for selecting variables that actively participate in the system's information transmission process. In doing

so, it embeds the connectivity structure of the market directly into model design, offering a conceptually grounded and empirically tractable route toward more generalisable forecasting systems.

2.3 Data and Model Specification

Our dataset comprises eleven time series that capture the behaviour of the Brent and WTI benchmarks, together with variables that summarise global market dynamics. All data were obtained from Bloomberg, and Table 2.1 reports the corresponding tickers and construction details. Specifically, the dataset includes:

- **Spot Prices**: Daily spot prices for Brent (BSP) and WTI (WSP), obtained from the ICE (London) and NYMEX (U.S.) exchanges.
- **Market Variables (Isolated Systems)**: Measures of market *Expansion*, *Liquidity*, and *Regime* for each benchmark, denoted as BE, BL, and BR for Brent, and WE, WL, and WR for WTI. These variables are derived from the futures term structure following the methodology of Mustanen et al. (2022).
- **Global Market Variables**: Corresponding measures of expansion (E), liquidity (L), and regime (R) are derived using Principal Component Analysis (PCA) applied to the combined set of spot-price series and market variables from the isolated systems.

Table 2.1 Bloomberg tickers and construction of variables used in the spillover analysis.

Variable	Description	Derivation Method	Bloomberg Ticker
BSP	Brent spot price	Observed series	CO1 Comdty
WSP	WTI spot price	Observed series	CL1 Comdty
BE	Brent market expansion	Area under Brent futures curve	CO1–CO36 Comdty
BL	Brent liquidity	Area under Brent open interest curve	CO1–CO36 Comdty
BR	Brent regime	Slope of Brent futures curve	CO1–CO36 Comdty
WE	WTI market expansion	Area under WTI futures curve	CL1–CL36 Comdty
WL	WTI liquidity	Area under WTI open interest curve	CL1–CL36 Comdty
WR	WTI regime	Slope of WTI futures curve	CL1–CL36 Comdty
E	Global expansion factor	PCA of (BE, WE)	Constructed
L	Global liquidity factor	PCA of (BL, WL)	Constructed
R	Global regime factor	PCA of (BR, WR)	Constructed

Note: Table presents the list of variables selected for the analysis. Futures-based variables (BE, BL, BR for Brent and WE, WL, WR for WTI) follow the construction in Mustanen et al. (2022) and use the 1–36 month futures and open interest curves (CO and CL). The global factors (E, L, R) are principal components of the corresponding Brent–WTI variable pairs and therefore have no direct Bloomberg tickers.

The combined use of spot and futures-derived variables is essential because these markets encode different but complementary dimensions of information. Spot prices capture the realised state of physical supply–demand conditions, whereas the futures curve reflects market expectations, convenience yields, hedging activity, and inventory fundamentals (Hamilton, 2009; Alquist and Gervais, 2013; Singleton, 2014). This informational content has been shown to lead spot price movements, particularly around episodes of structural tightness and regime shifts (Kilian and Murphy, 2014). Integrating both sources enables the variable set to reflect the two-way interaction between realised conditions and expectations, which is central to the spillover dynamics examined in this chapter.

The dataset spans twenty years, from 02-01-2003 to 12-01-2023, yielding a total of 5014 daily observations. To ensure consistency across variables, we apply *Z-score* normalisation:

$$Z_{i,t} = \frac{x_{i,t} - \mu_i}{\sigma_i}, \tag{2.1}$$

where $x_{i,t}$ is the original value of variable i at time t, μ_i is its sample mean, and σ_i is its sample standard deviation. Summary statistics for the dataset are presented in Table 2.2.

TABLE 2.2 : Summary statistics.

Market	Variable	Mean	SD	Median	Max	Min	Skew.	Kurt.	Jarque-Bera	P-Val
Spot	BSP	0.00	1	-0.16	2.76	-1.96	0.30	2.14	233.48	0.00
	WSP	0.00	1	-0.15	3.24	-4.39	0.31	2.45	145.73	0.00
Brent	BE	0.00	1	-0.14	3.17	-1.99	0.19	2.48	86.09	0.00
	BL	0.00	1	-0.36	2.20	-1.74	0.46	1.76	498.82	0.00
	BR	0.00	1	0.00	3.31	-3.79	0.02	3.08	1.52	0.47
WTI	WE	0.00	1	-0.12	3.53	-1.96	0.20	2.73	47.44	0.00
	WL	0.00	1	-0.05	2.21	-2.00	0.00	2.18	138.81	0.00
	WR	0.00	1	-0.05	2.94	-4.08	-0.21	3.49	85.68	0.00
Global	E	0.00	1	-0.12	2.71	-2.57	0.30	2.11	239.97	0.00
	L	0.00	1	-0.10	2.36	-2.33	0.23	1.95	272.58	0.00
	R	0.00	1	0.11	2.12	-2.77	-0.57	2.58	310.81	0.00

Note: Summary statistics of the z-score normalised variables used in the spillover analysis. The sample covers daily data from 02.01.2003 to 12.01.2023. Variable definitions and Bloomberg tickers are reported in Table 2.1.

2.3.1 Model Specification

To capture the multifaceted dynamics of oil markets, we construct six distinct models that reflect both isolated and global perspectives. Table 2.3 details the specific model configurations. These models are developed within a cointegration framework to ensure that long-run equilibrium relationships among variables are appropriately captured. The models are subsequently converted into a VAR representation to facilitate spillover analysis.

TABLE 2.3 : Model Selection Summary.

Model	List of Variables	Model Context
1	BSP; BE; BL; BR	Isolated market system
2	WSP; WE; WL; WR	
3	BSP; BE; BL; BR; WSP; WE; W L; WR	
4	BSP; E; L; R	Global market system
5	WSP; E; L; R	
6	BSP; WSP; E; L; R	

Note: Variable List: BSP - Brent spot; BE - Brent market expansion; BL - Brent market liquidity; BR - Brent market Regime; WSP - WTI spot; WE - WTI market expansion; WL - WTI market liquidity; WR - WTI market regime; E - global market expansion; L - global market liquidity; R - global market regime.

The primary aim of this model specification is to identify systems where the variables exhibit the strongest dynamic interdependencies. By comparing models based on variables describing isolated systems with those incorporating global variables, we assess the extent to which broader market interconnections contribute to forecast performance. The integrated models further highlight the mutual influence between the Brent and WTI markets, providing insights into how combining variable sets can improve forecast robustness.

2.4 Empirical Methodology

This section outlines the methodological framework used to evaluate the explanatory power of our variable selection approach. Our framework comprises three interrelated components: (i) preliminary tests, (ii) spillover analysis, and (iii) a forecastability validation exercise.

2.4.1 Preliminary Tests and VECM Specification

We begin by assessing the order of integration of all variables, as cointegration-based modelling requires that the underlying series behave as $I(1)$ processes. To ensure robustness, we apply two complementary tests with opposing null hypotheses: the Elliott–Rothenberg–Stock DF–GLS test (Elliott et al., 1996), which tests the null of a unit root, and the Kwiatkowski–Phillips–Schmidt–Shin (KPSS) test (Kwiatkowski et al., 1992), which tests the null of level stationarity. The DF–GLS procedure offers improved size and power relative to conventional ADF tests, particularly in the presence of persistence, heteroskedasticity and near–unit–root behaviour, while the KPSS test guards against spurious non-rejection of unit-root tests. Using tests with opposing nulls provides a diagnostic cross-check: failure to reject a unit root in DF–GLS together with

TABLE 2.4 : Unit Root Test Results: DF–GLS vs KPSS.

Market	Spec / Type	Test / Model	DF–GLS (ERS) Constant				KPSS Constant			
	Variable	Transform	Stat	[5%]	Reject?	Opt. Lag	Stat	[5%]	Reject?	Opt. Lag
SPOT	BSP	Level	-1.107	-1.948	N	30	1.040	0.463	Y	30
		First Diff	-2.172	-1.948	Y	32	0.098	0.463	N	30
	WSP	Level	-1.342	-1.948	N	29	0.827	0.463	Y	30
		First Diff	-2.435	-1.948	Y	32	0.068	0.463	N	30
BRENT	BE	Level	-0.736	-1.948	N	30	1.268	0.463	Y	30
		First Diff	-3.979	-1.948	Y	32	0.172	0.463	N	30
	BL	Level	-1.619	-1.948	N	30	6.432	0.463	Y	30
		First Diff	-13.151	-1.948	Y	29	0.170	0.463	N	30
	BR	Level	-2.765	-1.948	Y	18	0.696	0.463	Y	30
		First Diff	-4.821	-1.948	Y	17	0.035	0.463	N	30
WTI	WE	Level	-0.779	-1.948	N	32	1.235	0.463	Y	30
		First Diff	-9.904	-1.948	Y	32	0.172	0.463	N	30
	WL	Level	-0.487	-1.948	N	32	7.501	0.463	Y	30
		First Diff	-4.824	-1.948	Y	32	0.122	0.463	N	30
	WR	Level	-3.077	-1.948	Y	32	0.920	0.463	Y	30
		First Diff	-4.856	-1.948	Y	31	0.035	0.463	N	30
GLOBAL	E	Level	-1.420	-1.948	N	32	0.757	0.463	Y	30
		First Diff	-3.013	-1.948	Y	32	0.071	0.463	N	30
	L	Level	-1.593	-1.948	N	23	7.003	0.463	Y	30
		First Diff	-4.445	-1.948	Y	22	0.066	0.463	N	30
	R	Level	-0.782	-1.948	N	32	1.977	0.463	Y	30
		First Diff	-5.437	-1.948	Y	31	0.154	0.463	N	30

Note: DF-GLS denotes the Elliott-Rothenberg-Stock unit-root test under a constant-only specification (null: unit root). KPSS denotes the Kwiatkowski-Phillips-Schmidt-Shin test under a constant-only specification (null: stationarity). For DF=GLS, optimal lag lengths are selected using the automatic information criterion procedure. For KPSS, the reported lag corresponds to the Newey-West truncation parameter, set using the standard automatic rule with a maximum bandwidth of 31. Rejection of the KPSS null at the 5% level indicates evidence against stationarity. Mixed evidence for BR and WR reflects near unit-root behaviour under persistence and structural shifts. Overall, the joint evidence supports treating all series as $I(1)$ for Johansen cointegration and VECM modelling.

rejection of stationarity in KPSS supports an $I(1)$ interpretation. The results, summarised in Table 2.4, show that the majority of variables fail to reject the unit-root null in levels but reject it strongly after first differencing, while KPSS rejects stationarity in levels and typically fails to reject it after differencing. Two variables, the Brent and WTI regime slope series (BR and WR), exhibit marginal DF–GLS rejections in levels under the constant-only specification. However, the literature documents that term structure slope measures in commodity markets are characterised by strong persistence, episodic mean reversion and structural breaks (Fama and French, 1987; Hamilton, 2009; Alquist and Gervais, 2013; Kilian and Murphy, 2014), features that are known to induce spurious rejections in unit-root testing even when the underlying process is effectively integrated (Perron, 1989; Ng and Perron, 2001). We therefore interpret the level rejections for BR and WR as reflecting near unit–root dynamics rather than genuine stationarity.

Jointly, the evidence supports treating all variables as $I(1)$ for the purposes of Johansen cointegration and VECM modelling.

We then estimate a VECM to capture the long-run equilibrium structure among the variables. Lag lengths are selected using Akaike's Information Criterion (AIC) and validated with multivariate Ljung–Box Q-tests to ensure residual serial independence. The number of cointegrating relationships is determined using Johansen's trace and maximum-eigenvalue tests (Johansen and Nielsen, 2012a), and further Q-tests confirm the stability of each specification. The resulting lag and rank selections are reported in Table 2.5.

TABLE 2.5 : VECM Model Parameters and Complexity Summary.

Model	k **vars**	p **lag**	r **rank**	**No. Params**	Q_ϵ	**P-val**	**AIC**	**BIC**	**LogL**
Model 1	4	7	1	136	66.53	1.00	-66065.8	-65257.4	33156.92
Model 2	4	6	2	112	77.76	1.00	-49466.1	-48729.4	24846.07
Model 3	8	6	4	448	309.89	1.00	-124729.4	-121854.0	62805.68
Model 4	4	6	2	112	68.93	1.00	-65505.1	-64768.4	32865.57
Model 5	4	6	2	112	81.78	1.00	-59697.8	-58961.1	29961.92
Model 6	5	6	2	170	115.73	1.00	-85691.5	-84570.1	43017.76

Note: (i) k denotes the number of endogenous variables, p the lag order, and r the cointegration rank (Johansen). (ii) "No. Params" indicates the approximate number of freely estimated short- and long-run coefficients in the VECM representation. (iii) Q_ϵ: Ljung–Box Q-statistic for residual autocorrelation (p-value in brackets). (iv) AIC and BIC: Akaike and Bayesian information criteria; LogL: log-likelihood for each model. Model list — M1: BSP, BE, BL, BR; M2: WSP, WE, WL, WR; M3: BSP–WR combined; M4: BSP with (E, L, R); M5: WSP with (E, L, R); M6: BSP, WSP, (E, L, R).

Once a stable VECM specification is obtained, it is transformed into its implied VAR representation. This provides the infinite-order moving-average (MA) form required for computing the generalised forecast-error variance decompositions that underpin the spillover framework, including the TSI and related measures.

2.4.2 Spillover Analysis

Spillover analysis, as outlined by Diebold and Yılmaz (2014), forms the core of our variable selection approach by quantifying the degree of interconnectedness among market variables. In this study, we adopt a framework that comprises three components: static (time-domain) analysis, frequency domain analysis, and dynamic (rolling window) analysis. For complete derivations and formulations, see Appendix A.

Based on the normalised decompositions, we employ several measures of interconnectedness to characterise the transmission of shocks among the variables in our models:

- **Total Spillover Index:** The Total Spillover Index measures the overall degree of shock transmission across all variables in the system over a forecast horizon h:

$$S.O^g(h) = \frac{\sum_{\substack{i,j=1 \\ i \neq j}}^{N} \tilde{\tau}_{ij}^g(h)}{N} \times 100, \tag{2.2}$$

  where N denotes the number of variables in the system, i and j index the originating and receiving variables respectively, and $\tilde{\tau}_{ij}^g(h)$ represents the normalised generalised forecast error variance share from variable j to i at horizon h. The superscript g indicates that the decomposition is based on the generalised (Koop – Pesaran – Shin) approach, which is invariant to variable ordering.

- **Directional Spillovers:** These quantify the transmission of shocks in a specific direction. The spillover *transmitted by* variable i to all others ("TO") is defined as:

$$S.O_{i\cdot}^g(h) = \frac{\sum_{j \neq i} \tilde{\tau}_{ij}^g(h)}{\sum_{j=1}^{N} \tilde{\tau}_{ij}^g(h)} \times 100, \tag{2.3}$$

  while the spillover *received by* variable i from all others ("FROM") is given by:

$$S.O_{\cdot i}^g(h) = \frac{\sum_{j \neq i} \tilde{\tau}_{ji}^g(h)}{\sum_{j=1}^{N} \tilde{\tau}_{ji}^g(h)} \times 100. \tag{2.4}$$

  The "TO" measure therefore captures how much of the forecast–error variance in other variables is explained by shocks to i, whereas the "FROM" measure captures the extent to which shocks in other variables contribute to the forecast–error variance of i.

- **Net Spillover:** The net spillover indicates whether variable i acts as a net transmitter or receiver of volatility shocks:

$$S.O_i^g(h) = S.O_{i\cdot}^g(h) - S.O_{\cdot i}^g(h). \tag{2.5}$$

  Positive values imply that i is a net transmitter of shocks (its innovations influence other variables more than it is influenced by them), while negative values indicate a net receiver.

Frequency Domain Spillover Analysis. While the time domain spillover analysis provides an overall measure of interconnectedness, it does not capture how spillover effects vary across different time horizons. To address this, we extend the generalised variance decomposition (GVD) framework to the frequency domain using the spectral decomposition method proposed by Baruník and Křehlík (2018). This approach partitions the variance contributions into distinct frequency bands, allowing us to differentiate between short and long-term spillover effects. By analysing the spectral components,

we assess whether the bulk of the interdependencies arises from high-frequency (short-term) or low frequency (long-term) dynamics. The frequency domain analysis is especially valuable in energy markets, where both short-run speculative trading and long-run fundamental factors play crucial roles.

Dynamic (Rolling Window) Spillover Analysis. Recognising that interconnectedness evolves over time, we perform a dynamic spillover analysis using a rolling–window framework. Specifically, the VAR/VECM system is re-estimated over fixed windows of 490 and 1,225 observations (approximately two and five years at daily frequency). These horizons align with the typical duration of business cycles, which empirical research places between two and seven years across most advanced and emerging economies (Pedersen, 1998; Stock and Watson, 1999; Gonzalez et al., 2015; Kohlscheen, 2023). The shorter window captures transient shocks and high–frequency adjustment, whereas the longer one smooths local fluctuations and reflects broader cyclical reorganisation. This specification ensures that each rolling window spans approximately one full business cycle, balancing sensitivity and stability in connectedness dynamics. The resulting time–varying TSI is then analysed jointly with the Global Economic Policy Uncertainty (GEPU) Index to examine whether periods of heightened uncertainty coincide with intensified market interdependence.

2.4.3 Variable Forecastability Validation

The objective of this exercise is to assess whether the interconnectedness that defines spillover–based variable selection enhances forecasting performance, regardless of the forecasting framework employed. To this end, we evaluate the predictive ability of models with different degrees of system coherence, measured by their TSI, across two distinct modelling approaches: (i) an FCVAR model and (ii) a nonlinear, machine-learning LSTM model. The FCVAR experiment tests the internal consistency of the spillover, forecastability link within a traditional econometric setting, whereas the LSTM benchmark extends this validation to a data–driven, nonlinear context. The goal is therefore not to compare FCVAR and LSTM per se, but to demonstrate that stronger interconnectedness among variables consistently translates into superior predictive performance, independent of the model class.

FCVAR Approach. We estimate FCVAR models for three configurations, Model 1 (isolated Brent), Model 2 (isolated WTI), and Model 6 (hybrid global). This design enables a structured comparison between systems characterised by lower and higher interconnections. The FCVAR framework is employed to test whether variable sets

with stronger spillover linkages exhibit greater predictive efficiency within a cointegration environment. Detailed estimation and forecasting procedures are outlined in Appendix A.2.

Robustness Check - LSTM Benchmark. To verify that the hypothesised relationship between interconnectedness and predictability is not model-specific, we also employ an LSTM network following the design of Guo et al. (2023a). LSTMs are recurrent neural networks capable of learning nonlinear and long-range temporal dependencies (Hochreiter and Schmidhuber, 1997a; Fischer and Krauss, 2018), enabling them to approximate complex dynamics without imposing parametric restrictions. Model 6 serves as the test case and is benchmarked against an LSTM trained on the Guo et al. (2023a) variable set summarised in Table 2.6.

TABLE 2.6 : Guo et al. (2023b) variables.

Money market variables:	
ER	Exchange Rate
IR	Interest Rate
TR	Treasury Rate
Futures market variables:	
INE	Shanghai International Energy Exchange Commodity Futures
CF	South China Commodity Index
WTI	WTI Spot
Brent	Brent Spot
Stock market variables:	
SHCI	Shanghai Composite Index
EI	CSI Energy Index - representation of energy sector stock market
Economic variables:	
EPU	Economic Policy Uncertainty

Note: Data span: from 26-03-2018 to 12-01-2023. Data frequency: daily.

The network architecture and hyperparameter settings are reported in Table 2.7. To ensure comparability with the analogue study and to avoid look-ahead bias, the LSTM is trained using a chronological 90%–10% train/test split, a common design in neural-network forecasting where large training samples are required to learn stable temporal features (Goodfellow et al., 2016). A comparative spillover analysis is then conducted on both datasets to assess whether differences in their TSI values are reflected in forecast performance, thereby evaluating the robustness of the spillover-based variable selection framework across distinct modelling methodologies.

Rolling Window Forecasting and Evaluation. In both FCVAR and LSTM implementations, spot price forecasts for Brent and WTI are generated at multiple horizons

TABLE 2.7 : LSTM network architecture and hyperparameter settings across forecast horizons.

Config	Component	Spec	Activ.	Input/Output	Notes
Setup	LSTM Layer	100 units	tanh	Input: (10, 7)	Sequence length $k = 10$
	Dense Layer	4 neurons	Linear	Output: $n = \{1, 5, ...\}$	Multi-horizon forecasts
Training	Optimizer	Adam	–	–	Learning rate = 0.001
	Loss	`logcosh`	–	–	Robust to outliers
	Batch Size	32	–	–	Mini-batch training
	Epochs	200	–	–	Maximum iterations
	Val. Method	Test-set val	–	–	on 10% of total set
	Early Stop	Patience = 10	–	–	Prevents overfitting

Note: The LSTM model follows the analogue setup of Guo et al. (2023a) for comparability with the machine learning benchmark. Data are first split chronologically into 90% for training and 10% for testing. A further 20% of the training set is used as a validation holdout. The input tensor has shape $(10, 7)$, reflecting a sequence of ten trading days and seven predictor variables. Forecast horizons are $n = \{1, 5, 10, 20, 40, 60, 80, 100\}$. All input variables are normalised to the $[0, 1]$ range. Training employs the Adam optimiser (learning rate = 0.001) with a `logcosh` loss function and early stopping (patience = 10). The model is implemented in TensorFlow 2.10.

(1-day, 30-day, and 60-day) using a rolling window approach to capture evolving market conditions. Forecast performance is evaluated using Linear Exponential (LinExp) Loss, Mean Absolute Percentage Error (MAPE), and Relative Forecast Error (RFE). For LinExp Loss, we use an asymmetry coefficient $\alpha = -1.5$ to amplify over-prediction costs (Salisu and Vo, 2020b). RFE is defined as:

$$\text{RFE} = \left(\frac{\text{Loss}_{\text{Model A}}}{\text{Loss}_{\text{Model B}}} - 1 \right) \times 100, \tag{2.6}$$

where a negative RFE indicates that Model A outperforms Model B.

2.5 Empirical Results

In this section, we present the key empirical findings derived from our spillover analysis across different models and forecast horizons. We first summarise the TSI results in both static and frequency domain contexts, followed by a discussion of the dynamic (rolling window) analysis and its relationship with global economic policy uncertainty. Detailed model-by-model results and additional tables are provided in Appendix C.

2.5.1 Spillover Analysis

Table 2.8 reports the TSI values for Models 1 to 6 under different forecast horizons (30, 60, 90, 120, and 245 days) and for three decomposition ranges: static, 1–5 days (high frequency), and 5 days and above (low frequency). These figures provide a concise

comparison of spillover behaviour in the systems from the perspectives of an isolated and global market.

From the summary in Table 2.8, several observations stand out. First, static TSI values differ noticeably from those obtained under the 1–5 day and 5-day-to-infinity decompositions. This indicates that static analysis alone is insufficiently discriminating to determine an optimal variable combination when the forecast horizon is critical. To identify the most effective variable set, both short and long-term spillover effects must be taken into consideration.

From an isolated versus global variable perspective, models incorporating global market variables (M4, M5, M6) generally exhibit higher TSI values compared to their isolated market counterparts (M1, M2, M3). This difference is particularly pronounced in the 1 to 5 day decomposition, where models based on global variables more effectively capture the internal dynamics of the system than those based on isolated market data. These findings suggest that including global dynamics increases measured interconnectedness and enhances the system's ability to explain its overall dynamics.

The forecast horizon sensitivity analysis highlights distinct patterns between isolated and global market models. Specifically, as the forecast horizon increases from 30 to 245 days, the TSI for systems based on isolated market variables tends to rise, indicating that longer horizon forecasts capture a greater degree of latent, slow-moving interdependencies. This behaviour is consistent with evidence that cross-market linkages strengthen over longer horizons as information diffuses gradually through production, inventory, and hedging channels (Diebold and Yılmaz, 2014; Baruník and Křehlík, 2018). Studies such as Marcellino et al. (2006) and Antonakakis and Gabauer (2017) similarly report that forecast performance and connectedness at longer horizons reflect persistent common components rather than high-frequency volatility transmission.

Conversely, for models incorporating global market variables, the TSI generally decreases with increasing forecast horizon, implying that these systems are dominated by short-term contagion and event-driven effects, a pattern consistent with the transient nature of global shocks documented by Baruník and Křehlík (2018). This divergence suggests that while isolated market systems benefit from the accumulation of long-run equilibrium forces, globally integrated systems are primarily shaped by immediate, short-lived interactions. Hence, the appropriate choice of forecast horizon should align with the temporal scale of interdependencies captured by the selected set of variables.

2.5.2 Network Analysis

To visualise the direction and strength of spillovers, we construct network graphs for each model. Figure 2.1 presents a composite network graph summarising the spillover dynamics across all six models based on a 60-day forecast horizon. Each node represents

TABLE 2.8 : Spillover TSI Summary for Models M1 vs M4, M2 vs M5, and M3 vs M6.

Decomposition Type		**Static**		**1 to 5 days**		**5 to inf**	
Comparison	**Horizon**	**TSI 1**	**TSI 2**	**TSI 1**	**TSI 2**	**TSI 1**	**TSI 2**
M1 vs M4	30	40.19%	43.21%	32.12%	44.76%	40.19%	43.21%
	60	40.33%	42.56%	33.71%	41.30%	40.34%	42.56%
	90	40.42%	41.98%	34.42%	40.92%	40.42%	41.98%
	120	40.48%	41.53%	34.78%	41.79%	40.48%	41.53%
	245	40.58%	41.45%	35.34%	46.81%	40.58%	41.44%
M2 vs M5	30	34.44%	39.69%	11.70%	48.86%	34.69%	39.64%
	60	35.88%	37.88%	11.29%	51.03%	36.05%	37.82%
	90	37.14%	37.13%	11.33%	52.23%	37.28%	37.07%
	120	38.36%	37.19%	11.54%	52.44%	38.50%	37.13%
	245	43.38%	40.80%	12.52%	49.38%	43.49%	40.77%
M3 vs M6	30	55.01%	53.75%	23.61%	54.33%	55.28%	53.75%
	60	57.51%	52.12%	23.35%	57.80%	57.72%	52.11%
	90	59.39%	51.16%	22.96%	59.69%	59.56%	51.13%
	120	60.73%	50.68%	22.53%	60.68%	60.88%	50.65%
	245	63.80%	50.97%	21.04%	61.60%	63.90%	50.95%

Note: Following Diebold and Yilmaz (2012). "TSI" - Total Spillover Index. This table presents the TSI for Models M1 - M6 at various forecast horizons (30, 60, 90, 120, and 245 days), decomposed into *Static*, *1 to 5 days* (short-run), and *5 to infinity* (long-run) frequency domains. All values are expressed as percentages, reflecting the level of interconnectedness within each model configuration. Higher TSI values indicate greater spillovers (stronger interdependencies) among the variables in the system. Models M1, M2, and M3 represent isolated market perspectives, whereas M4, M5, and M6 incorporate global market variables.

a variable and the edges denote directional spillovers. The thickness and colour intensity of the edges correspond to the magnitude of the spillover effect. Overall, the composite graph provides an intuitive visual summary of the dynamic relationships among variables and highlights key differences between isolated and global market perspectives.

Several key observations emerge from the composite network graph analysis. First, from the perspectives of the isolated versus global models, the models that incorporate global market variables (Models 4, 5, and 6, right-hand side graphs) exhibit denser interconnections compared to their isolated market counterparts (Models 1, 2, and 3, left-hand side graphs). This suggests that global dynamics enhance the overall interconnectedness of the system.

When considering benchmark specific dynamics, the isolated models (Models 1 and 2) show that the spot prices (BSP for Brent and WSP for WTI) play pivotal roles either as major transmitters or receivers of shocks. In contrast, when global variables are included (Models 4 and 5), the network structure shifts, with global expansion (E) frequently emerging as a dominant transmitter.

From an integrated market perspective, the bottom panel graphs, which compare Model 3 (combined isolated) and Model 6 (combined global), illustrate that incorporating global

FIGURE 2.1 : Network Analysis Summary of Intra-Market Dynamics in the Brent and WTI Oil Markets.

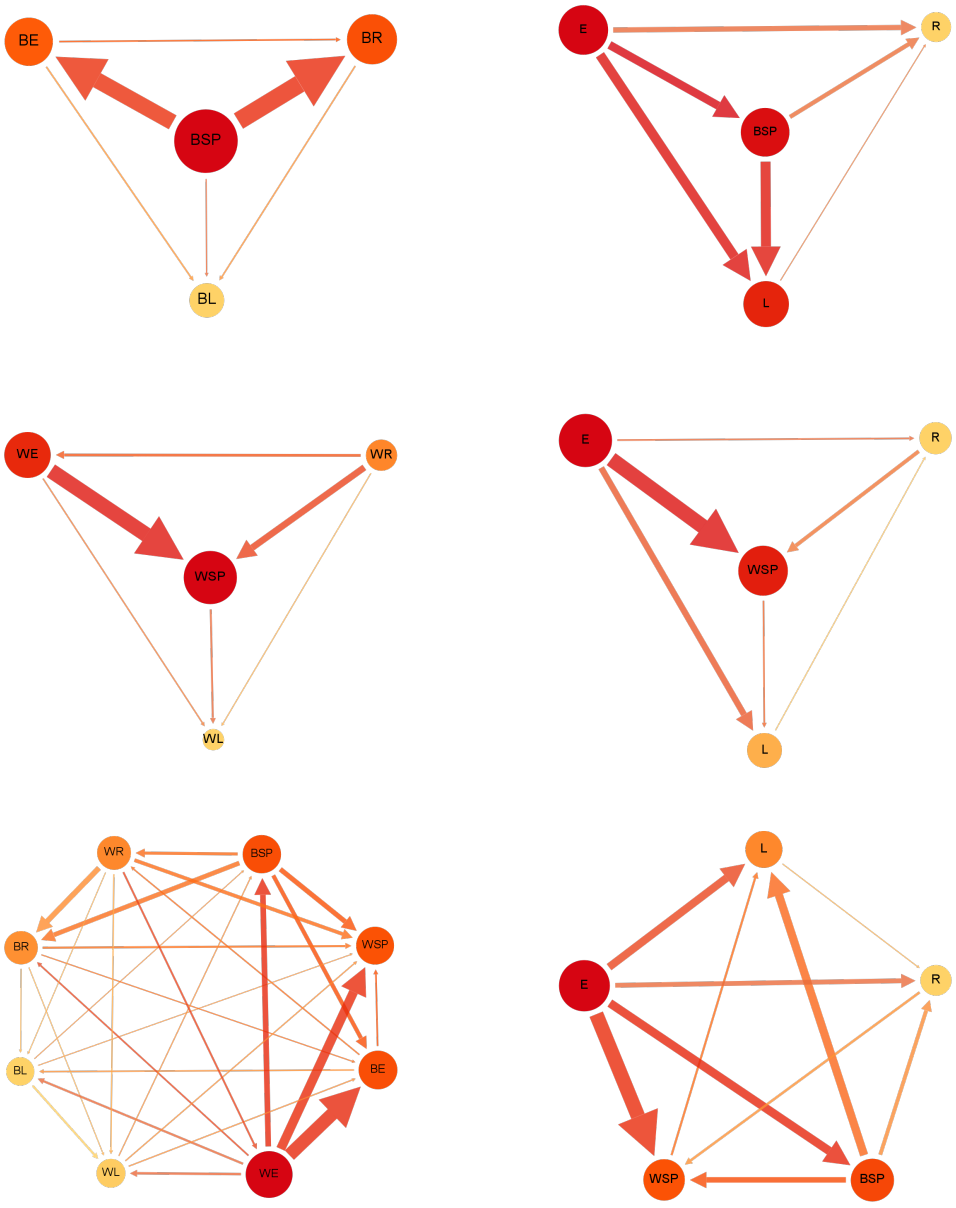

Note: This figure presents the network graph summaries for the six model specifications based on a 60-day forecast horizon. From left to right, the top panel displays Model 1 (isolated Brent market) and Model 4 (Brent with global market variables), the middle panel shows Model 2 (isolated WTI market) and Model 5 (WTI with global market variables), and the bottom panel illustrates Model 3 (combined isolated markets) and Model 6 (combined global markets). The thickness and direction of the edges represent the magnitude and direction of spillover effects between the variables.

market indicators leads to a more interconnected system. Model 6, in particular, exhibits the highest overall density of spillovers, consistent with its elevated TSI values. Although Model 3 appears to have a higher TSI than Model 6, as observed in Table 2.8, this is achieved by including a greater number of variables. From a modelling standpoint,

such an increase in variables is not desirable, as it may lead to overfitting and spurious accuracy. In contrast, the parsimonious structure of Model 6 avoids these pitfalls, making it a more robust representation of the system dynamics.

Overall, the network graph summary reinforces the quantitative spillover findings, suggesting that incorporating global market dynamics increases the measured interconnectedness among variables and, in doing so, enhances the endogenous explanatory power of the system.

2.5.3 Dynamic Spillover Analysis

To capture the evolution of interdependencies over time, we perform a rolling window analysis using two window sizes: two years and five years. Figure 2.2 displays the temporal behaviour of the TSI for each model, with the Global Economic Policy Uncertainty (GEPU) Index overlaid to illustrate the impact of global economic dynamics.

Our results indicate that models excluding global variables (Models 1–3) display relatively stable TSI values, whereas models incorporating global market factors (Models 4–6) exhibit both higher average TSI and greater sensitivity to major macroeconomic disruptions such as COVID-19 and the Russia–Ukraine conflict (Table 2.9). The two-year window captures short-term market reactions and transient volatility bursts, leading to more variable TSI trajectories. In contrast, the five-year window smooths these high-frequency fluctuations and shows a tendency for TSI values to converge toward a longer-run mean. This convergence implies that, over extended horizons, transient shocks dissipate and the overall degree of system interdependence stabilises around a persistent structural equilibrium. In other words, as short-lived idiosyncratic effects are averaged out, the rolling estimates increasingly reflect the enduring, equilibrium-level connectedness of the market.

Such behaviour is consistent with the literature on volatility persistence and long horizon connectedness (Diebold and Yilmaz, 2012; Antonakakis and Gabauer, 2017; Baruník and Křehlík, 2018), which shows that when spillover measures are computed over longer samples, they tend to stabilise around their structural mean, capturing systemic linkages rather than cyclical or event-driven fluctuations. From an economic perspective, this indicates that although short-term crises temporarily amplify cross-market transmission, the long-run network of interdependencies among oil market variables remains resilient and anchored to a stable equilibrium intensity.

Moreover, co-movements between the TSI and the GEPU Index suggest that periods of elevated global uncertainty coincide with higher spillover intensity, implying that uncertainty amplifies cross-market linkages. Notably, the WTI market (Model 2) appears more responsive to global events than the Brent market (Model 1), reflecting tighter integration of North America with global financial conditions.

FIGURE 2.2 : Total Spillover Index vs Global Economic Policy Uncertainty Index.

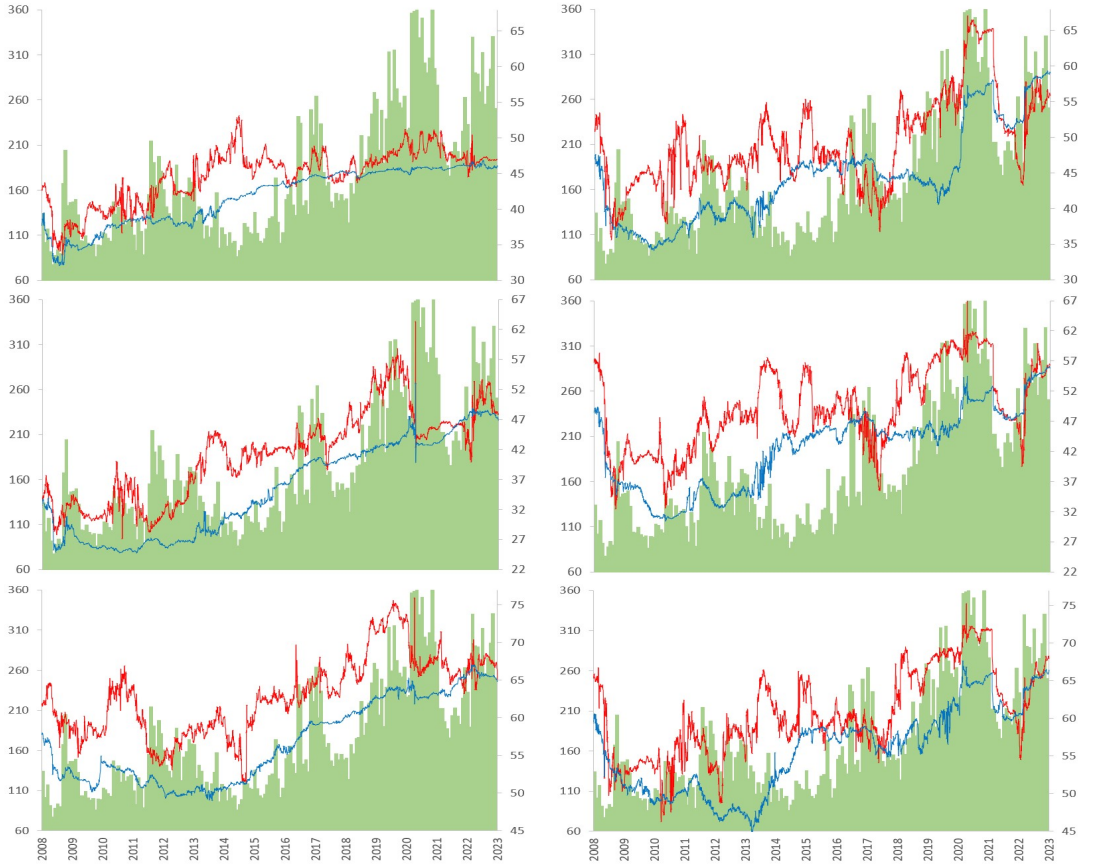

Note: Total Spillover Index vs GEPU. (i) Left Y-Axis - GEPU Index; Right Y-Axis - TSI percentage. (ii) Sample period: 12.01.2007 - 12.01.2023 [daily frequency]. (iii) Colour convention: Red line - TSI calculation based on moving window of 490 observations (equates to 2 years business cycle); Blue line - TSI calculation based on moving window of 1225 observations (equates to 5 years business cycle); Green area - GEPU index. (iv) List of graphs: a) Top row: (Left) Model 1 - BSP, BE, BL, BR; (Right) Model 4 - BSP, E, L, R; b) Middle row: (Left) Model 2 - WSP, WE, WL, WR; (Right) Model 5 - WSP, E, L, R; c) Bottom row: (Left) Model 3 - BSP, BE, BL, BR, WSP, WE, WL, WR; (Right) Model 6 - BSP, WSP, E, L, R.

Overall, the dynamic spillover analysis demonstrates that interconnectedness in the oil market has been persistently high and rising since the early 2010s, reflecting a structural shift towards stronger global co-movement. Models integrating both benchmarks (Model 3) or incorporating global variables (Models 4–6) capture both local and systemic interdependencies, confirming that variable selection grounded in spillover dynamics enhances explanatory and predictive power.

2.6 Concept Validation

The primary motivation of this study is to demonstrate that a higher Total Spillover Index is associated with improved predictability of the variables. To validate this hypothesis, we adopt a twofold approach: (i) assessing the spot price forecasting performance

TABLE 2.9 : TSI Summary statistics. Dynamic analysis.

Window Size	Model	Mean	Median	Max Val	Min Val	Std. D
490 obs [2 Year]	1	44.71	45.89	53.10	33.56	3.76
	2	41.06	42.20	63.31	27.16	7.36
	3	62.91	62.93	75.92	51.33	5.12
	4	49.25	47.99	67.00	35.61	6.28
	5	49.18	48.52	71.26	32.45	6.76
	6	60.69	59.83	75.19	46.28	5.71
1225 obs [5 Year]	1	41.66	43.25	46.83	32.09	3.95
	2	35.20	33.64	52.99	24.76	7.82
	3	56.75	55.56	67.05	49.00	5.49
	4	44.41	44.21	59.35	34.10	6.47
	5	42.54	44.58	56.11	30.47	6.66
	6	55.78	56.82	67.59	44.95	5.64

Note: Summary statistics for Total Spillover Index dynamic analysis. (i) Sample period between 07/12/2007 and 12.01.2023 of daily frequency. (ii) Window sizes: 2 years business cycle calculated based on 490 observations; 5 years business cycle calculated based on 1225 observations (iii) Model List: Model 1 - BSP, BE, BL, BR; Model 2 - WSP, WE, W L, WR; Model 3 - BSP, BE, BL, BR, WSP, WE, WL, WR; Model 4 - BSP, E, L, R; Model 5 - WSP, E, L, R; Model 6 - BSP, WSP, E, L, R;

of Models 1, 2 and 6 via the FCVAR approach, and (ii) comparing spot price forecast accuracy based on our best-performing model with the results based on variables from an analogous study (Guo et al., 2023b) using LSTM networks.

2.6.1 Out-of-Sample FCVAR Forecasting Performance

We apply the FCVAR model to three selected specifications, Model 1 (isolated Brent), Model 2 (isolated WTI), and Model 6 (hybrid global), to examine how variable selection based on high TSI values affects out-of-sample spot price forecasts. Figure 2.3 compares the forecasts of Brent (Model 1 vs. Model 6) and WTI (Model 2 vs. Model 6) with the actual spot prices. Model 6's forecast lines (red) maintain closer proximity to the observed series (blue) relative to Models 1 and 2, indicating that incorporating global market variables (E, L, R) enhances the predictive power of the system. In the bottom panels, Model 6 similarly outperforms Models 1 and 2 across the entire forecast horizon.

Table 2.10 quantifies forecast performance using the LinExp Loss and MAPE metrics. Negative RFE values between Model 6 and Models 1 or 2 confirm that Model 6 delivers superior forecast accuracy across multiple horizons. For instance, at a 60-day horizon, the Brent RFE values indicate Model 6 outperforms Model 1 by more than 20%, while the WTI RFE values show Model 6 surpasses Model 2 by over 15%. These improvements align with the hypothesis that systems exhibiting higher TSI, reflecting stronger dynamic interdependencies, yield more accurate forecasts.

FIGURE 2.3 : Brent and WTI Spot Out-of-Sample Forecasts - FCVAR Model 1, 2 & 6.

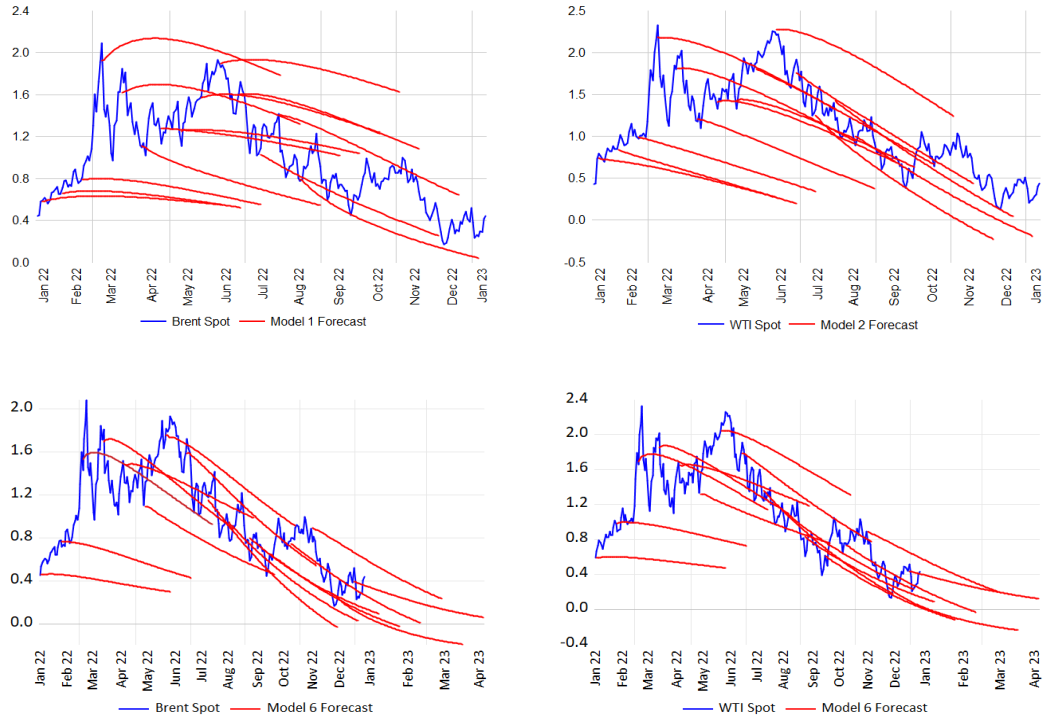

Note: Out-of-sample forecast fit of Brent and WTI spot price (red line) against the actual spot (blue line). Estimation model: FCVAR; Data frequency: daily; Estimation period: from 02-01-2003 to 12-01-2022; Out-of-Sample Forecasting Period: from 13-01-2022 to 12-01-2023. Spot prices source: Brent Spot - ICE; WTI Spot - NYMEX. Model 1 variables: BSP, BE, BL, BR; Model 2 variables: WSP, WE, WL, WR; Model 6 variables: BSP; WSP; E; L; R. The y-axis represents z-normalised representations of the prices; the x-axis represents dates.

Overall, the FCVAR forecasting exercise confirms the premise that integrating variables with strong interconnectedness emphasised by a high TSI enhances out-of-sample spot price forecasts. Model 6, which combines both benchmarks (Brent and WTI) alongside global market variables, consistently achieves lower forecast errors than Models 1 and 2, thereby substantiating the value of a global perspective in variable selection.

2.6.2 Forecast Comparison Tests and Model Dominance

While the previous subsection demonstrates that Model 6 achieves lower forecast losses than the isolated benchmarks (Models 1 and 2), formal statistical evaluation is required to determine whether these differences are significant and robust. We therefore apply three complementary forecast comparison procedures that assess: (i) equality of predictive accuracy, (ii) robustness of model inclusion in a superior set, and (iii) informational dominance. All tests are conducted on forecasts generated within the same FCVAR framework, ensuring that performance differences arise from the informational content of the variable sets rather than from differences in model class.

TABLE 2.10 : FCVAR Out-of-Sample Forecast Evaluation.

Loss Function	Model	Forecast Horison (days)							
		1	**5**	**10**	**20**	**40**	**60**	**80**	**100**
Brent Forecast Evaluation									
LinExp	**M6**	0.023	0.098	0.119	0.186	0.252	0.293	0.315	0.235
	M1	0.023	0.105	0.134	0.226	0.334	0.393	0.426	0.320
	RFE	-0.34	-6.86	-10.55	-17.76	-24.64	-25.49	-26.00	-26.48
MAPE	**M6**	0.0823	0.1716	0.2091	0.2613	0.3526	0.4483	0.4869	0.5026
	M1	0.0792	0.1686	0.2107	0.2752	0.3853	0.4744	0.5051	0.5647
	RFE	3.90	1.77	-0.76	-5.05	-8.47	-5.49	-3.60	-11.00
WTI Forecast Evaluation									
LinExp	**M6**	0.027	0.114	0.136	0.232	0.325	0.412	0.423	0.277
	M2	0.029	0.114	0.138	0.230	0.328	0.433	0.456	0.335
	RFE	-6.88	-0.21	-1.57	0.68	-0.86	-4.88	-7.15	-17.36
MAPE	**M6**	0.0798	0.1505	0.1893	0.2534	0.3381	0.4542	0.5020	0.4832
	M2	0.0828	0.1584	0.1994	0.2556	0.3468	0.4931	0.5629	0.6347
	RFE	-3.59	-5.01	-5.09	-0.89	-2.52	-7.89	-10.82	-23.87

Note: Following FCVAR methodology by Johansen and Nielsen (2012c). LinEx Loss: a=-1.5; RFE–Relative forecasting performance. Negative RFE value implies Model 6 superior forecast accuracy and vice versa.

TABLE 2.11 : Unified Forecast Evaluation Summary (BSP and WTI Spot Models)

Asset	Horizon	Model Pair	CPA p-val	MCS Survivors	p(f1 → f2)	p(f2 → f1)	Dominant Model
BSP	1	M1 vs M6	0.917	M1,M6	3.64e-02	1.33e-03	Tie
	5	M1 vs M6	0.572	M1,M6	2.22e-02	1.21e-02	Tie
	10	M1 vs M6	0.949	M1,M6	6.20e-03	3.27e-03	Tie
	20	M1 vs M6	0.688	M1,M6	1.10e-05	4.34e-06	Tie
	30	M1 vs M6	0.632	M1,M6	8.03e-08	2.88e-03	M6 (advantage)
	50	M1 vs M6	0.323	M1,M6	0.00000	9.81e-04	M6 (advantage)
	80	M1 vs M6	0.034	M1,M6	0.00000	0.00000	M6 (dominant)
	100	M1 vs M6	0.0006	M1,M6	0.00000	0.00000	M6 (dominant)
WTI	1	M2 vs M6	0.131	M2,M6	1.32e-03	8.59e-02	Tie
	5	M2 vs M6	0.677	M2,M6	3.90e-03	2.12e-02	Tie
	10	M2 vs M6	0.732	M2,M6	1.40e-03	1.55e-03	Tie
	20	M2 vs M6	0.096	M2,M6	1.09e-04	4.01e-03	M6 (advantage)
	30	M2 vs M6	0.282	M2,M6	3.19e-02	1.29e-02	M6 (advantage)
	50	M2 vs M6	0.161	M2,M6	3.06e-02	2.97e-02	M6 (advantage)
	80	M2 vs M6	0.0018	M2,M6	4.12e-14	6.91e-14	M6 (dominant)
	100	M2 vs M6	0.0004	M2,M6	4.44e-16	1.08e-12	M6 (dominant)

Note: CPA tests the null of equal predictive accuracy. MCS reports the models retained in the Hansen et al. (2011) superior set. Forecast encompassing tests are directional: $p(f_1 \to f_2)$ evaluates whether the isolated model encompasses Model 6, while $p(f_2 \to f_1)$ tests the reverse. The "Dominant Model" column summarises statistical evidence from these tests jointly with the relative forecast errors reported in Table 2.10. Labels indicate statistical parity (Tie), intermediate advantage for Model 6 (M6 (advantage)), or statistically robust superiority (M6 (dominant)). HAC–robust covariance (Newey–West, lag = $h-1$) is used throughout.

Conditional Predictive Ability (CPA). The Conditional Predictive Ability test of Giacomini and White (2006) evaluates whether the expected forecast loss of two competing models is statistically equivalent. Implemented with Newey–West HAC covariance (lag $h-1$), the CPA test nests the Diebold–Mariano statistic (Diebold and Mariano, 1995) as a special case when no conditioning variables are used. Results in Table 2.11 show no meaningful differences at short horizons ($h \leq 10$), consistent with the earlier finding that short-run volatility is largely dominated by high-frequency noise. As the horizon extends, the CPA p-values decline, and for horizons between 20 and 50 days, the evidence begins to favour Model 6. For horizons beyond 80 days the null of equal predictive accuracy is rejected at the 1% level for both Brent and WTI. The sign of the loss differential confirms that the rejection reflects superior accuracy of Model 6.

Model Confidence Set (MCS). To assess whether forecast differences remain robust to sampling variation, we employ the Model Confidence Set procedure of Hansen et al. (2011), using a stationary bootstrap (Politis and Romano, 1994) with 1,000 replications and confidence level $\alpha = 0.10$. At short horizons, both competing models remain within the superior set, indicating statistical parity. However, for horizons exceeding 50 days, the isolated models (M1 and M2) are iteratively eliminated, leaving only Model 6 in the confidence set. This behaviour indicates that the predictive advantage of the hybrid global model is systematic rather than driven by isolated episodes or sampling variability.

Forecast Encompassing. Finally, we assess informational dominance using the forecast encompassing regression of Chong and Hendry (1986):

$$y_t = \alpha + \beta f_{1,t} + \gamma f_{2,t} + u_t, \tag{2.7}$$

where $f_{1,t}$ and $f_{2,t}$ denote the forecasts from the competing models. HAC-robust tests of the restrictions $(\beta = 1, \gamma = 0)$ versus $(\beta = 0, \gamma = 1)$ indicate that at short horizons both models contribute marginally distinct information. Beyond 80 days, however, both null hypotheses are jointly rejected in favour of the specification where only Model 6 contributes explanatory content. This result confirms that the hybrid system encompasses the isolated specifications, meaning that its forecasts subsume the information embedded in Models 1 and 2.

The combined evidence from CPA, MCS, and forecast encompassing tests is summarised in Table 2.11. A consistent pattern emerges: (i) at short horizons all models display statistical parity; (ii) at medium horizons Model 6 begins to exhibit superior accuracy; and (iii) at horizons beyond 80 days, Model 6 shows clear statistical and informational dominance.

2.6.3 Robustness Check - Evaluation of Model 6 Forecast Performance Against Analogous Research

This section presents a comparative analysis of our best-performing Model 6 with the dataset from Guo et al. (2023b) to contrast the forecast accuracy of both variable sets. Spillover analysis, as reported in Appendix A.4, shows that Model 6 exhibits a higher TSI (63.14%) compared to 41.23% for the Guo et al. (2023b) dataset, indicating that Model 6 captures more robust interdependencies among the variables.

Out-of-sample forecasts are generated using LSTM networks trained on a chronological 90%–10% split, ensuring that the model is evaluated solely on unseen future observations. This design matches the structure used in Guo et al. (2023a), thereby supporting a fair comparison between Model 6 and the analogue benchmark. Forecast horizons extend up to 100 days. Table 2.12 summarises the forecast performance of both models.

TABLE 2.12 : LSTM Out-of-Sample Forecast Evaluation.

Loss Function	Model	Forecast Horizon (days)							
		1	**5**	**10**	**20**	**40**	**60**	**80**	**100**
Brent Forecast Evaluation									
LinEx Loss	**Model 6**	0.008	0.017	0.016	0.031	0.023	0.297	0.015	0.007
	Guo et al. (2023)	0.003	0.103	0.175	0.179	0.159	0.072	0.003	0.009
	RFE	134.03	-83.81	-91.02	-82.63	-85.76	315.41	352.94	-24.16
MAPE	**Model 6**	0.089	0.128	0.116	0.178	0.144	0.483	0.131	0.098
	Guo et al. (2023)	0.057	0.392	0.538	0.558	0.553	0.366	0.062	0.121
	RFE	56.20	-67.29	-78.51	-68.01	-74.03	32.16	113.51	-18.80
WTI Forecast Evaluation									
LinEx Loss	**Model 6**	0.006	0.005	0.006	0.013	0.010	0.009	0.004	0.007
	Guo et al. (2023)	0.005	0.007	0.067	0.096	0.084	0.013	0.008	0.009
	RFE	31.09	-30.89	-91.42	-85.98	-87.76	-30.54	-48.68	-25.13
MAPE	**Model 6**	0.071	0.059	0.064	0.105	0.086	0.079	0.058	0.077
	Guo et al. (2023)	0.062	0.070	0.298	0.378	0.363	0.122	0.106	0.123
	RFE	13.86	-16.09	-78.40	-72.12	-76.36	-34.95	-45.18	-37.47

Note: Following the LSTM methodology of Guo et al. (2023b). LinEx Loss asymmetry parameter: $\alpha = -1.5$. RFE denotes Relative Forecast Error; negative values indicate that Model 6 delivers lower loss (superior forecast accuracy) relative to the Guo et al. (2023) benchmark.

In the case of Brent forecasts, the Guo et al. (2023b) dataset demonstrates a short-term advantage at the 1-day horizon, with lower LinExp and MAPE values indicating that its variables may be more adept at capturing immediate market fluctuations. Positive Relative Forecast Error (RFE) values for this horizon confirm that the Guo et al. (2023b) model outperforms Model 6 in ultra-short-term predictions. However, as the forecast horizon extends beyond the first few days, Model 6 begins to exhibit significantly lower

LinExp and MAPE metrics, highlighting a clear distinction in predictive performance. A notable exception occurs at the 80-day horizon for Brent, where the Guo et al. (2023b) dataset briefly aligns more closely with actual price movements, as indicated by a lower MAPE value.

One particularly striking outlier occurs at the 60-day horizon for Brent forecasts using Model 6, where the predictions deviate substantially from the overall trend. Several LSTM-related factors may account for this anomaly, including the sensitivity of the model to the temporal structure of the data and the potential misalignment of historical patterns with sudden market events. Moreover, tuning of certain LSTM parameters, such as input sequence length or network architecture complexity, could also impact errors at these horizons.

Analysis of WTI forecasts further highlights the superiority of Model 6 across longer horizons, with consistently lower LinExp and MAPE values relative to the Guo et al. (2023b) dataset beyond the 20-day mark. Similar to the Brent results, the Guo et al. (2023b) dataset also shows competitive performance at the 1-day horizon, suggesting that certain variables in the setup of that study effectively capture very short-term price movements in WTI markets.

In general, this comparative analysis highlights the intricate role that variable selection and spillover-driven interdependencies play in forecasting oil prices. Although the Guo et al. (2023b) dataset exhibits an initial edge at the shortest horizons, Model 6, with its higher TSI, consistently outperforms at medium and long horizons. This finding aligns with the broader premise that a system displaying stronger dynamic interdependencies possesses a structural advantage in capturing complex, multi-period market behaviour. As a result, the broader interconnected market dynamics embedded in Model 6 appear to confer substantial benefits for longer-term forecasts, contributing valuable insights to ongoing discussions in financial modelling and forecasting.

2.6.4 Synthesis with Prior Research

Having established the spillover behaviour, dynamic patterns, and forecasting performance across all models, we now position these findings within the broader empirical and methodological literature. Studies based on VAR and VECM frameworks, such as Kilian (2009), Baumeister and Peersman (2013) and Zhang et al. (2019), report strong interactions between oil prices, macroeconomic activity and financial conditions. Within the Diebold–Yilmaz connectedness framework, Diebold and Yilmaz (2012), Antonakakis and Gabauer (2018) and Zhou and Duan (2021) find that oil markets are tightly linked to exchange rates, equity indices and uncertainty measures, with connectedness increasing around major crises. Our static and frequency domain TSI results confirm this general

picture of a densely interconnected system, but add a variable selection and forecastability dimension: models constructed around global expansion, liquidity and regime variables (Model 6) achieve comparable or higher connectedness with a more parsimonious set of predictors, and the accompanying FCVAR and LSTM exercises show that these higher TSI configurations translate into systematically better performance at medium and long forecast horizons.

The dynamic spillover analysis also fits naturally within the broader connectedness literature. Diebold and Yilmaz (2012) and Baruník and Křehlík (2018) emphasise that spillovers are time–varying and tend to spike during periods of financial stress, while Reboredo and Ugolini (2016), Antonakakis and Gabauer (2018) and Zhou and Duan (2021) document pronounced increases in oil–financial linkages around the global financial crisis, the European debt crisis and the COVID-19 shock. Our rolling window TSI estimates reproduce this pattern of crisis–driven surges and, in addition, show that models incorporating global factors are systematically more sensitive to episodes of heightened Global Economic Policy Uncertainty. At the same time, the tendency of TSI to stabilise when using longer windows is in line with evidence that long horizon connectedness reflects persistent structural linkages rather than short–lasting contagion effects (Baruník and Křehlík, 2018).

Finally, the forecast validation results complement the literature on oil price prediction using both econometric and machine learning approaches. Earlier work shows that models enhanced with economically meaningful drivers, such as demand indicators, financial conditions or volatility measures, tend to outperform purely statistical benchmarks (Malliaris and Malliaris, 2008; Johansen and Nielsen, 2012a; Salisu and Vo, 2020a; Guo et al., 2023a). Our FCVAR and LSTM experiments are consistent with this view, but refine it by demonstrating that it is not merely the inclusion of more information that matters, but the degree to which the selected variables are embedded in the spillover structure of the system ultimately determines their forecasting value. In particular, the superior performance of the high TSI hybrid model (Model 6) relative to both isolated benchmark models and the Guo et al. (2023a) variable set indicates that variables which occupy central positions in the connectedness network provide more robust forecasting signals. In this sense, the chapter extends existing studies by proposing a formal, spillover-based criterion for constructing the information set that provides a more consistent long-term forecast performance. These connections to the existing literature both validate the empirical patterns identified in this chapter and frame their broader relevance, motivating the policy and practical implications discussed next.

2.7 Policy and Practical Implications

The results of this study carry several implications for both forecasting practice and market oversight. The dynamic co-movement between the Total Spillover Index and global uncertainty measures indicates that models explicitly capturing interdependence provide earlier and more reliable signals of systemic stress. Rising spillovers have been shown to precede or accompany episodes of market disruption (Billio et al., 2012; Diebold and Yılmaz, 2014; Andries et al., 2019), and our findings confirm this pattern in oil markets, where peaks in TSI anticipate periods of heightened volatility and policy intervention. In energy markets, where prices transmit rapidly through physical trade, inventories, and financial contracts, monitoring spillover intensity offers policymakers and market participants an early indication that local disturbances are transitioning into broader systemic effects.

Second, the spillover-based framework provides a transparent and economically meaningful criterion for variable selection in forecasting and stress testing exercises. Rather than constructing the information set arbitrarily, forecasters can prioritise variables that contribute most to the internal coherence of the system. This principle extends naturally to macro-financial, climate-related, and high-dimensional forecasting contexts, where ensuring that predictive systems reflect genuine transmission channels is essential for credible scenario design.

Finally, the evidence that global market variables strengthen internal information circulation suggests that coordinated data integration across benchmarks, exchanges, and policy indicators improves not only forecast accuracy but also the capacity of models to generalise across regimes. From a governance perspective, this supports the use of interconnectedness metrics such as the TSI as early indicators of evolving market structure and as complementary tools for policy evaluation.

2.8 Conclusion

This study demonstrates that a spillover-based variable selection approach, quantified through the Total Spillover Index, can enhance both the empirical precision and the theoretical coherence of forecasting models. By analysing the dynamic interdependencies among variables drawn from isolated and global perspectives of the oil market, we show that systems exhibiting higher internal feedback, reflected in stronger spillover intensity, produce forecasts that are not only more accurate but also more robust across time horizons.

The empirical results derived from both the FCVAR and LSTM frameworks consistently demonstrate the superiority of the hybrid global specification (Model 6). The FCVAR

analysis shows that systems with higher TSI values capture long-run equilibrium relations more effectively, while the dynamic spillover patterns show that interdependencies intensify during episodes of macroeconomic and geopolitical stress. The co-movement between the TSI and the Global Economic Policy Uncertainty Index further illustrates that market information circulates through structured channels of feedback rather than diffusing randomly, reinforcing the internal coherence of the system.

The LSTM benchmark provides additional support for the framework. When compared to an analogous dataset from prior research (Guo et al., 2023b), Model 6 consistently yields lower forecast errors at medium and long horizons. This reinforces the central proposition of the study: forecastability is an emergent property of informational coherence, and models that reflect strong endogenous linkages inherit a greater ability to generalise across regimes.

Overall, the findings confirm that prioritising variables with strong dynamic interdependencies offers a robust, data-driven foundation for modelling complex economic systems. This approach integrates econometric structure with network intuition and demonstrates how the propagation of information governs predictive performance. While the analysis has focused on the oil market, the spillover-based framework extends naturally to other asset classes and to alternative machine learning architectures. Future research may explore these directions, further clarifying how system-level interdependence shapes predictive performance across diverse financial environments.

Chapter 3

Forward and Backward Memory in Fractionally Cointegrated Systems.

Abstract

Building a model within a memory–driven dynamic system framework, this paper studies the roles of historical persistence (backward memory) and anticipatory effects (forward memory) in the price formation process and their implications for forecast accuracy. Employing a novel bidirectional analysis of system memory in a dynamic setting, we compare two distinct market environments: the fundamentally driven Brent and the sentiment-driven Bitcoin markets. Empirical results show that Brent prices exhibit pronounced historical persistence but limited forward memory, reflecting robust historical trends yet rapid assimilation of new information, which restricts longer-term forecast accuracy. In contrast, Bitcoin demonstrates weaker backward memory but stronger forward memory, emphasising the significant role of anticipatory, speculative dynamics, thereby sustaining meaningful predictability for horizons of up to approximately 40 days. Robustness checks, including VAR Granger causality tests, further confirm that anticipatory signals are systematically incorporated into current price levels, thereby highlighting the importance of considering historical and anticipatory dimensions in financial forecasting and market modelling.

Keywords: Fractional Cointegration; Long Memory Processes; Anticipatory Effects; Forecastability; Price Formation; Brent Crude Oil; Bitcoin; Speculative Markets; Fundamentally Driven Markets; Investor Sentiment; Market Expectations; FCVAR Model.

JEL classification: C22; C32; G15; G17; G41; Q41

3.1 Introduction

Price formation in financial markets is traditionally modelled as a function of historical price movements and fundamental data. However, mounting evidence suggests that the price at time t_0 also reflects the anticipatory effects of investor expectations, how market participants foresee future developments and adjust their behaviour accordingly (Angeletos and Lian, 2018; King and Koutmos, 2021). In other words, the price emerges from the interplay between backward–looking dynamics and forward-looking anticipatory forces (Angeletos and Lian, 2018; King and Koutmos, 2021; Farmer et al., 2006).

Recent advances in time series econometrics, particularly the FCVAR framework (Johansen, 2008; Johansen and Nielsen, 2012b), enable a flexible characterisation of persistence within a system through the fractional integration parameter d, which may be non–integer and time varying (Jones et al., 2014). Although individual market variables often exhibit $I(1)$ behaviour, cointegrating relationships among them can produce a system with significantly lower memory, indicative of error correction and stabilising feedback loops (Granger and Joyeux, 1980; Hoskin, 1981).

This study is motivated by the idea that the price at t_0 reflects two intertwined forces: historical persistence (backward memory) and forward persistence (forward memory), indicative of anticipatory influences, that is, the extent to which expectations about future conditions are already embedded in current prices. Backward memory is estimated conventionally using observations from t_{-n} to t_0. Forward memory, by contrast, is estimated retrospectively using observations from t_0 to t_{+n}. It does not imply that future outcomes influence past prices; rather, it quantifies how strongly the subsequent trajectory aligns with what was already implicit in the system at t_0. In this sense, forward memory provides an ex-post measure of the anticipatory structure embedded in prices, consistent with theoretical accounts of forward-looking adjustment, dispersed information, and behavioural propagation mechanisms (Angeletos and Lian, 2018; King and Koutmos, 2021; Farmer et al., 2006).

Estimating fractional integration parameters d in both backward and forward directions using rolling and expanding windows therefore enables us to separate historical persistence from anticipatory dynamics. We apply this bidirectional framework to two contrasting environments: Brent crude oil, where price formation is dominated by slow-moving fundamentals (Morana, 2013), and Bitcoin, where speculative behaviour and investor sentiment play a central role (King and Koutmos, 2021; Gurdgiev and O'Loughlin, 2020).

Our findings show a contrasting pattern between *backward* and *forward* d. In the oil market, *backward* d is typically high (0.8-1), indicating that historical price dynamics are highly persistent. However, during periods of stress or structural adjustment, *backward* d occasionally drops into lower values, suggesting strong mean reversion or overreaction

effects (Morana, 2013). Simultaneously, the *forward d* remains generally low (often below 0.5) with occasional spikes, implying that anticipatory adjustments are quickly assimilated.

In contrast, Bitcoin exhibits consistently low *backward d* and high, volatile *forward d*. Although conventional wisdom holds that a d value below 0.5 indicates a covariance stationary and thus forecastable process (Granger and Joyeux, 1980; Hoskin, 1981), our results suggest that this may not hold when strong anticipatory effects are present. For instance, despite low *backward d* in the Bitcoin market, its high and volatile *forward d* implies that anticipatory dynamics can drive explosive price behaviour (Turatti et al., 2020; Blanc et al., 2017). While Turatti et al. (2020) previously highlighted mean–averting behaviour in Bitcoin returns, they did not explicitly examine *forward d* dynamics, an aspect directly explored in our empirical analysis. Moreover, studies by King and Koutmos (2021) and Gurdgiev and O'Loughlin (2020) support the notion of speculative dynamics in Bitcoin, highlighting the need to consider both backward and forward measures to fully contextualise its price evolution.

The core research question addressed in this study is: what does the relationship between backward and forward system memory imply for the price formation process, how do these dynamics differ between fundamentally and sentiment-driven markets, and what implications does this have for their forecastability? Our approach advances the methodology on fractional cointegration and offers insights into asymmetries in information flow that have a direct impact on forecasting and risk management (Lo, 2004; Shiller, 2003).

The remainder of this paper is organised as follows. Section 3.2 reviews the literature on fractional integration and anticipatory effects; Section 3.3 outlines our theoretical framework and empirical strategy; Section 3.4 describes the data and pre–processing checks; Section 3.5 presents the empirical results; Section 3.6 discusses implications of bidirectional memory dynamics; and Section 3.7 concludes with the summary of the results and directions for future research.

3.2 Literature Review

The investigation of anticipatory effects and long memory in financial markets has attracted considerable attention in both theoretical and empirical studies. This review synthesises the relevant literature, focusing on three core themes: long memory in financial time series, anticipatory effects in price formation, and the contrasting dynamics between fundamentally driven and speculative markets.

Long memory processes, characterised by persistent autocorrelation structures, have been widely explored in financial econometrics. Early works by Granger and Joyeux (1980) and Hoskin (1981) introduced fractional differencing to model such persistence,

capturing the notion that shocks may decay slowly relative to the rapid dissipation observed in short memory processes.

In commodity markets, for instance, Alvarez-Ramirez et al. (2002) and Morana (2013) document multifractal properties and long memory in crude oil prices, while minimal market impact models endogenising persistent order flow and price response have been developed by Donier et al. (2015) to explain non–linear impact dynamics. Conventional wisdom holds that when the fractional integration parameter d is below 0.5, the process is covariance stationary and thus forecastable. However, our research challenges this view by incorporating the role of anticipatory effects.

Anticipatory effects refer to the influence of market expectations on current asset prices. Angeletos and Lian (2018) and King and Koutmos (2021), supported by empirical analyses of meta–order impact (Gomes and Waelbroeck, 2015), argue that investor expectations and herding behaviour can generate self–reinforcing feedback loops, leading to price trajectories that deviate from fundamental values. In speculative markets, such effects may cause prices to diverge from the forecastability implied by covariance stationarity. For example, Turatti et al. (2020) provide evidence of mean–averting behaviour in Bitcoin returns, suggesting that anticipatory forces may drive prices away from historical norms. This highlights the need to examine both backward–looking (historical) and forward–looking (anticipatory) components of market dynamics.

The Brent crude oil and Bitcoin markets represent two distinct financial environments. Previous research has shown that oil prices are predominantly influenced by fundamental factors like supply–demand imbalances, geopolitical events, and macroeconomic conditions (Morana, 2013). Conversely, Bitcoin price movements are mainly driven by speculative sentiment, investor behaviour, and herding dynamics (King and Koutmos, 2021; Gurdgiev and O'Loughlin, 2020). While several studies have independently explored long memory and forecastability in these markets (Degiannakis and Filis, 2018; Deng et al., 2011), direct comparative analysis of anticipatory effects remains scarce, highlighting a gap addressed by our research.

The fundamentally anchored nature of Brent aligns with structural models of the crude oil market, where supply stiffness, inventory behaviour, storage costs, transportation constraints and geopolitical conditions jointly generate persistent price dynamics (Alvarez-Ramirez et al., 2002; Morana, 2013). These mechanisms also explain the long memory characteristics documented in crude oil returns, where shocks propagate slowly through the physical market. Consequently, the literature suggests that markets such as Brent, which are anchored in slow-moving fundamental forces, may be expected to display stronger backward persistence, whereas forward-looking adjustments are likely to be absorbed more rapidly. Whether this pattern holds empirically is examined in the subsequent sections of this chapter.

The FCVAR model provides a robust framework for analysing systems with long memory and cointegrated relationships (Johansen, 2008; Johansen and Nielsen, 2012b). By allowing fractional integration orders that capture both short and long–term dependencies, the FCVAR model is particularly suited for investigating the interplay between historical persistence and anticipatory effects. Although recent applications have demonstrated its versatility in various contexts (Jones et al., 2014), no studies have explored the dual estimation of backward and *forward* d values.

Despite the extensive literature on long memory and anticipatory effects, research comparing these phenomena across fundamentally driven and speculative markets remains scarce. Empirical evidence from oil markets indicates strong historical persistence (Morana, 2013), yet the impact of anticipatory adjustments remains underexplored. Similarly, while Turatti et al. (2020) highlight the explosive dynamics in Bitcoin returns, their approach is complex and rests on key assumptions, such as the stationarity of Bitcoin returns and the adequacy of Gibbs-sampling-augmented randomisation, to capture evolving market dynamics. Although their methodology accounts for heteroscedastic disturbances, it does not fully address structural breaks, regime shifts, or the influence of exogenous factors like regulatory changes and macroeconomic shocks. As such, a broader approach that fully captures and contextualises the dynamic interplay between historical persistence and anticipatory effects remains necessary.

To address this gap, our study employs a rolling and expanding window FCVAR approach to estimate d in both backward (t_{-n} to t_0) and forward (t_0 to t_{+n}) directions. This bidirectional estimation method is designed to capture the intricacies of price evolution in a dynamic and comprehensive manner, enabling us to determine whether the explosive dynamics observed in Bitcoin are primarily driven by strong anticipatory adjustments rather than solely by historical persistence. Moreover, this methodology allows for a direct comparison between the volatile, sentiment–driven dynamics of Bitcoin and the more stable, fundamentally driven dynamics observed in markets such as Brent crude oil. Our comparative analysis thus enables us to test the following hypothesis:

Hypothesis: In the Brent crude oil market, *backward* d will be significantly higher than *forward* d, reflecting the dominance of long-run fundamentals and the rapid assimilation of anticipatory adjustments. In contrast, in the Bitcoin market, *backward* d will be significantly lower while *forward* d will be significantly higher, indicating that speculative forces and anticipatory effects, driven by self–reinforcing behavioural biases, drive explosive price dynamics.

By testing this hypothesis, our research aims to contribute to the understanding of how anticipatory effects modulate the forecastability implied by conventional long memory estimates, thereby offering new insights into market efficiency and the dynamics of price formation.

3.3 Methodology

This section outlines the methodological approach and theoretical basis of our analysis, focusing on the estimation of the fractional integration parameter d using the FCVAR framework. For complete technical derivations and detailed model equations, readers are referred to Appendix B.

3.3.1 FCVAR Model and Bidirectional Memory Estimation

Our research employs the FCVAR model as the principal framework for capturing the interplay between historical persistence and anticipatory effects in price formation. Developed by Johansen (2008); Johansen and Nielsen (2012b), the FCVAR model extends conventional cointegration techniques by allowing the integration order d to assume non–integer and time–varying values. This flexibility is essential for accurately representing long memory properties, as the parameter d quantifies the persistence of shocks in a system. Although traditional theory stipulates that a d value below 0.5 indicates a covariance stationary and forecastable process, our analysis suggests that this may be insufficient when anticipatory effects are significant.

Within this framework, we distinguish between two measures of memory. The *backward* d is estimated from historical data, capturing the enduring influence of past shocks. In contrast, the *forward* d is derived from future observations, reflecting the anticipatory effects of investor expectations. This dual estimation allows us to examine whether fundamentally driven markets, such as Brent crude oil, exhibit higher *backward* d relative to *forward* d, while sentiment–driven markets like Bitcoin display the opposite pattern. It is important to note that although the *forward* d measure can only be estimated retrospectively, since it relies on future observations, it nonetheless provides critical insights into how anticipatory market behaviour shapes current price levels.

3.3.2 Empirical Strategy

Our empirical strategy estimates d in both backward and forward directions using rolling and expanding windows. Each estimation begins with a window of 510 observations, approximately two years of trading data. This length provides a balance between statistical stability and sensitivity to structural changes as shorter windows are known to yield noisy persistence estimates, while excessively long windows risk smoothing over the dynamics our framework aims to uncover.

The final choice of 510 observations is also supported by an empirical optimisation procedure described in Appendix B.2. Using a VECM, nested within the FCVAR framework when $d = 1$, we evaluated short-horizon forecast performance across a range of window

lengths. The optimisation consistently selected windows close to two years, which aligns with typical business-cycle frequencies in crude oil markets and the adjustment horizon of supply–demand fundamentals.

In both rolling and expanding specifications, t_0 denotes the final observation in the current window, dividing the data into backward and forward segments. The rolling window keeps both segments fixed at 510 observations, while the expanding window increases the backward segment over time as the forward segment contracts. Estimation terminates once the forward window reaches 510 observations, ensuring that each evaluation point is based on a full and symmetrically defined forward segment.

To ensure robustness, the empirical design incorporates PCA-based dimension reduction, spurious memory diagnostics, and systematic checks of cointegration ranks and lag structures. Additional analyses, including VAR Granger causality tests (Appendix B.3), further validate the interaction mechanisms underlying price formation in each market.

3.4 Data and Preprocessing checks

Our analysis begins with an extensive dataset of 299 time series collected from various market segments, encompassing commodity and cryptocurrency markets as well as key market indexes. In commodity markets, data are obtained from both the spot and futures segments. To manage the inherent complexity of futures data, we express them in terms of their term structure characteristics. Following the framework proposed by Mustanen et al. (2022), for each of the markets considered, we derive two key measures from the price term structure: the slope, which captures regime shifts in the market, and the area under the curve, which represents market expansion. In addition, the area under the curve of the open interest term structure is used to capture the liquidity component. This dimension reduction approach ultimately condenses our variable set to 23 individual series we use in our analysis (see Table 3.1). For oil market analysis, we focus on variables derived from the term structure of both the Brent and WTI oil benchmarks, including BE, BL, BR, WE, WL, and WR, as detailed in Section (a) of Table 3.1. Throughout this chapter, references to the Brent and WTI 'spot' prices correspond to their respective front-month futures contracts, which are the standard high-frequency benchmarks used in empirical energy-market analysis.

At each iteration, Principal Component Analysis (PCA) is applied to the futures-based variables from both oil markets, and the first three principal components are extracted. These components, which jointly capture the dominant equilibrium dynamics of market expansion, liquidity, and regime change, are then combined with the respective spot price series to form the input for the FCVAR estimation.

TABLE 3.1 : Summary of Analysis Variables.

Segment	Variable	Description	Derivation Method	Ticker/Source
		(a) Energy Market Variables		
Brent Oil	BSP	Spot (front-end future)	-	CO1 Comdty
	BE	Market expansion	Area Under Futures	CO1 → CO36 Comdty
	BL	Market liquidity	Area Under Open Int	CO1 → CO36 Comdty
	BR	Market regime	Slope of Futures	CO1 → CO36 Comdty
WTI Oil	WSP	Spot (front-end future)	-	CL1 Comdty
	WE	Market expansion	Area Under Futures	CL1 → CL36 Comdty
	WL	Market liquidity	Area Under Open Int	CL1 → CL36 Comdty
	WR	Market regime	Slope of Futures	CL1 → CL36 Comdty
Natural Gas	NGE	Market expansion	Area Under Futures	NG1 → NG12 Comdty
	NGL	Market liquidity	Area Under Open Int	NG1 → NG12 Comdty
	NGR	Market regime	Slope of Futures	NG1 → NG12 Comdty
		(b) Hedging Indicators		
Gold	GCSP	Spot	-	GOLDS Comdty
	GCE	Market expansion	Area Under Futures	GC1 → GC12 COMB Comdty
	GCL	Market liquidity	Area Under Open Int	GC1 → GC12 COMB Comdty
	GCR	Market regime	Slope of Futures	GC1 → GC12 COMB Comdty
		(c) Cryptocurrency Market Indicators		
Currency	BTC	Bitcoin spot	-	finance.yahoo.com
Currency	ETH	Ethereum spot	-	finance.yahoo.com
Currency	LTC	Litecoin spot	-	finance.yahoo.com
Index	CCI30	Crypto Curr Index 30	-	cci30.com
		(d) Risk Indicators		
Index	TEUSCA	Twitter-based Uncert.	-	policyuncertainty.com
Index	USEPU	US Econ Policy Uncert.	-	policyuncertainty.com
Index	VIX	Expected Vol Index	-	VIX Index
Index	GILI	Global Infl Linked Index	-	LF94TRUU

Note: Table presents the list of variables selected for the analysis of oil and cryptocurrency market dynamics. It is organised into several sections, each representing a specific market segment or type of variable. The 'Derivation Method' column describes the dimension reduction techniques applied to obtain particular variable, as proposed by Mustanen et al. (2022). For example, the variable 'BE' signifies the Brent market expansion and is calculated as the area under the futures curve, spanning the 1 to 36 month maturity range, with specific tickers referenced in the rightmost column.

In contrast, for the Bitcoin market, we use all the variables from the reduced set as shown in Table 3.1. Here, PCA is applied to the entire reduced dataset, and again the first three principal components are extracted. These components are then combined with the Bitcoin spot price series at each iteration to be used in the FCVAR estimation. To ensure consistency and comparability across different variables, we apply Z-Score normalisation, scaling each series to a mean of zero and a variance of one, thus mitigating issues arising from differences in scale.

This dual approach ensures that our models are both parsimonious and generalised, effectively capturing the critical market forces we target in this study, fundamental

TABLE 3.2 : Summary Statistic of Analysis Variables

Variable	Median	Maximum	Minimum	Skewness	Kurtosis	Jarque-Bera	Pval
BSP	-0.0469	3.3525	-2.6808	0.5910	3.5477	126.3	0.00
BE	-0.1774	2.8866	-2.0337	0.5457	2.8171	91.1	0.00
BL	0.3952	1.5718	-1.9777	-0.8836	2.2094	278.9	0.00
BR	-0.0499	2.2454	-3.6034	-0.3535	3.2112	40.5	0.00
WSP	-0.1733	3.3218	-5.0826	0.6459	3.9854	196.4	0.00
WE	-0.2263	3.1602	-2.2792	0.7014	3.2823	152.4	0.00
WL	0.0631	2.2080	-2.3658	-0.2615	2.7394	25.4	0.00
WR	-0.0111	2.8309	-3.6114	-0.4844	3.3214	77.5	0.00
NGE	-0.3173	4.1333	-1.0855	2.0667	6.9337	2423	0.00
NGL	0.0296	2.6302	-2.3022	-0.1509	2.5524	21.7	0.00
NGR	0.2307	1.8138	-4.8740	-1.9935	8.0648	3092	0.00
GCSP	-0.2283	2.0509	-1.6500	0.2414	1.4393	198.6	0.00
GCE	-0.3002	1.8821	-1.5372	0.2681	1.4187	207.5	0.00
GCL	-0.1616	3.9569	-2.0013	0.8776	4.5646	411.4	0.00
GCR	-0.0405	3.0406	-2.4032	0.1681	2.9507	8.6	0.01
BTC	-0.3944	3.1948	-0.9113	1.3368	3.7108	569.5	0.00
ETH	-0.5143	3.5589	-0.7816	1.5466	4.4824	875.5	0.00
LTC	-0.2959	4.5506	-1.1484	1.2537	4.6867	679.6	0.00
CCI30	-0.3880	3.7617	-0.9698	1.4357	4.2909	737.5	0.00
TEUSCA	-0.1991	27.386	-1.0093	12.730	320.79	7563540	0.00
USEPU	-0.2838	7.0149	-1.2343	2.4311	10.809	6297	0.00
VIX	-0.2178	7.8309	-1.2229	2.3155	13.344	9558	0.00
GILI	-0.3029	2.3089	-2.1636	0.6392	2.3296	155.1	0.00

Note: This table summarises the statistical attributes of z-normalised variables (Eq: 2.1) used in the analysis of oil and cryptocurrency markets. Measures of central tendency and variability such as the median, maximum, and minimum are reported alongside skewness and kurtosis to describe the distribution characteristics. The Jarque-Bera test is used to assess the normality of the distributions with corresponding p-values (Pval) indicating the significance levels. A p-value of 0.00 suggests that the null hypothesis of normality is rejected for the given variable at conventional significance levels. Period covered: Start Date - 04-01-2016; End Date - 16-03-2023; Number of Observations - 1786.

dynamics in oil markets and speculative sentiment in the Bitcoin market. Summary statistics for the transformed series are presented in Table 3.2.

3.4.1 Pre-Estimation Memory Validation

To verify that the persistence observed in our series reflects genuine long memory rather than artefacts of structural breaks or deterministic trends, we employ the Shimotsu test (Shimotsu and Phillips, 2006) alongside PP, KPSS and modified Wald diagnostics. As shown in Table 3.3, the estimated fractional parameters remain broadly stable across subsample splits, and the Wald statistics fall well below critical values for most variables. The only notable deviation occurs for the Brent series in the four-way split, coinciding with the COVID–19 and Russia–Ukraine shocks; such episodic breaks are well known to trigger spurious rejections in subsample-based tests without undermining the presence of underlying long memory.

TABLE 3.3 : Summary Statistic of Analysis Variables

Market	Variable	$\hat{d}$							Average $\hat{d}$			PP crit			KPSS crit			Test Stat		Wald Stat	
		1	2		4				1	2	4	1	2	4	1	2	4	PP	KPSS	2	4
Bitcoin	BTC	1.16	0.98	1.24	1.02	0.94	1.51	1.17	1.16	1.11	1.16	-2.56	-2.85	-3.43	0.35	0.46	0.73	-2.98	0.04	2.70	6.47
	EGY	1.06	0.93	1.13	0.68	1.15	1.04	1.22	1.06	1.03	1.02	-2.56	-2.85	-3.42	0.35	0.46	0.74	-1.76	0.09	1.75	5.84
	INFL	1.05	0.95	1.12	0.67	1.22	1.17	1.12	1.05	1.03	1.04	-2.56	-2.85	-3.42	0.35	0.46	0.74	-2.64	0.10	1.20	6.70
	LQT	0.92	0.88	0.97	0.68	1.26	0.96	1.00	0.92	0.92	0.98	-2.56	-2.85	-3.43	0.35	0.46	0.74	-2.63	0.14	0.36	5.80
Oil	BSP	1.07	1.01	1.16	0.71	1.53	1.20	1.23	1.07	1.09	1.17	-2.56	-2.85	-3.42	0.35	0.46	0.74	-2.52	0.06	0.86	11.94
	E	1.02	0.97	1.09	0.71	1.34	1.13	1.17	1.02	1.03	1.09	-2.56	-2.85	-3.42	0.35	0.46	0.74	-1.91	0.06	0.61	7.24
	L	0.78	0.71	0.90	0.83	0.63	1.26	0.97	0.78	0.80	0.92	-2.56	-2.85	-3.43	0.35	0.46	0.74	-2.81	0.06	1.40	7.24
	R	1.03	1.07	1.09	0.89	1.15	1.01	1.31	1.03	1.08	1.09	-2.56	-2.85	-3.42	0.35	0.46	0.74	-2.42	0.21	0.01	3.39

Note: The table presents spurious memory analysis for Brent and Bitcoin market associated series using the Shimotsu test. The test examines if the long memory properties are genuine or spurious by comparing the estimated values of d across different frequency bands. The PP and KPSS tests check for the presence of unit roots and stationarity, respectively. The Wald statistic evaluates the stability of the memory parameter across different sample periods against critical value of 7.815. Series E, L, and R are global market representations for market expansion, liquidity and regime extracted from term structure related to Brent and WTI futures markets per methodology by Mustanen et al. (2022). Series ENG (energy), INFL (inflation) and LQT (liquidity) represent PC1, PC2 and PC2 derived using PCA from the entire dataset considered in this study (Table 3.1). Period covered: Start Date - 04-01-2016; End Date - 16-03-2023; Total number of Observations - 1786.

We complement this analysis with DF–GLS (Elliott et al., 1996) and KPSS (Kwiatkowski et al., 1992) unit root tests, which carry opposite null hypotheses. For the majority of variables reported in Table 3.4, DF–GLS fails to reject the unit-root null in levels but rejects it in first differences, while KPSS rejects stationarity in levels but not after differencing. This joint pattern is consistent with $I(1)$ dynamics or fractional integration with a non-integer d behaviour (including cases where d exceeds one), and helps rule out spurious stationarity or trend-driven artefacts. One exception is the liquidity variable L in the oil market section, which appears stationary under DF–GLS but not under KPSS. This mixed outcome is common for high-frequency microstructure indicators and does not affect the subsequent FCVAR analysis, as the model accommodates combinations of $I(0)$ and $I(1)$ variables and relies on system-wide equilibrium relations rather than the integration order of any single series.

Jointly, the Shimotsu results, subsample stability, and complementary unit-root evidence support the conclusion that the persistence observed in our data reflects genuine long-range dependence, validating the use of FCVAR for the bidirectional memory analysis that follows.

TABLE 3.4 : DF–GLS and KPSS Unit Root Tests.

Market	Variable	Transform	DF–GLS (H_0: unit root)				KPSS (H_0: stationarity)			
			Stat	Crit [5%]	Reject	Lag	Stat	Crit [5%]	Reject	Lag
Oil	BSP	Level	-0.802	-1.956	N	25	2.552	0.461	Y	27
		First Diff	-3.733	-1.956	Y	24	0.072	0.461	N	8
	E	Level	-0.604	-1.956	N	25	2.883	0.461	Y	27
		First Diff	-3.345	-1.956	Y	24	0.071	0.461	N	9
	L	Level	-4.732	-1.956	Y	23	0.586	0.461	Y	26
		First Diff	-3.940	-1.956	Y	22	0.102	0.461	N	10
	R	Level	-1.015	-1.956	N	24	0.850	0.461	Y	27
		First Diff	-3.275	-1.956	Y	23	0.226	0.461	N	5
Bitcoin	BTC	Level	-0.790	-1.956	N	14	3.907	0.461	Y	27
		First Diff	-10.773	-1.956	Y	13	0.080	0.461	N	6
	EGY	Level	-0.337	-1.956	N	20	4.227	0.461	Y	27
		First Diff	-8.638	-1.956	Y	19	0.105	0.461	N	12
	INFL	Level	-0.833	-1.956	N	3	1.650	0.461	Y	27
		First Diff	-13.354	-1.956	Y	2	0.126	0.461	N	13
	LQT	Level	-1.750	-1.956	N	14	9.208	0.461	Y	27
		First Diff	-1.884	-1.956	N	13	0.140	0.461	N	13

Note: This table reports Elliott–Rothenberg–Stock DF–GLS and KPSS unit root tests for the variables entering the Brent and Bitcoin systems. DF–GLS is estimated under a constant-only specification with automatic lag selection and a unit-root null. KPSS is run with a constant and the null of level stationarity. 'Reject' indicates rejection of the respective null at the 5% level. Evidence of $I(1)$ behaviour is interpreted as a DF–GLS non-rejection in levels paired with rejection in first differences, and a KPSS rejection in levels but not in differences. Reported lags correspond to the automatic lag (DF–GLS) or bandwidth (KPSS) selected endogenously by each procedure.

3.5 Empirical Results

The empirical analysis tests the hypothesis that fundamentally driven markets, such as Brent crude oil, exhibit higher backward memory relative to forward memory, while speculative markets, such as Bitcoin, display the opposite structure. Fractional integration parameters d are estimated using both expanding and rolling window frameworks, as illustrated in Figure 3.1 with summary statistics reported in Table 3.5. Before presenting the estimates, it is useful to recall that, from an economic standpoint, forward memory provides an ex-post view of anticipatory behaviour, showing how much of the later price movement was already implicit in conditions at time t_0.

Brent Crude Oil

Under the expanding window framework (top-left panel of Figure 3.1), the backward memory estimates for Brent remain consistently high, generally between 0.8 and 1, indicating strong persistence in the historical component of price formation. This stability aligns with the slow-moving fundamental forces shaping the Brent market, including

production cycles, inventory adjustments, and OPEC behaviour. In contrast, the forward memory series frequently drops to much lower values and exhibits several abrupt drops, including episodes where the estimated value approaches zero. These collapses occur precisely during periods in which the forward window contains major structural breaks, such as the 2018 oil price correction and the 2020 COVID-19 demand shock. In such cases, the future trajectory is dominated by developments not yet forming at time t_0, resulting in minimal anticipatory structure and hence low forward memory.

FIGURE 3.1 : Forward and backward system long-memory profiles of Brent and Bitcoin markets - obtained via FCVAR framework using expanding and rolling window estimation.

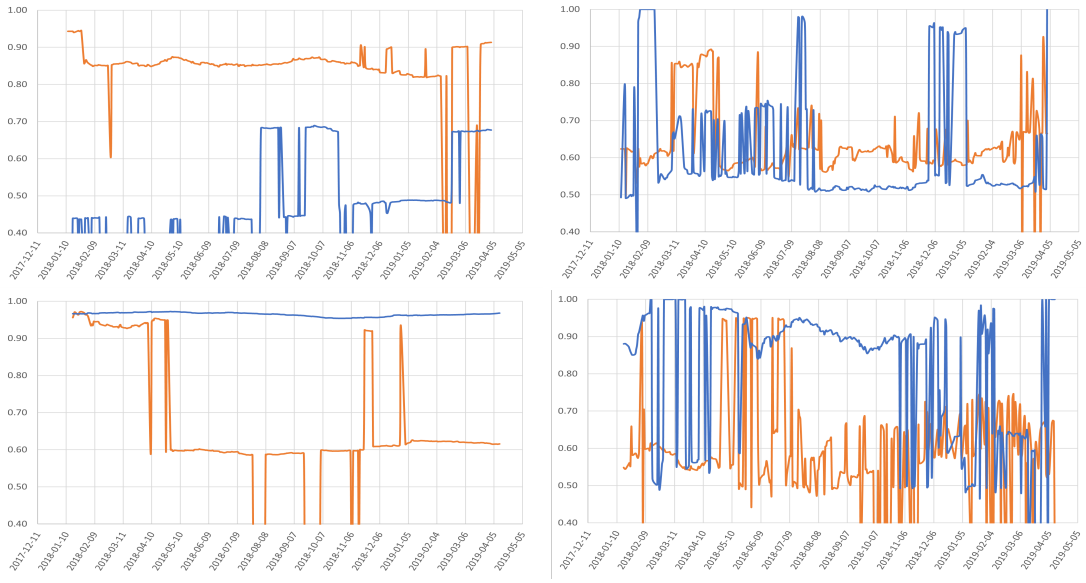

Note: The figure reports backward (orange) and forward (blue) fractional memory estimates for Brent (top panels) and Bitcoin (bottom panels). Backward memory is estimated using observations from t_{-n} to t_0, while forward memory is estimated retrospectively using observations from t_0 to t_{+n} and provides an ex-post measure of anticipatory structure at time t_0. The left panels correspond to expanding-window estimation and the right panels to rolling windows. Abrupt spikes or drops, particularly in forward-memory estimates, arise when the estimation window spans major structural breaks or rapid regime shifts.

The rolling window results for Brent (top-right panel) display a different dynamic. Here, backward memory declines, while forward memory becomes substantially more volatile. Short-lived spikes in forward memory correspond to brief intervals in which anticipatory factors, such as expectations of supply disruptions, geopolitical tensions, or macroeconomic announcements, temporarily exert a disproportionate influence on price dynamics. The heightened volatility of the forward estimates reflects the greater sensitivity of rolling windows to recent information and short-run market conditions. This contrast between expanding and rolling windows highlights that anticipatory component of Brent is episodic rather than persistent, whereas its historical component remains stable over longer horizons.

Bitcoin

The Bitcoin results (bottom panels of Figure 3.1) show a markedly different memory structure. In the expanding window analysis, forward memory remains persistently high, often close to one, indicating that speculative expectations and momentum-based trading patterns dominate the forward-looking behaviour of the system. Backward memory, however, exhibits pronounced instability, with sharp declines occurring frequently throughout the sample. These declines reflect susceptibility of Bitcoin to regime shifts, sentiment reversals, and liquidity-driven cycles, all of which disrupt its historical structure. In economic terms, price formation in Bitcoin market increasingly reflects expectation-driven behaviour rather than persistence grounded in past dynamics.

TABLE 3.5 : Summary Statistic of Memory Profiles

System	Memeory	Mean	Median	S.D.	Skew	Kurt	Jarque-Bera	Pval
Brent	EW d_b	0.837	0.855	0.113	- 4.183	19.998	4,561	0.000
	EW d_f	0.460	0.444	0.141	0.409	1.936	22.88	0.000
	RW d_b	0.635	0.616	0.094	1.150	5.902	174.3	0.000
	RW d_f	0.610	0.539	0.150	1.565	4.205	143.0	0.000
Bitcoin	EW d_b	0.662	0.610	0.178	0.294	2.787	4.98	0.083
	EW d_f	0.965	0.966	0.005	- 0.626	2.393	24.60	0.000
	RW d_b	0.595	0.577	0.142	0.537	4.323	36.9	0.000
	RW d_f	0.818	0.894	0.178	- 1.077	2.926	59.0	0.000

Note: This table summarises the statistical attributes memory profiles derived from z-normalised variables used in the analysis of oil and cryptocurrency markets. Measures of central tendency and variability such as the mean and median are reported alongside skewness and kurtosis. The Jarque-Bera test is used to assess the normality of the distributions with corresponding p-values (Pval) indicating the significance levels. Notation: EW - expanding window; RW - rolling window; d_b - historical observations based memory; d_f - future observations based memory.

The rolling window analysis amplifies these features. Both backward and forward memory display extreme volatility, with the forward component oscillating rapidly between high and low persistence. This instability reflects the absence of a stable long-run cointegrating structure in the Bitcoin system. When rolling windows include abrupt speculative surges or rapid unwinding episodes, the FCVAR estimator temporarily loses the underlying equilibrium relationships, leading to rapid swings in the estimated memory parameters. Such behaviour is characteristic of markets driven primarily by sentiment and speculative dynamics rather than by fundamental anchors.

In summary, the results highlight a clear structural contrast between the two markets. Brent exhibits strong and stable backward memory, consistent with long-run dynamics anchored in slow-moving fundamentals, while its forward memory is episodic and sensitive to isolated anticipatory events. Bitcoin displays the opposite configuration. Here, forward memory is dominant but highly unstable, reflecting the speculative and expectation-driven nature of its price formation, whereas backward memory is weaker and frequently disrupted by regime shifts. These patterns reinforce the importance of

modelling both historical persistence and the anticipatory structure of the system, particularly in markets characterised by speculative dynamics or frequent structural breaks.

3.5.1 Relation of Memory Dynamics to Prior Research

The memory profiles documented above are broadly consistent with, but also extend, existing evidence on crude oil and cryptocurrency markets. For Brent, the dominance of backward over forward memory aligns with studies that emphasise the role of slow-moving fundamentals in driving oil price dynamics, including supply constraints, inventory adjustment, and geopolitical risk (Alvarez-Ramirez et al., 2002; Morana, 2013). The high and relatively stable backward d estimates are in line with long memory and multifractal features reported for crude oil, where shocks propagate gradually through the physical and storage networks. At the same time, the episodic surges and collapses in forward memory around events such as the 2018 price correction and the COVID-19 demand shock echo the view that sharp revisions in expectations are absorbed quickly once new information becomes common knowledge, limiting longer-horizon forecastability in fundamentally anchored markets.

The Bitcoin results speak more directly to the literature on expectation-driven dynamics, herding and speculative feedback loops. The combination of persistently high, volatile forward memory and weaker, unstable backward memory complements evidence of explosive or mean-averting behaviour in Bitcoin returns (Turatti et al., 2020) and the strong role of sentiment and herding documented in King and Koutmos (2021); Gurdgiev and O'Loughlin (2020). Our findings suggest that these mechanisms manifest not just in short-run return dynamics, but also at the system level through an anticipatory structure in which future price paths are tightly linked to conditions already forming at time t_0. In this sense, the bidirectional memory approach provides a time series counterpart to models in which dispersed information and speculative trading generate self-reinforcing price trajectories (Angeletos and Lian, 2018; Donier et al., 2015).

The contrast between Brent and Bitcoin also revisits the link between long memory and forecastability. Standard results imply that a process with $d < 0.5$ should be covariance stationary and, in principle, predictable (Granger and Joyeux, 1980; Hoskin, 1981). However, our evidence shows that low backward memory in Bitcoin coexists with forward memory close to unity and sizeable forecast bias, indicating that expectation-driven dynamics can undermine the practical gains from covariance stationarity. Conversely, in Brent, strong backward persistence does not translate into sustained forecastability once forward memory collapses around major structural breaks. This pattern complements forecast studies that document horizon-dependent performance in oil and energy markets (Degiannakis and Filis, 2018; Deng et al., 2011), while highlighting that the balance between backward and forward memory is a key determinant of when long memory structure supports, or fails to support, usable predictability.

In summary, the results place this chapter at the intersection of three strands of literature: long memory and fractional cointegration, anticipatory and sentiment-driven price formation, and forecast performance in commodity and crypto markets. The bidirectional FCVAR framework adds to this work by making the asymmetry between historical and anticipatory components explicit, and by showing how this asymmetry differs systematically between fundamentally anchored and speculative market environments.

3.6 Implications of Bidirectional Memory Dynamics

Our empirical findings highlight critical differences in how forward and backward memory influence forecastability in Brent and Bitcoin markets. Building on these differences, this section discusses the practical implications of bidirectional memory dynamics, particularly concerning forecast reliability, market predictability, and risk management.

3.6.1 Forecast Performance: Brent Oil vs. Bitcoin

Our FCVAR model, applied to Brent oil using both expanding and rolling window approaches, yields comparatively low forecast errors at short horizons. Detailed formulations of error metrics used in this analysis are provided in the Appendix B.4. For example, at a 1–day horizon, the RMSE and LinEx values are modest, and the TS remains within a reasonable range, suggesting that the model slightly outperforms a naive no–change forecast. However, as the horizon extends beyond 10-20 days, RMSE and LinEx escalate, while TS values increase above the ± 4 threshold, indicating systematic forecast bias and growing divergence from actual observed prices. This shift indicates an increase in systematic bias at medium–to–long horizons. The deterioration in forecast accuracy aligns with the empirical memory findings. Although Brent exhibits relatively high backward memory, its forward memory is weaker. In other words, well–defined historical trends in Brent market do not translate into robust long–term predictability once new shocks or structural shifts emerge. However, under the expanding window approach, the forecast errors eventually level off beyond roughly 50 days, reaching a stable but relatively elevated plateau. By contrast, under rolling windows, performance becomes more volatile, and LinEx can reach extreme values (“inf”), illustrating how discarding older data can make the model unresponsive to sudden market movements.

In contrast, forecast performance in the Bitcoin market presents a different picture. Although the RMSE and LinEx values are low in short horizons (1–5 days), the Tracking Signal is often large and negative, indicating systematic underpredictions. Despite this bias, forecast errors remain relatively contained up to around 50 days compared to a naive benchmark. This behaviour suggests that Bitcoin prices are driven more by anticipatory signals, reflecting a more persistent forward memory and lower backward

FIGURE 3.2 : Error profiles of FCVAR out-of-sample of Brent and Bitcoin spot price forecasts in forward and backward direction. Expanding vs rolling window analysis.

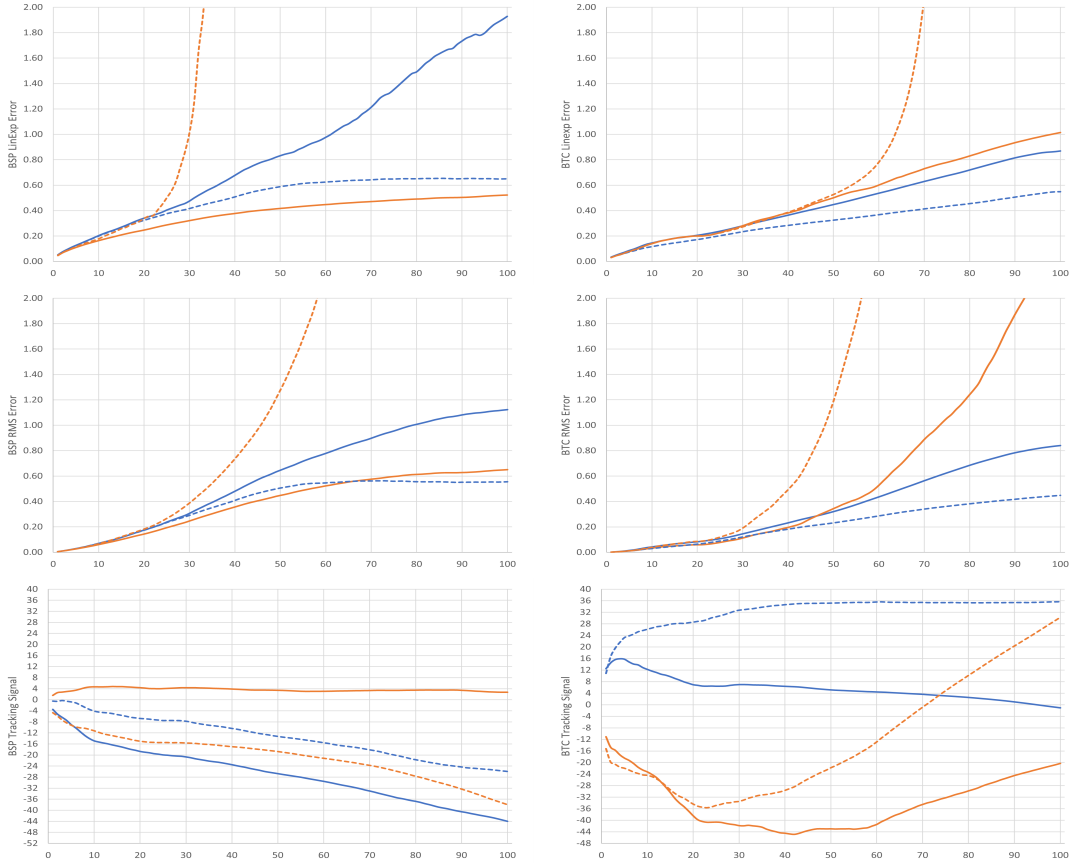

Note: Figure presents error profiles of forward and backward facing Brent and Bitcoin spot price forecasts obtained via FCVAR model using expanding and rolling window analysis. (i) X-Axis - forecast horizon [days]; Y-Axis - Error captured corresponding loss function. (ii) Period analysed: Start Date - 17-01-2018; End Date - 18-02-2021; Total number of Observations - 766 of daily frequency. (iii) Colour convention: Orange line - forward facing forecast error; Blue line - backward facing forecast error. (iv) Line structure: Solid line - expanding window analysis; Dashed line - rolling window analysis. (iv) List of graphs: a) Left Panel Graphs: Brent spot price error analysis; b) Right Panel Graphs: Bitcoin spot price forecast error analysis.

memory, where expectations of future events are rapidly integrated into current prices, enabling the model to capture emerging trends with reasonable accuracy in the short to medium term. However, this forward–looking bias may lead to systematic underpredictions during sudden market shifts. Beyond 50 days in the expanding window context, errors begin to climb more rapidly, and in the rolling window context, errors rise sharply beyond roughly 40 days, emphasising the challenge of balancing immediate momentum with long–term forecast stability.

3.6.2 Interpretation and Broader Implications

The contrasting forecast performance of Brent oil and Bitcoin fundamentally stems from the way each market assimilates information into prices. Manifesting as bidirectional

TABLE 3.6 : Out-of-sample forecast error profiles of Brent and Bitcoin Sport price forecasts for 100 observations in forward and backward directions.

(a) Expanding Window Analysis

Loss	Type	BRENT								BTC							
		1	**5**	**10**	**20**	**30**	**50**	**70**	**100**	**1**	**5**	**10**	**20**	**30**	**50**	**70**	**100**
LinEx	**F**	0.047	0.111	0.165	0.247	0.322	0.418	0.472	0.523	0.031	0.079	0.141	0.200	0.282	0.501	0.730	1.014
	B	0.051	0.126	0.203	0.340	0.475	0.832	1.212	1.929	0.034	0.086	0.145	0.206	0.280	0.448	0.630	0.870
RMSE	**F**	0.004	0.026	0.060	0.143	0.246	0.447	0.575	0.650	0.002	0.011	0.033	0.059	0.113	0.344	0.888	2.575
	B	0.005	0.029	0.070	0.175	0.306	0.646	0.898	1.123	0.002	0.016	0.043	0.084	0.146	0.321	0.564	0.841
TS	**F**	1.555	3.135	4.632	4.328	4.304	3.380	3.304	2.688	-11.10	-18.48	-23.27	-38.49	-41.89	-43.00	-34.54	-20.31
	B	-3.487	-9.012	-14.87	-18.70	-20.69	-26.73	-33.02	-43.98	12.46	15.75	12.22	6.910	6.963	5.122	3.592	-1.092

(b) Rolling Window Analysis

Loss	Type	BRENT								BTC							
		1	**5**	**10**	**20**	**30**	**50**	**70**	**100**	**1**	**5**	**10**	**20**	**30**	**50**	**70**	**100**
LinEx	**F**	0.048	0.116	0.179	0.339	1.015	368.2	inf	inf	0.031	0.080	0.142	0.201	0.272	0.527	2.082	inf
	B	0.052	0.125	0.200	0.324	0.416	0.588	0.642	0.650	0.032	0.074	0.116	0.172	0.234	0.325	0.414	0.549
RMSE	**F**	0.005	0.028	0.067	0.183	0.387	1.276	4.085	31.92	0.002	0.013	0.040	0.086	0.191	1.190	6.727	83.91
	B	0.005	0.029	0.071	0.174	0.291	0.506	0.562	0.555	0.002	0.012	0.030	0.065	0.122	0.232	0.341	0.449
TS	**F**	-4.716	-9.221	-11.24	-15.02	-15.66	-18.75	-23.76	-37.95	-15.20	-21.95	-24.40	-34.36	-33.42	-21.77	-0.798	30.29
	B	-0.450	-0.751	-4.071	-6.738	-7.737	-13.24	-18.09	-25.93	10.84	23.22	26.13	28.59	32.74	35.19	35.42	35.67

Note: The table presents results of forward and backward facing Brent and Bitcoin spot price forecasts obtained using FCVAR model in context of expanding and rolling window analysis, as presented in Figure 3.2. (i) Period analysed: Start Date - 17-01-2018; End Date - 18-02-2021. (ii) Total number of Observations - 766 of daily frequency. (iii) Type description: F - forward facing forecast; B - backward facing forecast; (iv) Loss Functions: a) LinExp - Linear Exponential Loss with symmetry coefficient of - 1; b) RMSE - Root Mean Squared Error; c) TS - Tracking Signal.

persistence, the asymmetry in this process plays a key role in determining forecast accuracy. In the Brent market, the combination of strong backward memory and comparatively weak forward memory implies that, although historical trends are deeply entrenched, new or unforeseen information can rapidly erode predictability over longer horizons. Consequently, forecast accuracy deteriorates once the market diverges from established patterns. However, under an expanding window framework, errors eventually stabilise at a higher level after approximately 50 days. In the Bitcoin market, on the other hand, robust forward memory paired with lower backward memory indicates that anticipatory dynamics dominate price formation. Here, the FCVAR model retains relatively low error rates up to about 50 days, beyond which forecast performance deteriorates more noticeably.

Under expanding windows, Brent forecast errors stabilise at approximately 50–day horizon, showing resilience to slowly evolving fundamental trends. In contrast, Bitcoin errors

continue to increase steadily after a similar horizon, reflecting its exposure to abrupt speculative shifts. Rolling window analysis further highlights these differences, showing earlier degradation in forecast accuracy for both markets due to its sensitivity to recent changes and short–lived memory structures.

These findings indicate a potential forecast threshold of approximately two months. At this threshold, the expanding window forecasts for Brent, although not highly accurate, slow down the deterioration rate, whereas the rapid post 50–day increase in forecast errors for Bitcoin highlights its greater susceptibility to volatility bursts. Under the rolling window approach, the horizon at which forecast errors begin to surge is even shorter, often manifesting a rapid increase around the 30–day mark.

From a practical standpoint, in more mature, fundamentally driven Brent market, a heavy reliance on historical persistence may not be sufficient for extended horizon forecasts, given that forward–looking factors, even if short–lived, appear strong enough to cause sudden regime shifts. In contrast, in the more speculative environment of the Bitcoin market, the FCVAR model retains a degree of predictive power over slightly longer horizons, capturing momentum effects before eventually yielding to sharper fluctuations.

In summary, the results demonstrate that the interplay between historical persistence and anticipatory effects directly shapes the forecast performance. Markets with robust forward memory, such as Bitcoin, maintain better predictability up to a threshold of about 40–50 days (albeit with persistent directional bias). Conversely, markets that rely predominantly on backward memory, such as Brent oil, experience forecast errors that plateau at a higher level but remain relatively stable once that plateau is reached. Therefore, a careful appraisal of both memory dimensions is essential for developing adaptive forecasting approaches capable of handling regime shifts and evolving market sentiment.

3.6.3 VAR Granger Causality Tests

To complement the memory–based analysis, we examine whether components of the Brent and Bitcoin systems contain predictive information in a standard VAR framework. Granger causality tests assess whether the lagged values of one variable, or a group of variables, help forecast another. These tests do not measure the strength or stability of a relationship; they simply indicate whether additional lagged information improves forecasting performance relative to a restricted model.

Applied here, the tests allow us to assess whether elements associated with forward or backward-looking dynamics appear as useful predictors within a linear VAR structure. The results, reported in Table 3.7, show distinct patterns across the two markets. In Brent, predictive links are sporadic and often short-lived, consistent with a fundamentally anchored market in which new information is absorbed quickly. In Bitcoin,

TABLE 3.7 : VAR Granger Causality Tests analysis of causality structure within Brent & Bitcoin markets.

Dependent	Excluded	Expanding Window			Rolling Window		
Variabe	Variabe	Chi-sq	df	Prob.	Chi-sq	df	Prob.
BSP	BRENT d_f	0.241	3	0.963	3.505	5	0.623
	BRENT d_b	6.070	3	0.080***	3.323	5	0.650
	All	6.320	6	0.298	6.715	10	0.752
BRENT d_f	BSP	0.351	3	0.869	2.924	5	0.712
	BRENT d_b	1.432	3	0.853	4.416	5	0.491
	All	1.899	6	0.955	7.380	10	0.689
BRENT d_b	BSP	1.137	3	0.698	19.356	5	0.002*
	BRENT d_f	1.982	3	0.810	5.068	5	0.408
	All	3.058	6	0.888	26.032	10	0.004*
BTC SPOT	BTC d_f	4.688	3	0.196	1.031	3	0.794
	BTC d_b	9.427	3	0.024**	4.993	3	0.172
	All	13.428	6	0.037**	6.811	6	0.339
BTC d_f	BTC SPOT	7.164	3	0.067***	1.659	3	0.646
	BTC d_b	3.916	3	0.271	3.727	3	0.293
	All	11.284	6	0.080***	5.105	6	0.530
BTC d_b	BTC SPOT	4.187	3	0.242	3.576	3	0.311
	BTC d_f	2.074	3	0.557	11.618	3	0.009*
	All	6.880	6	0.332	14.505	6	0.025**

Note: This table displays the results from the VAR Granger Causality Tests, and report the directional causality between Brent and Bitcoin spot prices, as well as their system memory measured in both directions along the time axis where d_f represents forward facing system memory, and BTC d_b depicts backward facing system memory. The analysis was conducted through expanding and rolling windows to capture dynamic interrelationships over time. The chi-squared (Chi-sq) statistics and their corresponding degrees of freedom (df) are reported alongside the probabilities (Prob.) that allow us to infer the significance of the causality. Asterisks denote levels of statistical significance with * indicating $p < 0.01$, ** indicating $0.01 < p < 0.05$, and *** indicating $0.05 < p < 0.1$. The term 'All' refers to the joint hypothesis that all coefficients on the lagged values of the excluded variable are zero. Results with high significance suggest a stronger evidence of causality between corresponding variables.

forward-facing components display clearer predictive content, reflecting the more anticipatory and sentiment-driven nature of its price formation. Jointly, these findings reinforce the interpretation that the forward memory estimates capture economically meaningful anticipatory structure, while also showing how these dynamics manifest within conventional econometric relationships.

3.6.3.1 Brent Market Analysis

In the Brent market, using the expanding window approach, we observe a marginally significant effect ($p < 0.10$) from forward memory on the Brent spot price (BTC panel

in Table 3.7). This finding aligns closely with our previous analyses of memory profiles, where forward memory showed occasional spikes indicative of anticipatory factors becoming more pronounced during certain episodes, notably periods of geopolitical tension or abrupt changes in global demand.

In contrast, backward memory appears to lack significant predictive power over spot prices, consistent with our forecast analyses showing limited long–term predictability despite robust historical persistence. Furthermore, in the rolling window framework, the Brent spot price emerges as a highly significant predictor of forward memory (p-value < 0.01), implying that short–term price fluctuations quickly feed into market expectations. This dynamic is consistent with our forecast results, where Brent exhibited rapid degradation in forecast accuracy beyond short horizons, reflecting how rapidly anticipatory dynamics are reshaped by incoming market information.

3.6.3.2 Bitcoin Market

In the Bitcoin market, the expanding window analysis provides strong evidence (p-value < 0.05) that forward memory significantly predicts Bitcoin spot prices. This finding resonates with our earlier empirical analysis of memory profiles, in which forward memory consistently dominated the Bitcoin market, highlighting its speculative and anticipatory nature. This also corroborates our forecast analysis, where predictive performance in the Bitcoin market remained comparatively robust over longer horizons, indicating that forward–looking signals drive its price formation process. Furthermore, a marginally significant effect (p-value < 0.10) from Bitcoin spot on backward memory further supports the notion that shifts in current market conditions occasionally realign historical persistence patterns.

In the context of the rolling window analysis, the Bitcoin market demonstrates a distinctive causal dynamic, with backward memory significantly influencing forward memory (p-value < 0.01). This effect reflects the interaction between historical momentum and speculative anticipation in shorter–term scenarios. Specifically, recent trends or historical price behaviours appear to actively shape forward–looking market expectations. This aligns well with our observation of intermittent volatility spikes in forward memory. Notably, forward memory does not appear to be strongly predictive of the spot price within rolling windows, emphasising the volatility of shorter–term forecasts in highly speculative markets and confirming our earlier observations of increased forecast error as horizons extend.

In summary, the VAR Granger Causality analyses highlight key differences between the Brent and Bitcoin markets with respect to the interactions between memory components and price formation. For Brent, historical persistence remains influential yet does not strongly forecast future prices, particularly when short–term fluctuations quickly

alter forward–looking dynamics, thus explaining the observed rapid deterioration in longer–horizon forecasts. In contrast, Bitcoin displays a pronounced role of anticipatory factors, strongly reflected in forward–facing memory, with historical price behaviour dynamically contributing to shaping these forward–looking expectations, particularly in the short–term. These distinct causal structures reaffirm our earlier conclusions drawn from memory profile dynamics and forecast performance analyses: the Brent market predictability erodes rapidly amid evolving information flows and structural shifts, while Bitcoin's speculative nature allows anticipatory dynamics to sustain forecastability, albeit with a heightened risk of volatility–induced prediction errors.

3.7 Conclusion

In this study, we investigated the interplay between historical persistence and anticipatory effects in financial markets by employing an FCVAR model with a bidirectional estimation of the fractional integration parameter d. By distinguishing between backward memory, which captures the persistent influence of historical price movements, and forward memory, which by logical symmetry reflects anticipatory adjustments, our analysis provides novel insights into the mechanisms that underpin price formation in two contrasting market environments: the fundamentally-driven Brent and the sentiment-driven Bitcoin markets.

Our empirical findings indicate that Brent oil prices exhibit strong backward memory, coupled with relatively subdued forward memory. This suggests that, although historical trends in oil are robust, rapid assimilation of new information leads to a swift degradation of forecast accuracy over longer horizons. In contrast, Bitcoin prices are characterised by lower backward memory and significantly higher, more volatile forward memory, highlighting the dominant role of speculative, anticipatory forces. These differences are supported by our VAR Granger causality tests, which demonstrate that forward memory is systematically embedded in current prices, albeit with differing intensities across the two markets.

The implications of these results are multifaceted. In mature, fundamentally-driven markets such as oil, reliance on historical persistence alone does not ensure robust long–term forecastability, particularly when unexpected shocks or regime changes occur. Conversely, in speculative markets like Bitcoin, a strong forward memory component can offer a forecasting advantage over short to medium horizons, even amidst heightened volatility. These findings emphasise the need for forecasting models to consider both historical trends and forward–looking signals to improve predictive performance and support effective risk management.

Although our study provides a robust framework for analysing bidirectional memory dynamics, further research is warranted to explore the effects of structural breaks and

to develop adaptive models that can respond to rapid shifts in market sentiment. Expanding this analysis to additional asset classes may also offer broader insights into the generalisability of our findings and contribute to a more comprehensive understanding of market efficiency.

Overall, this research contributes to the growing literature on long memory and anticipatory effects by highlighting how the balance between historical persistence and forward–looking adjustments critically shapes forecast performance. In doing so, it offers new perspectives on the determinants of predictability in financial markets and paves the way for the development of more resilient forecasting methodologies.

Chapter 4

A Memory-driven Multi-period Capital Asset Pricing Model

Abstract

In this paper we present a unified framework for asset prices by integrating dynamic effects of memory - broadly, varied persistence d profiles of assets with concavity-based multi-period compounding. We derive empirically meaningful closed-form effective variance and skewness under a standard $ARFIMA(0,\ d,\ 0)$ model. Application to sixty large-cap US equities shows that assets with negative d (antipersistence) yield mean-reversion benefits, circa 15% lower effective variance and 25% higher Sharpe ratios, whereas assets with positive d (persistent) incur greater variance drag but provide momentum-style skewness bonuses when exploited. Sorting and aligning portfolios by memory profiles, and incorporating both variance and skewness adjustments into a memory-tilted optimisation delivers additional improvements in expected log-growth and risk-adjusted performance. Augmenting CAPM with a memory HML (mHML) factor reduces out-of-sample forecast error by up to 44%. Memory-aware efficient frontier analysis expands the frontier by over 1% at a 20% volatility target, with further uplift from skewness-informed tilts. Robustness checks conducted on the Japanese equity market confirm the cross-market validity and consistency of our empirical results. These results establish that profiling assets by memory enables explicit mean-reversion and momentum strategies, systematically mitigating volatility drag and harnessing skewness for superior long-term performance under realistic market conditions.

Keywords: Memory-augmented CAPM; effective variance; effective skewness; momentum and antipersistence: compound growth optimisation.

JEL Classification: G12; C32; C58; G11: C61

4.1 Introduction

Since the seminal work of Markowitz (1952), mean-variance analysis has remained central to portfolio theory and asset allocation practice. Markowitz's framework elegantly simplifies the complex investment problem into a single-period optimisation that trades off expected returns against variance as a proxy for risk. However, despite its widespread acceptance in both academia and practice, the classical mean-variance approach relies heavily on restrictive assumptions that often diverge significantly from empirical financial market realities (Campbell et al., 1997; Fernholz and Fernholz, 2002). Critically, the standard mean-variance paradigm assumes that investment decisions are made within a single-period context, implicitly ignoring the compounding effects of returns over multiple periods; real-world investors instead focus on wealth accumulation over extended horizons, where compounding growth dynamics and volatility drag become significant (Luenberger, 1998). Single-period optimisation, repeated iteratively, fails to explicitly address these multi-period compounding effects and thus provides suboptimal guidance for long-term portfolio construction.

Another central limitation of classical mean-variance analysis is twofold: first, it assumes asset returns are independent and identically distributed (i.i.d.), typically Gaussian, thus excluding the serial correlation and long-range dependence (memory) that empirical research has shown to be pervasive in financial markets (Ding et al., 1993; Baillie, 1996; Granger and Joyeux, 1980; Lo, 1991). Since single-period variance accumulates linearly only under independence, whereas positive or negative autocorrelation causes risk to inflate or deflate non-linearly, neglecting persistence phenomena ultimately leads to mismeasurement of risk, inefficient diversification (in some cases, naive equal-weighting was found to outperform optimised portfolios built on poorly estimated inputs (DeMiguel et al., 2009)). Second, the framework rests on mean-variance preferences, equivalent to quadratic utility (or logarithmic utility under restrictive assumptions), which fail to capture the economic significance of volatility drag and the compounding of returns over multi-period horizons (Merton, 1975; Cover, 1999). A more appropriate theoretical foundation would need to explicitly integrate logarithmic utility, directly addressing compound growth dynamics, while retaining variance as a pragmatic risk measure (Kelly, 1956; Pratt, 1978).

While theoretical portfolio selection traditionally assumes frictionless markets, such as absent transaction costs, taxes, and other frictions, practical portfolio implementation must inevitably confront such real-world complexities. Therefore, a robust portfolio theory should not only present frictionless derivations but also validate its conclusions under realistic market conditions, quantifying the impact of transaction costs and other frictions on investment performance (Brandt et al., 2009; Gârleanu and Pedersen, 2013). Lastly, classical mean-variance theory assumes a static investment universe characterised

by fixed asset covariances. Empirical evidence, however, suggests that risk characteristics, including asset volatilities and correlations, evolve significantly over time (Engle, 2002). Thus, any realistic portfolio selection framework must incorporate dynamic estimation of risk measures, accommodating evolving temporal and cross-sectional dependencies.

Motivated by these limitations and by the gaps in the extant literature, this chapter introduces a unified theoretical and empirical framework that explicitly captures multi-period volatility drag and temporal dependence through fractional memory dynamics modelled via an $ARFIMA(0, d, 0)$ process. We analytically demonstrate how the fractional integration parameter d governs the hidden cost or benefit of volatility, providing a rigorous economic foundation for embedding memory effects directly into portfolio optimisation. Empirically, we show that augmenting the CAPM with a memory-based HML factor yields statistically significant alphas and materially improves risk-adjusted performance relative to the conventional framework. Efficient-frontier analysis further demonstrates that incorporating memory shifts the frontier into lower-risk and higher-return regions, with the largest point-wise improvements occurring at investor-relevant volatility targets.

We validate these findings through an international robustness exercise using Japanese equity data. Despite differences in market structure, liquidity, and institutional participation, the empirical patterns persist: memory-sorted portfolios produce systematic differences in variance, skewness, and compounding efficiency, and the memory-augmented CAPM continues to deliver superior forecasting accuracy. This confirms that the economic relevance of memory is not market-specific, enhancing the generalisability of the framework.

Jointly, the results show that serial dependence in returns is not merely a statistical artefact but a systematically priced component of risk that shapes long-horizon wealth accumulation. By linking concavity-based compounding with fractional integration, the chapter delivers two principal contributions: (i) closed-form expressions for effective variance and volatility drag under serial dependence, and (ii) empirical evidence that antipersistence constitutes a priced characteristic that can be exploited through dynamic portfolio strategies.

Although this chapter employs sorted portfolios and a factor-style construction for the memory spread, the economic interpretation differs fundamentally from that of the traditional anomaly literature. Classical factors, value, size, momentum, profitability and investment are static cross-sectional characteristics designed to explain average return differences across firms (Fama and French, 1993, 2015). Fractional memory, by contrast, is a *structural time series property* of the return process, reflecting how shocks propagate across horizons through long-range dependence (Granger and Joyeux, 1980; Hoskin, 1981; Baillie, 1996). Its economic impact arises through the temporal geometry of wealth,

effective variance and volatility drag (Kelly, 1956; Luenberger, 1998), skewness accumulation under serial correlation (Hoskin, 1981), and sequence of returns-driven drawdown risk (Grossman and Zhou, 1993; Chekhlov et al., 2005). These mechanisms materially influence long-horizon investor outcomes yet do not appear in characteristic-based factor models, whose risk structure is inherently cross-sectional (Cochrane, 2005). Memory therefore introduces an intertemporal dimension of risk that complements, rather than replicates, the cross-sectional structure of existing anomalies.

The remainder of the chapter is structured as follows. Section 4.2 reviews the relevant academic literature, and Section 4.3 develops the theoretical foundations of the memory–based multi-period framework. Section 4.5 describes the data and sample construction, while Section 4.4 outlines the empirical methodology. Section 4.6 presents the main empirical findings, including the asset-pricing tests, portfolio performance results and efficient frontier analysis. Section 4.7 reports the replication study using Japanese equities to assess robustness. Finally, Section 4.9 concludes with limitations and directions for future research.

4.2 Literature

Research on portfolio optimisation and asset pricing has evolved considerably from the foundational single–period mean–variance framework of Markowitz (1952) toward richer multiperiod settings. Early extensions sought to improve parameter estimation (Jorion, 1986), incorporate estimation uncertainty (DeMiguel et al., 2009), and account for frictions and state dependence through dynamic allocation models (Gârleanu and Pedersen, 2013; Celikyurt and Özekici, 2007; Dai et al., 2021). More recent work examines long horizon portfolio rules that combine gradual capital injection, market timing and sparse investment structures (Shi et al., 2025). Despite these advances, most frameworks retain a cross–sectional focus in which risk is evaluated period by period, and thus abstract from the temporal ordering of returns and the compounding mechanisms that shape long horizon wealth.

A parallel strand studies the role of persistence and long memory in asset returns. Classical contributions by Granger and Joyeux (1980) and Hoskin (1981) show that fractional integration generates hyperbolic decay in autocorrelations, implying that shocks propagate across long horizons. Empirical evidence confirms that such dependence structures appear in both returns and volatility (Ding et al., 1993; Baillie, 1996; Pavlova et al., 2014), and that ignoring long-run dependence can lead to horizon–specific mispricing or excessive risk-taking (León and Revéiz, 2010; Rodríguez, 2011). More recent heterogeneous–frequency and multi-scale models (Delpini et al., 2012; Chikhi et al., 2019) further highlight that persistence shapes risk dynamics over multiple horizons. Yet these contributions remain largely reduced-form: they model persistence statistically but do

not provide an analytical mapping between memory and the moments of compounded wealth.

Sequence-of-returns effects, a phenomenon widely recognised in retirement and decumulation studies (Pfau, 2011; Doran et al., 2012; Clare et al., 2020), demonstrate that identical average returns can produce markedly different long-horizon outcomes once shocks arrive in different orders. Traditional mean–variance and intertemporal models (Markowitz, 1952; Merton, 1973) abstract from this channel by assuming period independence. Recent work in intertemporal pricing, including recursive utility frameworks (Epstein and Zin, 1989) and long-run-risks models (Bansal and Yaron, 2004; Zhou et al., 2015), acknowledges that low-frequency dynamics are priced, but they do not incorporate long memory structures explicitly. Related studies on downside and drawdown risk (Ang et al., 2006; Anthonisz et al., 2017; Grossman and Zhou, 1993; Chekhlov et al., 2005) show that path-dependent risk bears independent economic significance, yet these frameworks do not link such risk to systematic temporal dependence.

Recent developments in modelling further emphasise the predictive and economic value of sequential information. Deep sequence architectures (Cong et al., 2021) and text-derived intertemporal factors (Cong et al., 2024) demonstrate that ordered return histories contain dynamic signals missed by static factor models. Complementary evidence from volatility targeting and intertemporal risk management strategies (Moreira and Muir, 2017; Barroso and Santa-Clara, 2015) shows that time varying serial dependence materially alters realised Sharpe ratios and optimal allocation rules, highlighting the role of temporal structure in long horizon decisions. Multi-horizon sequence risk and drawdown analyses (Grossman and Zhou, 1993; Chekhlov et al., 2005; Rockafellar et al., 2000) further highlight that path-dependent losses cannot be captured within single period mean–variance frameworks. Collectively, these strands strengthen the motivation for models that explicitly incorporate temporal dependence, providing a conceptual foundation for a memory-augmented, multi-period approach.

Although we employ sorted portfolios and construct a factor-style memory spread, the economic interpretation of fractional memory differs fundamentally from the static characteristic-based anomalies considered in the Fama-French framework (Fama and French, 1993, 2015). As discussed in Section 4.1, memory is a structural time series property that shapes the propagation of shocks and therefore affects multi-period compounding and sequence risk. Its role in the asset-pricing framework developed here is intertemporal rather than cross-sectional, complementing rather than replicating the behaviour of conventional factors.

In summary, these strands point toward a conceptual gap. Existing models acknowledge dynamic risk, but few integrate fractional memory into a coherent framework that links persistence to the effective moments of compounded wealth. The literature recognises sequence effects empirically, but lacks a structural explanation of how temporal

dependence shapes multi-horizon risk-adjusted performance. This chapter addresses this gap by deriving closed-form expressions for effective variance and skewness under fractional integration, embedding memory directly within the compounding process, and examining its empirical implications for wealth preservation, cross-sectional pricing and portfolio efficiency. In doing so, the chapter positions memory as an intertemporal state variable that governs long-horizon economic outcomes rather than merely a statistical descriptor of return autocorrelation.

4.3 Theoretical Framework - Compounding-Adjusted Risk

In this section, we develop the mathematical foundations for analysing how log concavity induces volatility drag in compound growth, and how the control of temporal dependencies can mitigate this effect. We first formalise risk aversion under logarithmic utility and, via Jensen's inequality, show how return variability reduces expected log-wealth. We then derive a second-order Taylor expansion to quantify the single-period variance drag and extend it to a multi-period variance decomposition that incorporates serial correlation. In particular, we demonstrate that the fractional memory parameter d governs aggregate variance, and hence the implicit cost or benefit of volatility, by weighting each lag-k autocovariance. Finally, we embed these insights in an $ARFIMA(0,\ d,\ 0)$ model (covariance stationary for $-0.5 < d < 0.5$) and derive a closed-form, multi-period expression for volatility drag that explicitly captures the impact of each autocorrelation.

4.3.1 Log-Utility, Jensen's Inequality and Volatility Drag

Let an investor's utility be

$$U(W) = \ln(W), \quad W > 0 \tag{4.1}$$

and suppose initial wealth W_0 is invested over single-period yielding random return r with $r > -1$ almost surely and $\mathbb{E}[|\ln(1+r)|] < \infty$. Then

$$W_1 = W_0(1+r), \quad \ln W_1 = \ln W_0 + \ln(1+r) \tag{4.2}$$

so that

$$\mathbb{E}[\ln W_1] = \ln W_0 + \mathbb{E}[\ln(1+r)] \tag{4.3}$$

Since $\ln(\cdot)$ is strictly concave on $(0, \infty)$, Jensen's inequality implies:

$$\ln\Big(\mathbb{E}[W_1]\Big) = \ln\Big(W_0(1+\mathbb{E}[r])\Big) \ \geq \ \mathbb{E}[\ln W_1] \tag{4.4}$$

with strict inequality whenever $\mathrm{Var}(r) > 0$. Hence the gap

$$\ln\Big(1 + \mathbb{E}[r]\Big) - \mathbb{E}\Big[\ln(1+r)\Big] \;=\; \underbrace{\Delta_{\text{drag}}}_{\geq 0} \tag{4.5}$$

quantifies the loss in expected log-wealth caused solely by return variability. Hence, a risk-averse investor strictly prefers, in terms of expected utility, the certain return $\mathbb{E}[r]$ to any volatile prospect with the same mean. Figure 4.1 illustrates how volatility affects the investor's expected log–return.

FIGURE 4.1 : Illustration of The Effect of Volatility Drag on Investor's Expected Log-Returns.

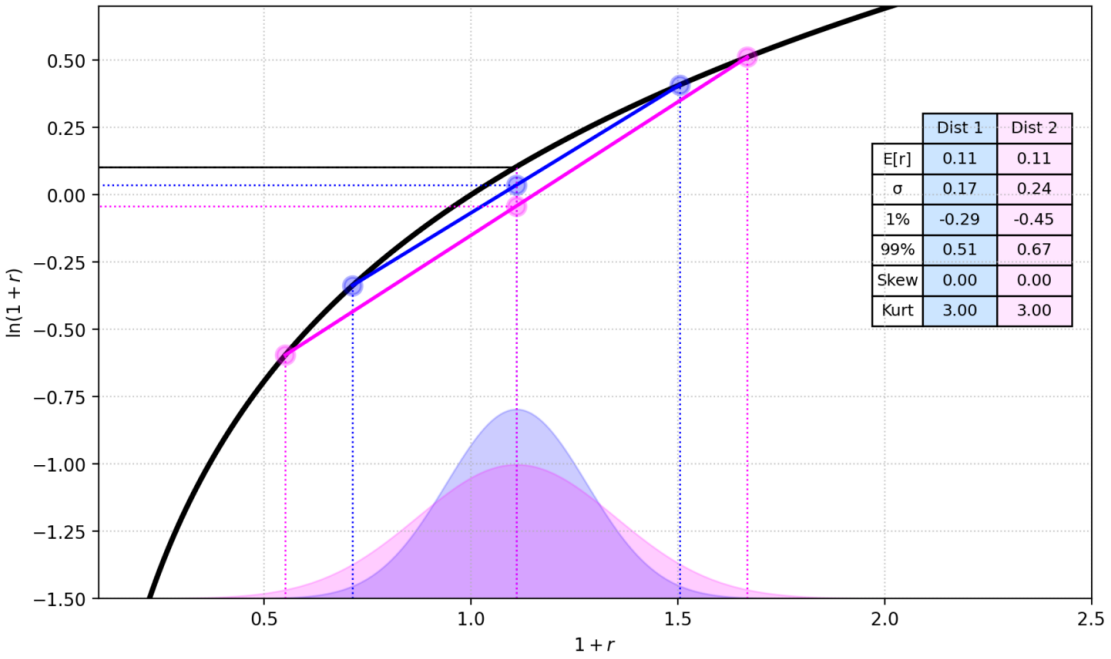

Note: The solid black curve depicts $y = \ln(1 + r)$, illustrating its strict concavity. Two dashed secant lines, blue for the low-volatility case and magenta for the high-volatility case, are drawn between their respective extreme returns, both having the same arithmetic mean $r = 0.11$ (i.e. $1 + r = 1.11$). At this common midpoint, the vertical bars measure the Jensen volatility drag, namely the gap between the chord midpoint (expected log under the assumption of normally distributed returns) and the true ln value on the curve. Horizontal dashed projections to the y-axis show the absolute log-return levels of the chord midpoints (coloured lines) versus the actual log-return (black line). The markedly larger magenta gap compared to the blue gap visually demonstrates how higher return dispersion amplifies the expected-log shortfall, whereas lower dispersion attenuates it.

The application of Jensen's inequality to concave utility functions, such as logarithmic utility $U(W) = \ln(W)$, is well documented. Jensen (1906) originally proved the concavity-based bound on expectations, Pratt (1978) showed how this concavity underpins measures of investor risk aversion, and Kelly (1956) demonstrated that maximising expected logarithmic utility yields the optimal growth portfolio. Merton (1975) extended these insights into a continuous-time framework deriving optimal consumption and investment rules under logarithmic utility. Luenberger (1998) provided a comprehensive

discrete-time treatment of volatility effects and their implications for portfolio choice.

4.3.2 Taylor Expansion and Single–Period Approximation

To quantify drag under moderate volatility we employ a second–order Taylor expansion of $\ln(1+r)$ about $r = \mu := \mathbb{E}[r]$, valid when $|r-\mu|$ is not too large[1]. Writing

$$\ln(1+r) = \ln(1+\mu) + \frac{r-\mu}{1+\mu} - \frac{(r-\mu)^2}{2(1+\mu)^2} + O\Big((r-\mu)^3\Big) \tag{4.6}$$

Taking expectations and noting $\mathbb{E}[r-\mu] = 0$ and $\mathbb{E}[(r-\mu)^2] = \sigma^2$ gives

$$\mathbb{E}[\ln(1+r)] \approx \ln(1+\mu) - \frac{\sigma^2}{2(1+\mu)^2}, \tag{4.7}$$

where single-period half variance drag term, also known as 'variance drain' (Messmore, 1995), emerges transparently from the concave nature of logarithmic utility. Fixed by the return variance σ^2, this term quantifies the extent to which randomness in r reduces the expected log–return and follows from Jensen's inequality applied to the concave ln function (Cover, 1999, Ch. 2). Note that the higher-order cumulant corrections, proportional to the return skewness and higher moments are omitted from Eq (4.7) for tractability, since they enter at $O\Big(\sigma^3\Big)$ and beyond and are typically negligible under moderate volatility.

4.3.3 Multi-Period Drag and Effective Variance

For illustrative simplicity we assume identically distributed returns $r_1, \ldots, r_T$ with common mean μ and variance σ^2. Therefore, terminal wealth under passive buy-and-hold strategy is

$$W_T = W_0 \prod_{t=1}^{T}(1+r_t), \qquad \ln W_T = \ln W_0 + \sum_{t=1}^{T} \ln\Big(1+r_t\Big) \tag{4.8}$$

Linearity of expectation yields

$$\begin{aligned} \mathbb{E}\Big[\ln W_T\Big] &= \ln W_0 + \sum_{t=1}^{T} \mathbb{E}\Big[\ln(1+r_t)\Big] \\ &\approx \ln W_0 + T\left(\ln(1+\mu) - \frac{\sigma^2}{2\,(1+\mu)^2}\right) \end{aligned} \tag{4.9}$$

[1] See the expansion in Campbell et al. (1997), and the general log approximation in Hamilton (1994, Ch. 1).

Thus the average per-period log-growth rate is

$$\frac{1}{T}\,\mathbb{E}[\ln W_T] \approx \ln(1+\mu) - \frac{\sigma^2}{2\,(1+\mu)^2} \tag{4.10}$$

This linear accumulation of drag over T periods is discussed at length in the stochastic portfolio theory literature (Fernholz and Fernholz, 2002, Ch. 3).

To extend Eq. (4.10) to the serially-dependent case, start with the variance of the T-period cumulative return:

$$\mathrm{Var}\Big(\sum_{t=1}^{T} r_t\Big) = \sum_{t=1}^{T} \mathrm{Var}(r_t) \;+\; 2\sum_{i=1}^{T}\sum_{j=i+1}^{T} \mathrm{Cov}(r_i, r_j) \tag{4.11}$$

Under homoscedastic assumption $\mathrm{Var}(r_t) = \sigma^2$ and $\mathrm{Cov}(r_i, r_j) = \sigma^2\rho(j-i)$, thus

$$\mathrm{Var}\Big(\sum_{t=1}^{T} r_t\Big) \;=\; \sigma^2\Big[T + 2\sum_{k=1}^{T-1}(T-k)\rho(k)\Big] \tag{4.12}$$

Dividing by T gives the *effective* per-period variance which governs how serial dependence inflates or deflates the variance under independence[2]

$$\sigma^2_{\text{eff}}(T) \;=\; \frac{1}{T}\,\mathrm{Var}\Big(\sum_{t=1}^{T} r_t\Big) \;=\; \frac{\sigma^2}{T}\Big[T + 2\sum_{k=1}^{T-1}(T-k)\rho(k)\Big]. \tag{4.13}$$

so that the effective volatility drag over T periods

$$\Delta^{\text{eff}}_{\text{drag}}(T) = \frac{\sigma^2_{\text{eff}}(T)}{2\,(1+\mu)^2} \tag{4.14}$$

depends critically on the sign and magnitude of the serial correlations. If $\rho(k) = 0$ for all k, then

$$\sigma^2_{\text{eff}}(T) = \sigma^2 \quad \text{and} \quad \Delta^{\text{eff}}_{\text{drag}}(T) = \frac{\sigma^2}{2(1+\mu)^2}, \tag{4.15}$$

naturally collapse to the *i.i.d.* case. From economic interpretation perspective this implies that the positive serial correlation amplifies compounding drag, while negative serial correlation reduces it.

A full derivation of the variance decomposition, with relaxed *i.i.d.* assumption to allow for autocorrelation in returns and demonstrating its effect on aggregate variance and volatility drag, is provided in Section 4.3.5.

[2] A similar multi–period weighting of autocorrelations is derived in Campbell and Viceira (2002, Ch. 3)

4.3.4 Fractional Memory and the $ARFIMA(0,\ d,\ 0)$Model

Consider the pure fractional $ARFIMA(0,\ d,\ 0)$ model in the covariance stationary regime $-0.5 < d < 0.5$, first introduced by Granger and Joyeux (1980). One may derive its properties via spectral methods or directly from the binomial series definition of $(1-B)^{-d}$ (see Hoskin 1981 or Beran 1994, Ch. 2). Define

$$\gamma(k) \;=\; \mathrm{Cov}\Big(r_t,\ r_{t+k}\Big), \qquad \gamma(0) \;=\; \mathrm{Var}(r_t). \tag{4.16}$$

Then the univariate lag–k autocorrelation is

$$\rho(k) \;=\; \frac{\gamma(k)}{\gamma(0)}\,, \tag{4.17}$$

which satisfies the power-law asymptotic

$$\rho(k) \;=\; \frac{\gamma(k)}{\gamma(0)} \;\sim\; \frac{\Gamma(1-2d)}{\Gamma(d)\,\Gamma(1-d)}\; k^{2d-1} \quad \text{as } k \to \infty. \tag{4.18}$$

In particular, when $-0.5 < d < 0$, one has $\Gamma(d) < 0$ while $\Gamma(1-2d)$ and $\Gamma(1-d)$ remain positive, so the prefactor

$$\frac{\Gamma(1-2d)}{\Gamma(d)\,\Gamma(1-d)} \tag{4.19}$$

is negative and hence $\rho(k) < 0$ for large k. This negative, power-law decay is the mathematical signature of antipersistence (Baillie, 1996).

At the boundary $d = |0.5|$, the simple power-law in Eq. (4.18) transitions to a logarithmic form; see Sowell (1992) for details. For $|d| \geq 0.5$, covariance stationarity no longer holds and these asymptotic expressions break down.

4.3.5 Variance Decomposition in Presence of Autocorrelation in Returns

Define $\ell_t = \ln(1 + r_t)$. Applying a second-order Taylor expansion to $\ln(1 + r_t)$ around $r_t = \mu$ yields Eq (4.6). Since

$$\ell_t \approx \underbrace{\ln(1+\mu) - \Delta_{\text{drag}}}_{C} + \underbrace{\frac{r_t - \mu}{1+\mu}}_{L_t} \tag{4.20}$$

and the constant term C has zero variance, all the variation in ℓ_t arises from the zero-mean component L_t. In this decomposition, C represents the expected log-return after accounting for the variance-induced drag Δ_{drag}, effectively shifting the mean downward, while L_t captures fluctuations around that shifted mean. Therefore, the shortfall in

compounded growth is first due to the mean reduction in C and second from any additional widening of the log-return distribution due to serial correlation, which enters via the effective variance of the sum of returns.

Assuming a weak stationarity of $\{r_t\}$, so $\mathrm{Cov}(r_i, r_j) = \rho(j-i)\,\sigma^2$ (Brockwell and Davis, 1991, Ch. 1), the variance of the sum of log–returns is expressed as

$$\mathrm{Var}\Big(\sum_{t=1}^{T} \ell_t\Big) \;\approx\; \frac{1}{(1+\mu)^2}\,\mathrm{Var}\Big(\sum_{t=1}^{T} r_t\Big) \tag{4.21}$$

But

$$\mathrm{Var}\Big(\sum_{t=1}^{T} r_t\Big) = T\,\sigma^2 + 2\sum_{i=1}^{T}\sum_{j=i+1}^{T}\mathrm{Cov}(r_i, r_j) \tag{4.22}$$

Hence

$$\mathrm{Var}\Big(\sum_{t=1}^{T} \ell_t\Big) \;\approx\; \frac{1}{(1+\mu)^2}\Big[T\,\sigma^2 + 2\sum_{i<j}\mathrm{Cov}(r_i, r_j)\Big] \tag{4.23}$$

Equivalently, since $\mathrm{Cov}(r_i, r_j) = \rho(j-i)\,\sigma^2$ the relationship can also be expressed as:

$$\mathrm{Var}\Big(\sum_{t=1}^{T} \ell_t\Big) \approx \frac{\sigma^2}{(1+\mu)^2}\Big[T + 2\sum_{k=1}^{T-1}(T-k)\,\rho(k)\Big] \tag{4.24}$$

This clearly shows that negative $\rho(k)$ reduces aggregate variance, and thus volatility drag, relative to the *i.i.d.* case.

In contrast, when returns follow an $ARFIMA(0,\ d,\ 0)$ with $0 < d < 0.5$, the double sum term inside square brackets diverges faster than linearly in T because $\sum_k \rho(k) = \infty$ (see Beran 1994, Baillie 1996). Empirical evidence of such long memory behaviour in absolute and squared returns appears in Ding et al. (1993). Hence, persistent long–memory amplifies aggregate variance, and so amplifies volatility drag, compared to the antipersistent or short memory cases.

4.3.6 Portfolio-Level Variance and Drag

To extend our univariate variance decomposition to a portfolio of N assets, we must account for both within and cross asset serial dependence. For two assets i and j, we follow the derivations provided by Brockwell and Davis (1991) and Hamilton (1994) and define their lag-k autocovariance and cross-covariance as

$$\gamma_{ii}(k) = \mathrm{Cov}\Big(r_{i,t}, r_{i,t+k}\Big), \quad \gamma_{ij}(k) = \mathrm{Cov}\Big(r_{i,t}, r_{j,t+k}\Big), \tag{4.25}$$

and obtain their corresponding lag-k correlation through the normalisation

$$\rho_{ij}(k) = \frac{\gamma_{ij}(k)}{\sqrt{\gamma_{ii}(0)\,\gamma_{jj}(0)}}, \tag{4.26}$$

with $\sigma_i^2 = \gamma_{ii}(0)$ denoting the variance of asset i.

Letting $w = (w_1, \ldots, w_N)^\top$ denote the portfolio weight vector, with $\sum_i w_i = 1$, the return on the portfolio at time t is

$$R_{p,t} = \sum_{i=1}^{N} w_i \, r_{i,t}, \tag{4.27}$$

and the cumulative return over T periods is $\sum_{t=1}^{T} R_{p,t}$. The variance of this cumulative return is given by

$$\mathrm{Var}\Big(\sum_{t=1}^{T} R_{p,t}\Big) = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i \, w_j \, \mathrm{Cov}\Big(\sum_{t=1}^{T} r_{i,t}, \sum_{t=1}^{T} r_{j,t}\Big), \tag{4.28}$$

where each covariance term may be decomposed as

$$\mathrm{Cov}\Big(\sum_{t=1}^{T} r_{i,t}, \sum_{t=1}^{T} r_{j,t}\Big) = \sigma_i \, \sigma_j \Big[T + 2\sum_{k=1}^{T-1} (T-k)\, \rho_{ij}(k)\Big], \tag{4.29}$$

Substituting yields the total portfolio variance:

$$\mathrm{Var}\Big(\sum_{t=1}^{T} R_{p,t}\Big) = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i \, w_j \, \sigma_i \, \sigma_j \Big[T + 2\sum_{k=1}^{T-1} (T-k)\, \rho_{ij}(k)\Big]. \tag{4.30}$$

This multivariate extension builds directly on the univariate long-horizon analysis of Campbell and Viceira (2002).

The per-period effective portfolio variance is obtained by dividing by T, giving

$$\sigma^2_{\mathrm{eff},p}(T) = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i \, w_j \, \sigma_i \, \sigma_j \Big[1 + \frac{2}{T}\sum_{k=1}^{T-1} (T-k)\, \rho_{ij}(k)\Big]. \tag{4.31}$$

In the special case of equal weights $w_i = 1/N$, due to variance being homogeneous of degree two, the constant weight $1/N$ factors out as $1/N^2$, yielding

$$\sigma^2_{\mathrm{eff,p}}(T) = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} \sigma_i \, \sigma_j \Big[1 + \frac{2}{T}\sum_{k=1}^{T-1} (T-k)\, \rho_{ij}(k)\Big]. \tag{4.32}$$

In both cases associated effective portfolio-level volatility drag is

$$\Delta^{\mathrm{eff}}_{\mathrm{drag,p}}(T) = \frac{\sigma^2_{\mathrm{eff,p}}(T)}{2\,(1+\mu_p)^2}, \quad \mu_p = \sum_{i=1}^{N} w_i \, \mu_i. \tag{4.33}$$

This unified framework highlights how both within asset memory, captured through $\rho_{ii}(k)$, and cross asset temporal dependencies, represented by $\rho_{ij}(k)$ for $i \neq j$, jointly contribute to the growth of aggregate portfolio variance.

Of particular note is the role of serial correlation in producing what may be interpreted as temporal diversification, a phenomenon through which the dynamic structure of return dependence over time complements traditional cross-sectional diversification. The autocorrelation structure therefore emerges not just as a statistical feature but also as an economically significant determinant of the compound return profile, with clear implications for portfolio construction and risk management.

4.3.7 Memory-Augmented Mean-Variance Optimisation

Building on our multi-period variance decomposition, we now incorporate a *memory risk premium* on the expected return side. Let

$$\mu_i = \mathbb{E}[r_i], \quad d_i \in (-0.5,\, 0.5) \tag{4.34}$$

denote the mean return and fractional memory parameter of each asset. Following the impact of memory on returns through volatility drag, we propose a linear tilt in mean returns due to persistence dynamics:

$$\tilde{\mu}_i = \mu_i + \kappa\,(-d_i) + \epsilon_i, \tag{4.35}$$

where $\kappa > 0$ is estimated by regressing μ_i on $-d_i$. This memory tilt formula is fundamentally an *asset–level* adjustment. The fractional memory d_i if each asset generates a bespoke shift $\kappa(-d_i)$ in its expected return. These shifted expectations $\{\tilde{\mu}_i\}$ are then carried into the familiar mean-variance optimiser. For any portfolio weight vector w, the resulting portfolio-level expectation is

$$\tilde{\mu}_p \;=\; \sum_i w_i\,\tilde{\mu}_i \;=\; \sum_i w_i\,\mu_i \;+\; \kappa \sum_i w_i\,(-d_i). \tag{4.36}$$

Because the tilt operates on individual assets, its impact on the portfolio arises endogenously via the chosen weights. One could equivalently define a *composite* portfolio memory $d_p = \sum_i w_i\, d_i$ and write $\tilde{\mu}_p = \mu_p + \kappa(-d_p)$, but only the asset-level formulation preserves the granularity needed when w itself depends on d.

Simultaneously, serial dependence enters the variance side through the per-period effective covariance:

$$\Big(\sigma_{\text{eff}}(T)\Big)_{ij} = \sigma_i\,\sigma_j \frac{1}{T}\Big[T + 2\sum_{k=1}^{T-1}(T-k)\,\rho_{ij}(k)\Big]. \tag{4.37}$$

The resulting memory-augmented mean-variance programme is

$$\max_{w\geq 0,\,\mathbf{1}^\top w=1} \;\; \tilde{\mu}^\top w \;-\; \tfrac{1}{2}\, w^\top \sigma_{\text{eff}}(T)\, w. \tag{4.38}$$

Ignoring the non-negativity constraint for illustration, the first–order condition yields the familiar closed-form:

$$w^* \propto \sigma_{\text{eff}}(T)^{-1}\,\tilde{\mu}, \tag{4.39}$$

with the proportionality factor chosen so that $\sum_i w_i^* = 1$. Thus, antipersistence reduces both the variance penalty and increases the expected return, jointly pushing the efficient frontier into regions unattainable by conventional mean-variance or by variance-only adjustments alone.

This unified treatment provides a direct bridge from the single–asset *ARFIMA* dynamics to practical portfolio management insights, by quantifying exactly how antipersistence reduces and long–memory amplifies compound growth drag in a diversified setting.

4.4 Empirical Methodology

4.4.1 Economic Significance of Memory

Memory estimation. To empirically validate the economic significance of fractional memory, we estimate the fractional integration parameter d for each asset using the Feasible Exact Local Whittle (FELW) method (Shimotsu, 2010). The FELW estimator is a semiparametric frequency domain technique that consistently recovers the memory parameter in $ARFIMA(0,\ d,\ 0)$ without requiring specification of short-run ARMA dynamics, making it robust to model misspecification and well suited to finite samples (Robinson, 1995; Shimotsu, 2010). Daily returns are calculated from price data, and each series is demeaned within rolling windows to remove shifts in mean. Following established practice, we set the bandwidth $m = \left\lfloor T^{0.6} \right\rfloor$, which balances bias and variance in long–memory estimation (Shimotsu, 2010). We employ a 30-day rolling window (approximately six weeks), calibrated so that the mean and median of the market level d are near zero, an empirical hallmark of market neutrality, thus isolating economically meaningful asset-centric memory dynamics.

Interpretation of the memory parameter. It is important to clarify how we interpret the fractional memory parameter d within this framework. Although the FELW estimator provides asymptotically consistent estimates, our economic interpretation does not rely on formal hypothesis testing of d against the null $d = 0$. The role of d in this chapter is structural: even small deviations from zero alter the weighting of autocovariances and higher-order cumulants across horizons, producing economically meaningful shifts in effective variance, skewness accumulation, and sequence-of-returns risk. Consequently, the analysis focuses on the economic magnitude of the implied multi-period adjustments rather than on statistical significance of d in isolation. Rolling-window estimates, cross-sectional dispersion, and stability diagnostics ensure that the variation

in d used for portfolio construction reflects persistent dependence patterns rather than estimation noise.

Memory and mean–reversion link examination. To test whether our long–run memory estimates d_i manifest as observable short–run dynamics, either mean reversion or momentum, and to quantify their forecasting potential, we first apply a set of complementary tests on asset returns obtained from memory sorted quintile portfolios; AR(1) regressions, Augmented Dickey-Fuller (ADF) tests, and rolling window forecasts. We conduct these tests on four series; raw returns, returns adjusted by the risk-free rate, market returns, and both. The AR(1) model evaluates the autoregressive coefficient ϕ in $r_{i,t} = \alpha_i + \phi_i r_{i,t-1} + \varepsilon_{i,t}$ to identify mean-reversion or persistence. The ADF test formally checks stationarity, while rolling-window one-step ahead forecasts defined by $\hat{r}_{i,t+1|t} = \hat{\alpha}_i + \hat{\phi}_i r_{i,t}$ assess the temporal stability of persistence parameters using Mean Squared Error (MSE) and Mean Absolute Error (MAE) as performance measures.

CAPM extension. To gauge how memory affects asset prices, we construct a memory-based High–Minus–Low (HML) factor following the methodology of Fama and French (1993). Each trading day, we estimate the fractional memory parameter $\hat{d}_i$ for every asset in our universe and sort the cross section into five quintiles of equal size (Q1–Q5). We then form equally weighted portfolios for each quintile, where Q1 contains the lowest memory assets, Q5 the highest memory assets, and compute their daily returns $R_{\mathrm{Q1},t}, \ldots, R_{\mathrm{Q5},t}$. The memory HML factor is defined as the daily spread

$$\mathrm{HML}_{m,t} = R_{\mathrm{Q5},t} - R_{\mathrm{Q1},t},$$

which isolates the return premium associated with low versus high memory. This framework serves as the empirical analogue of the linear memory tilt $\kappa(-d_i)$ introduced in Eq. (4.35), which we implement on portfolio level when deriving the efficient frontier via the memory-augmented mean-variance algorithm following Eq. (4.37).

We integrate the derived memory factor into the Capital Asset Pricing Model (CAPM) and empirically evaluate its pricing significance using the GRS test (Gibbons et al., 1989).

$$R_{it} - R_{ft} = \alpha_i + \beta_{iM}(R_{Mt} - R_{ft}) + \epsilon_{it} \tag{4.40}$$

Specifically, we estimate the augmented CAPM:

$$R_{it} - R_{ft} = \alpha_i + \beta_{iM}(R_{Mt} - R_{ft}) + \beta_{iHMLm}\,\mathrm{HML}m,t + \epsilon_{it}, \tag{4.41}$$

where the statistical significance of intercepts α_i are jointly tested using the GRS statistic:

$$\text{GRS} = \frac{(T-N-K)}{N} \frac{\hat{\alpha}^\top \hat{\Sigma}^{-1} \hat{\alpha}}{1 + \bar{f}^\top \hat{F}^{-1} \bar{f}}, \tag{4.42}$$

with T observations, N assets, and K factors (market and HML). Significant alphas indicate that memory characteristics indeed represent a priced dimension of systematic risk, beyond market exposure alone.

Out-of-sample return forecasts. To further validate the significance of memory in the return formation process, we examine the predictive power of CAPM augmented with HML factor, against conventional CAPM. To do this, we divide the entire sample of excess daily portfolio returns (Q1-Q5) and factor series into a 80% training set and a 20% test set. For each quintile portfolio we fit two regressions on conventional and extended CAPM defined by Eq. (4.40) and Eq. (4.41) respectively.

We then generate a set of forecasts on the test data producing a *nowcast* for the same day's return. Forecast accuracy is evaluated for both models and both horizons using Mean Squared Error (MSE), Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). Relative performance between the two models is assessed using:

$$\text{RFE} = \left(\frac{\text{Loss}_{CAPM_{HML}}}{\text{Loss}_{CAPM}} - 1 \right) \times 100, \tag{4.43}$$

where a negative RFE indicates that CAPM extended with memory factor outperforms conventional CAPM.

4.4.2 Theoretical Framework Validation

Portfolio Performance Analysis Building on Section 4.3, we sort assets into quintile portfolios by their cross-sectional memory estimates $\hat{d}_i$ (averaged over the full sample) and examine two rebalancing rules: frictionless daily updates versus a 30-day schedule (matching the estimation window) with a 0.2 % transaction cost per portfolio rebalancing. Risk is measured by the effective variance $\sigma^2_{\text{eff}}(T)$ from Eq. (4.13), capturing serial dependence and long–memory effects. Performance is summarised by annualised return, volatility (both conventional and square root of Eq. (4.13) adjusted by 255 days to convert it to annual standard deviation), Sharpe & Sortino ratios, beta, alpha, Treynor ratio, ROI, VaR, ES, and the average pairwise correlation among holdings to demonstrate diversification benefits within each portfolio.

Efficient–Frontier Diagnostics Beyond portfolio analysis, we compare the *shape* of memory-augmented efficient frontiers defined by Eq. (4.37) with their conventional counterparts from the mean-variance framework proposed by Markowitz (1952). To

ensure that the results reflect true in-sample performance, we calculate the average memory of each asset and then sort assets into quintiles based on cross-sectional increase in memory values. For each 30-day window we first apply the static memory-tilt as defined by EQ. (4.35)

$$\boldsymbol{\mu}^{\text{tilt}} = \boldsymbol{\mu}^{\text{conv}} + \kappa\,(-\mathbf{d}),$$

with κ estimated via cross–sectional OLS. On this tilted return vector we overlay two covariance structures: 1) the full effective covariance $\Sigma_{\text{eff}}^{\text{full}}(T)$ from Eq. (4.37), which retains both own series memory and cross-serial terms; and 2) the diagonal variant $\Sigma_{\text{eff}}^{\text{diag}}(T)$, which preserves only each asset's memory-inflated variance and zeros all off-diagonals.

In each case we solve the memory-augmented algorithm of Eq. (4.38) constrained by long–only criteria, yielding tangency portfolios $\mathbf{w}_t^{\text{full}}$ and $\mathbf{w}_t^{\text{diag}}$ and their frontiers. Following Jobson and Korkie (1981) we apply a paired t–test to the two daily excess return streams to assess the Sharpe ratio uplift. Complementarily, we compute the signed area for each covariance mode (full vs. diagonal)

$$\Delta\text{AUC} = \int_{\sigma_{\min}}^{\sigma_{\max}} \Big[\mu_{\text{mem}}(\sigma) - \mu_{\text{conv}}(\sigma)\Big]\, d\sigma \tag{4.44}$$

using a common volatility grid. This is consistent with DeMiguel et al. (2009) approach. A positive $\Delta\text{AUC}^{\text{mode}}$ indicates that the memory–adjusted frontier (full or diagonal) dominates its conventional counterpart across all attainable risk levels. We repeat these diagnostics for each memory quintile and for the full universe, thereby pinpointing where memory tilt plus full covariance modelling yields the greatest benefit.

To assess investor-relevant improvements along the frontier, we measure the return differential at 20% target volatility level ($\sigma^\star$) using point-wise uplift:

$$\text{Uplift}(\sigma^\star) = \mu_{\text{mem}}(\sigma^\star) - \mu_{\text{conv}}(\sigma^\star) \tag{4.45}$$

where $\mu_{\text{mem}}(\sigma^\star)$ and $\mu_{\text{conv}}(\sigma^\star)$ are the expected returns read off the memory augmented and conventional frontiers at volatility $\sigma^\star$. By fixing risk at $\sigma^\star$ (e.g. 20%), this metric directly quantifies the additional return that memory dynamics contribute at the precise risk level investors choose. Such pointwise diagnostics are standard in volatility-targeting and risk-budgeting contexts (e.g. Moreira and Muir, 2017; Barroso and Santa-Clara, 2015), as they translate global frontier shifts into actionable estimates of return improvement at the risk levels that practitioners actually employ.

4.5 Data

Our empirical analysis utilises daily closing price data for sixty large-capitalisation US equities, each exceeding USD 10 billion in market value and having continuous trading histories spanning January 2004 to November 2024. Daily returns are computed using the standard formulation: $r_t = \frac{P_t - P_{t-1}}{P_{t-1}}$, consistent with common practice in the empirical literature on asset-pricing and risk management (Campbell et al., 1997). Given the small magnitude of daily returns, the approximation $\ln(1 + r_t) \approx r_t$ remains accurate to second-order precision, meaning that memory estimates, CAPM regressions, and risk metrics remain effectively unchanged when using simple versus logarithmic returns.

For CAPM estimations, we additionally include the one-month US Treasury bill rate as the risk-free rate and employ the S&P 500 index returns to proxy for the market portfolio. The full dataset comprises 5,277 trading day observations, which is reduced to 5,247 observations after accounting for the 30-day rolling window required for memory parameter estimation.

Table 4.1 summarises the essential characteristics of asset returns, including mean returns, volatility, skewness, and kurtosis. Complementing this, Table 4.2 provides descriptive statistics for the FELW-estimated fractional memory parameter d, capturing measures of central tendency, dispersion, and the results of Jarque–Bera tests for non-normality. The statistics reveal pronounced cross-sectional variation in the estimated memory parameters, with distinct subsets of assets exhibiting antipersistence ($d < 0$) and others showing long–memory persistence ($d > 0$). Such cross-sectional heterogeneity provides a robust empirical foundation for our investigation into how memory dynamics influence asset returns, effective variance, and risk premia, thus directly informing both the theoretical framework and subsequent portfolio analyses.

4.6 Empirical Analysis

This section presents the empirical results structured to validate the theoretical propositions outlined in Section 4.3. Firstly, portfolio-level summary statistics are detailed, highlighting distinctions across memory-ranked quintiles to verify theoretical predictions on volatility drag and memory dynamics. Secondly, temporal dynamics are assessed through mean-reversion tests, explicitly linking short-term autocorrelation patterns to fractional memory parameters and the theoretical constructs in Section 4.3.5. Finally, we empirically evaluate the extended CAPM incorporating a memory risk premium using the GRS test, thus substantiating the theoretical model of compounding-adjusted risk presented in Section 4.3.6 and Section 4.3.7.

TABLE 4.1 : Asset Universe Returns - Descriptive Statistics.

Ticker	Mean	Median	Max	Min	Std. Dev.	Skew	Kurt	Jarque-Bera
GSPC	0.0002	0.0007	0.1158	- 0.1198	0.0137	- 0.2403	13.526	15,733
RFR	0.0001	0.0000	0.0002	- 0.0000	0.0001	0.9700	2.473	573
AAPL	0.0015	0.0013	0.1390	- 0.1792	0.0217	0.1315	8.130	3,740
ABT	0.0005	0.0006	0.1094	- 0.0979	0.0149	- 0.0752	8.788	4,751
ADBE	0.0008	0.0011	0.1772	- 0.1903	0.0228	0.0203	10.079	7,101
AEP	0.0002	0.0009	0.1321	- 0.1158	0.0145	0.1234	12.270	12,186
AMD	0.0008	0.0009	0.2238	- 0.1686	0.0368	0.0678	6.038	1,311
AMZN	0.0008	0.0006	0.2695	- 0.2182	0.0254	0.6194	14.899	20,281
APD	0.0004	0.0007	0.1465	- 0.1555	0.0179	- 0.3758	11.182	9,566
BA	0.0004	0.0004	0.2432	- 0.2385	0.0240	0.2350	18.853	35,646
BAC	- 0.0000	0.0003	0.3098	- 0.2897	0.0315	0.5603	24.819	67,638
BMY	0.0002	0.0004	0.1144	- 0.1599	0.0162	- 0.2615	12.184	11,990
C01	- 0.0001	- 0.0003	0.5782	- 0.2641	0.0338	1.8753	46.885	274,906
CAT	0.0006	0.0009	0.1472	- 0.1428	0.0215	- 0.0072	7.356	2,689
CL	0.0002	0.0005	0.1050	- 0.0978	0.0128	0.0422	10.521	8,016
CMCSA	0.0003	0.0005	0.2452	- 0.1448	0.0193	0.3429	16.659	26,506
COP	0.0004	0.0007	0.2521	- 0.2484	0.0239	- 0.0319	15.031	20,512
CSCO	0.0002	0.0004	0.1595	- 0.1621	0.0196	- 0.0386	13.589	15,890
CVX	0.0004	0.0008	0.2274	- 0.2212	0.0199	0.1249	22.705	55,033
D01	0.0002	0.0005	0.1057	- 0.1231	0.0143	- 0.1529	12.868	13,812
DD	0.0003	0.0004	0.1843	- 0.1903	0.0233	0.0453	10.894	8,831
DUK	0.0004	0.0009	0.1331	- 0.1150	0.0136	0.0639	14.644	19,217
EMN	0.0004	0.0004	0.1219	- 0.1360	0.0218	- 0.2007	7.780	3,261
F	0.0002	0.0006	0.2512	- 0.2500	0.0282	0.0662	15.902	23,592
FDX	0.0004	0.0005	0.1553	- 0.2140	0.0214	- 0.4822	12.649	13,326
GILD	0.0004	0.0005	0.1382	- 0.1434	0.0193	- 0.0788	9.184	5,423
GIS	0.0002	0.0005	0.0943	- 0.1141	0.0128	- 0.4978	10.827	8,823
HAL	0.0002	0.0004	0.2653	- 0.3764	0.0300	- 0.5222	18.089	32,419
HD	0.0005	0.0007	0.1407	- 0.1979	0.0179	- 0.0453	13.313	15,072
HON	0.0003	0.0007	0.1507	- 0.1209	0.0175	0.0313	9.589	6,153
HPQ	0.0002	0.0010	0.1710	- 0.2003	0.0223	- 0.2675	11.533	10,358
IBM	0.0003	0.0006	0.1152	- 0.1285	0.0155	- 0.2245	10.790	8,627
INTC	0.0000	0.0008	0.1952	- 0.2606	0.0225	- 0.5244	14.544	19,042
JNJ	0.0002	0.0003	0.1223	- 0.1004	0.0118	0.2142	13.774	16,475
JPM	0.0004	0.0006	0.2510	- 0.2073	0.0252	0.8433	20.402	43,316
KMB	0.0002	0.0004	0.0970	- 0.0883	0.0129	- 0.1807	10.444	7,872
KO	0.0002	0.0007	0.1388	- 0.0967	0.0126	0.0055	15.222	21,167
LMT	0.0005	0.0006	0.1065	- 0.1276	0.0152	- 0.1386	11.553	10,378
LOW	0.0008	0.0006	0.1650	- 0.2477	0.0203	- 0.1509	15.004	20,433
MCD	0.0005	0.0008	0.1813	- 0.1588	0.0140	0.3850	21.162	46,826
MDT	0.0002	0.0005	0.1030	- 0.1282	0.0159	- 0.1481	9.824	6,612
MET	0.0003	0.0008	0.2800	- 0.2677	0.0282	0.4261	22.733	55,283
MMM	0.0003	0.0007	0.2299	- 0.1295	0.0167	0.5370	19.371	38,143
MS	0.0002	0.0006	0.8698	- 0.2589	0.0337	5.2052	144.988	2,872,264
MSFT	0.0005	0.0004	0.1860	- 0.1474	0.0186	0.1909	12.181	11,965
ORCL	0.0005	0.0008	0.2043	- 0.1350	0.0194	0.5204	13.658	16,251
OXY	0.0006	0.0005	0.3370	- 0.5201	0.0297	- 0.7419	41.853	214,228
PEP	0.0002	0.0003	0.1050	- 0.1193	0.0123	- 0.5187	17.348	29,326
PFE	- 0.0000	- 0.0005	0.1086	- 0.1062	0.0158	0.1371	8.819	4,809
PG	0.0003	0.0006	0.1201	- 0.0874	0.0123	0.0217	12.522	12,848
PPG	0.0002	0.0008	0.1483	- 0.1250	0.0186	- 0.0714	9.307	5,641
PPL	0.0001	0.0009	0.1480	- 0.1371	0.0158	- 0.1179	15.466	22,028
SHW	0.0005	0.0006	0.1526	- 0.1868	0.0185	- 0.2973	14.823	19,860
SLB	0.0003	0.0001	0.1991	- 0.2742	0.0265	- 0.2899	12.085	11,745
SO	0.0004	0.0006	0.1344	- 0.1177	0.0132	0.0839	16.617	26,279
T	0.0005	0.0009	0.1628	- 0.1041	0.0156	0.3273	13.144	14,642
TGT	0.0002	0.0002	0.2043	- 0.2493	0.0207	0.0815	19.650	39,287
TXN	0.0006	0.0011	0.1193	- 0.1183	0.0196	- 0.1065	5.771	1,095
VZ	0.0004	0.0007	0.1463	- 0.0807	0.0141	0.4148	11.391	10,076
WFC	0.0002	0.0002	0.3276	- 0.2382	0.0270	0.9335	22.956	56,929
WMT	0.0003	0.0007	0.1107	- 0.1138	0.0133	0.0828	14.084	17,413
XOM	0.0004	0.0004	0.1719	- 0.1395	0.0185	0.2275	12.570	13,008

Note: This table presents the descriptive statistics of the asset universe and the S&P 500 index benchmark returns. It also includes statistics for the Risk-Free Rate (RFR) and GSPC index selected to represent the overall market performance. The period covered spans from 13-01-2004, to 14-11-2024, with a total of 5,247 observations.

TABLE 4.2 : Asset Universe Memory - Descriptive Statistics.

Ticker	Mean	Median	Max.	Min.	Std. Dev.	Skew.	Kurt.	Jarque-Bera
GSPC	0.039	0.001	0.116	-0.120	0.012	-0.263	16.000	37,010.0
AAPL	0.017	0.047	1.499	-1.645	0.373	-0.272	4.046	303.59
ABT	0.023	0.030	1.642	-1.452	0.372	-0.087	3.998	224.52
ADBE	-0.001	0.013	2.182	-1.672	0.387	-0.083	3.911	187.46
AEP	-0.008	0.015	2.906	-1.683	0.382	-0.220	4.716	685.74
AMD	0.057	0.071	1.962	-1.376	0.379	-0.046	3.934	192.75
AMZN	0.018	0.040	2.083	-1.534	0.356	-0.271	4.327	449.24
APD	0.011	0.013	2.103	-1.529	0.384	-0.018	4.036	234.95
BA	0.048	0.055	2.012	-1.601	0.360	-0.112	4.456	474.71
BAC	0.079	0.085	1.968	-1.599	0.365	-0.199	4.461	501.11
BMY	-0.020	0.013	1.532	-1.615	0.383	-0.352	3.667	205.61
C01	0.072	0.085	1.626	-1.593	0.345	-0.080	4.433	454.27
CAT	0.062	0.065	1.540	-1.627	0.335	-0.046	4.268	353.27
CL	-0.027	-0.017	1.853	-1.836	0.381	-0.238	4.180	353.77
CMCSA	0.020	0.035	1.717	-1.578	0.355	-0.137	4.078	270.38
COP	0.029	0.040	1.969	-1.442	0.360	-0.035	4.287	363.28
CSCO	0.030	0.038	1.500	-1.758	0.350	-0.207	4.255	381.65
CVX	0.016	0.029	1.989	-1.500	0.377	-0.113	4.214	333.17
D01	-0.004	0.019	1.564	-1.468	0.399	-0.255	3.573	128.57
DD	0.048	0.058	1.607	-1.891	0.355	-0.278	4.353	467.65
DUK	-0.009	-0.008	2.051	-1.458	0.412	-0.032	3.474	50.06
EMN	0.033	0.046	2.225	-1.380	0.393	-0.230	3.774	177.02
F	0.070	0.084	2.062	-1.285	0.349	-0.148	3.893	193.54
FDX	-0.000	0.013	1.689	-1.788	0.375	-0.250	3.954	253.48
GILD	0.006	0.015	1.777	-1.591	0.373	0.096	4.234	340.77
GIS	-0.015	0.001	1.824	-1.578	0.345	-0.246	4.158	345.88
HAL	0.032	0.049	1.865	-1.655	0.370	-0.280	3.922	254.18
HD	0.039	0.037	2.132	-1.831	0.382	0.036	3.748	123.50
HON	0.023	0.038	1.802	-1.617	0.378	-0.137	3.755	141.09
HPQ	0.017	0.038	1.605	-1.547	0.360	-0.228	3.786	180.36
IBM	0.065	0.068	1.945	-1.747	0.360	-0.108	4.362	415.72
INTC	0.060	0.065	1.677	-1.399	0.357	-0.210	3.595	116.18
JNJ	0.019	0.040	1.756	-1.895	0.392	-0.206	3.747	159.00
JPM	0.052	0.062	2.223	-1.517	0.372	-0.007	5.140	1,001.4
KMB	-0.013	0.006	1.686	-1.544	0.384	-0.089	3.486	58.59
KO	0.028	0.035	2.373	-1.528	0.389	-0.001	4.507	496.45
LMT	0.016	0.032	1.730	-1.509	0.381	-0.106	3.725	124.79
LOW	0.005	0.026	1.744	-1.510	0.371	-0.257	4.203	374.19
MCD	0.000	0.004	1.957	-1.383	0.357	0.080	3.904	184.37
MDT	0.003	0.020	1.593	-1.868	0.383	-0.323	4.177	394.04
MET	0.015	0.023	2.389	-1.635	0.376	-0.021	3.830	151.08
MMM	0.027	0.031	1.913	-1.453	0.374	0.029	4.613	569.85
MS	0.036	0.049	2.597	-1.768	0.387	-0.094	5.248	1,112.3
MSFT	0.011	0.012	1.749	-1.762	0.381	-0.071	4.156	296.63
ORCL	0.017	0.039	1.810	-1.426	0.358	-0.411	4.000	366.56
OXY	-0.017	0.010	2.099	-1.701	0.384	-0.273	4.082	321.54
PEP	0.003	0.009	1.811	-1.834	0.365	0.030	4.228	330.35
PFE	0.055	0.049	1.823	-1.464	0.368	0.146	4.457	482.73
PG	-0.034	-0.021	1.873	-1.502	0.388	-0.178	3.817	173.71
PPG	0.014	0.021	2.248	-1.495	0.366	-0.143	4.307	391.56
PPL	0.015	0.017	1.944	-1.635	0.412	-0.099	3.745	129.97
SHW	-0.053	-0.038	1.513	-1.577	0.376	-0.262	3.524	120.15
SLB	0.047	0.062	1.910	-1.615	0.349	-0.097	4.207	326.65
SO	-0.046	-0.045	1.916	-1.732	0.404	0.021	4.248	340.70
T	0.038	0.058	1.732	-1.659	0.384	-0.243	4.040	288.09
TGT	-0.033	-0.007	1.931	-1.828	0.393	-0.275	3.799	205.53
TXN	-0.031	-0.017	1.568	-1.911	0.356	-0.272	3.947	260.56
VZ	0.044	0.072	2.332	-1.734	0.390	-0.256	4.166	354.55
WFC	-0.006	0.005	1.720	-1.822	0.397	-0.128	4.008	236.29
WMT	0.019	0.016	2.276	-1.573	0.390	-0.081	3.847	162.53
XOM	-0.013	-0.000	1.576	-1.677	0.364	-0.304	4.216	403.64

Note: This table presents the descriptive statistics of the memory of individual assets in the asset universe, estimated using the return series of the corresponding assets. It also includes statistics for the GSPC index selected to represent the overall market performance. The period covered spans from 13-01-2004 to 14-11-2024, with a total of 5,247 observations.

4.6.1 Portfolio Level Statistics Overview

Table 4.3 reports the key return distribution moments for five memory-sorted portfolios. Consistent with the theoretical expectation from Section 4.3.5, higher memory quintiles (Q5) demonstrate significantly greater realised volatility and tail risk compared to the antipersistent quintiles (Q1), which exhibit lower volatility and kurtosis. This observation validates our theoretical proposition that antipersistence attenuates effective variance and tail risk, whilst persistence amplifies them.

TABLE 4.3 : Portfolio Returns - Descriptive Statistics.

Portfolio	Mean	Median	Max.	Min.	Std. Dev.	Skew.	Kurt.	Jarque-Bera
Q1	0.054	0.0007	0.129	-0.098	0.012	0.298	15.21	32,668.5
Q2	0.066	0.0008	0.110	-0.129	0.012	-0.008	14.43	28,580.6
Q3	0.055	0.0010	0.111	-0.129	0.013	-0.295	16.69	41,025.2
Q4	0.050	0.0008	0.144	-0.118	0.013	0.021	19.28	57,930.7
Q5	0.048	0.0008	0.197	-0.132	0.013	0.356	26.30	118,832.1
MKT	0.039	0.0007	0.116	-0.120	0.012	-0.263	16.00	37,010.2

Note: This table presents the descriptive statistics of the returns for the quintile portfolios and the market. The quintile portfolios are constructed based on daily variations of memory parameters estimated using a 30-observation rolling window applied to the corresponding asset return series. The Q1 portfolio comprises assets with the lowest memory values, whereas Q5 includes those with the highest memory values on the corresponding day. The data spans from 13-01-2004 to 14-11-2024, forming a total of 5,247 observations.

Memory parameter statistics presented in Table 4.4 confirm that Q1 assets exhibit strong antipersistence, with mean $d \approx -0.47$, whilst the Q5 assets demonstrate pronounced persistence mean $d \approx 0.48$. The significant dispersion and non-normality of these estimates justify our quintile-based sorting approach, directly supporting out theoretical discussions in Section 4.3.4 and the subsequent empirical testing of memory risk.

TABLE 4.4 : Portfolio Memory - Descriptive Statistics.

Portfolio	Mean	Median	Max.	Min.	Std. Dev.	Skew.	Kurt.	Jarque-Bera
Q1	-0.473	-0.469	0.203	-1.338	0.172	0.022	3.527	61.2
Q2	-0.140	-0.141	0.509	-0.863	0.141	-0.007	4.044	238.4
Q3	0.030	0.027	0.702	-0.606	0.133	0.109	4.292	375.1
Q4	0.189	0.184	0.860	-0.347	0.138	0.307	4.132	362.6
Q5	0.477	0.474	1.650	-0.165	0.165	0.282	4.323	452.6
MKT	-0.002	-0.002	2.353	-1.569	0.378	-0.026	4.419	440.7

Note: This table presents the descriptive statistics for portfolio-level memory values for the quintile portfolios and the market. The quintile portfolios are constructed based on daily variations in the estimated memory parameter using a 30-observation rolling window. The Q1 portfolio comprises assets with the lowest memory values, while Q5 includes those with the highest memory values on the corresponding day. The data spans from 13-01-2004 to 14-11-2024, forming a total of 5,247 observations.

4.6.2 Mean-Reversion Analysis

To test whether short-run autocorrelation ϕ_i mirrors the long-run memory dynamics captured by d_i, we estimate AR(1) models on four return series of our quintile portfolios; raw, excess (net of R_f) (RFR), market adjusted (MKT), as well as idiosyncratic when both RFR and MKT effects are removed (Table 4.5).

Under the null $H_0 : \phi_i = 0$, raw and excess returns reject H_0 in every quintile ($\hat{\phi} \approx -0.072$ to -0.111, $p < 0.01$), confirming strong short-run mean-reversion consistent with antipersistence. After controlling for market effects or both market and risk-free (ALL), Q2, Q4, and Q5 exhibit insignificant $\hat{\phi} \approx 0$, implying their daily returns show no net short-run drift despite markedly non-zero long-run memory d. In contrast, once market influences are removed, the AR(1) coefficients for Q1 and Q3 flip from $\hat{\phi} \approx -0.072$ & -0.083 to 0.037 & 0.071 respectively (both $p < 0.01$). This suggests that part of their raw mean-reversion was driven by market reversals, whereas idiosyncratically both exhibit mild momentum. These patterns emphasise that long-run memory d captures the decay of autocorrelations over many lags, while short-run dynamics captured by $\hat{\phi}$ can differ materially once systematic effects are removed.

Conducted ADF tests reject the unit-root null for all quintiles under every specification (ADF $t \ll -11$, $p < 0.01$), confirming stationarity of all return series. Finally, 30-day rolling one-step forecasts yield the lowest errors for antipersistent Q1 (MSE = 0.00015, MAE = 0.0082) and the highest for persistent Q5 (MSE = 0.00019, MAE = 0.0086), demonstrating superior short-term predictability in the presence of strong antipersistence. After removing market (and risk-free) effects, these error rates fall further, to MSE = 0.00003, MAE = 0.0039 for Q1 and MSE = 0.00004, MAE = 0.0041 for Q5, highlighting that pure idiosyncratic memory signals yield even sharper one-step forecasts.

The AR(1) results confirm that short-run autocorrelations (ϕ_i) align with our long-run memory estimates (d_i): strongly antipersistent portfolios (Q1, Q2) exhibit negative ϕ_i (mean-reversion), while strongly persistent portfolios (Q3-Q5) show neutral or positive ϕ_i (momentum or no drift) once market and risk-free influences are removed. ADF tests uniformly reject non-stationarity, ensuring these serial-dependence patterns are stable. Rolling one-step forecasts deliver the largest error reductions in Q1 (up to 44% MSE drop) and Q5 (up to 40% MSE drop), with only marginal gains for mid-memory portfolios. Together, these findings validate that our estimated d_i parameters capture real, economically meaningful short-run dynamics and confer tangible predictive advantages at both ends of the memory spectrum.

TABLE 4.5 : Mean Reversion Analysis Results.

		AR(1)				ADF		Rolling Forecast	
Portf	**Adjustment**	ϕ	**S.E.**	**Tval**	**Pval**	**Tstat**	**Pval**	**MSE**	**MAE**
Q1	Raw	-0.072	0.014	-5.243	0.00	-17.018	0.00	0.00015	0.0082
Q2	Raw	-0.111	0.014	-8.054	0.00	-16.432	0.00	0.00016	0.0084
Q3	Raw	-0.083	0.014	-6.004	0.00	-16.203	0.00	0.00017	0.0085
Q4	Raw	-0.096	0.014	-6.961	0.00	-13.932	0.00	0.00018	0.0087
Q5	Raw	-0.091	0.014	-6.651	0.00	-13.852	0.00	0.00019	0.0086
Q1	RFR	-0.072	0.014	-5.242	0.00	-17.022	0.00	0.00015	0.0082
Q2	RFR	-0.110	0.014	-8.047	0.00	-16.412	0.00	0.00016	0.0084
Q3	RFR	-0.083	0.014	-6.005	0.00	-16.209	0.00	0.00017	0.0085
Q4	RFR	-0.096	0.014	-6.950	0.00	-13.904	0.00	0.00018	0.0087
Q5	RFR	-0.091	0.014	-6.636	0.00	-13.790	0.00	0.00019	0.0086
Q1	MKT	0.037	0.014	2.676	0.01	-42.360	0.00	0.00003	0.0039
Q2	MKT	-0.011	0.014	-0.817	0.41	-14.386	0.00	0.00003	0.0037
Q3	MKT	0.071	0.014	5.183	0.00	-18.620	0.00	0.00003	0.0036
Q4	MKT	0.012	0.014	0.894	0.37	-11.999	0.00	0.00003	0.0038
Q5	MKT	-0.013	0.014	-0.953	0.34	-16.276	0.00	0.00004	0.0041
Q1	ALL	0.037	0.014	2.676	0.01	-42.360	0.00	0.00003	0.0039
Q2	ALL	-0.011	0.014	-0.814	0.42	-14.380	0.00	0.00003	0.0037
Q3	ALL	0.071	0.014	5.182	0.00	-18.621	0.00	0.00003	0.0036
Q4	ALL	0.012	0.014	0.893	0.37	-12.000	0.00	0.00003	0.0038
Q5	ALL	-0.013	0.014	-0.953	0.34	-16.276	0.00	0.00004	0.0041

Note: This table presents the results of our mean reversion analysis across quintile portfolios, using three adjustment approaches (Raw, RFR, MKT, and ALL (RFR+MKT)). For each portfolio, the table reports the AR(1) coefficient (ϕ), its standard error, t-statistic, and p-value; the ADF test t-statistic and p-value; as well as rolling window forecast performance metrics (MSE and MAE). The quintile portfolios (Q1 to Q5) are sorted by the estimated memory parameter, with Q1 representing assets with the lowest memory (strong anti-persistence) and Q5 those with the highest memory (strong persistence). The analysis is based on 5,247 observations spanning the study period from 13-01-2004 to 14-11-2024.

4.6.3 Memory-Factor CAPM and GRS Test Results

The empirical validity of the memory-based HML factor is tested using the GRS framework (Table 4.6). Consistent with theoretical predictions, low-memory portfolios (Q1, Q2) exhibit insignificant alphas over the full sample. Here Q1 t-statistic of 0.51 ($p = 0.77$) and Q2 of 1.93 ($p = 0.09$) imply that antipersistence is fully captured by the memory factor. Conversely, portfolios with higher memory (Q3–Q5) exhibit significant alphas (t-statistics 5.79, 4.92 and 3.44 respectively; $p \approx 0$), indicating residual abnormal returns are not explained by the additional factor, reinforcing the notion that persistence needs to be explicitly accounted for within asset-pricing models.

Subsample analyses covering periods of market stress (2008–10) indicate intensified pricing of memory risk, particularly in Q3 and Q4. These findings suggest that memory dynamics interact significantly with market conditions to influence volatility and portfolio performance. In an out-of-sample split, the training period shows robust significance for Q3 – Q5, while the test period continues to register marginal significance for Q3 ($t = 1.87$, $p = 0.10$).

In all tests, Q1 maintains high average returns without abnormal returns, reinforcing that strong antipersistence is predictable and adequately priced.

TABLE 4.6 : HML CAPM Extension - GRS Test Results.

Factor	**HML**	**Q1**		**Q2**		**Q3**		**Q4**		**Q5**	
Period	**Interval**	**tStat**	**Pval**	**tStat**	**Pval**	**tStat**	**Pval**	**tStat**	**Pval**	**tStat**	**Pval**
Full	2004-01-13 to 2024-11-14	0.512	0.77	1.931	0.09	5.795	0.00	4.925	0.00	3.436	0.00
1	2004-01-13 to 2006-01-12	0.612	0.69	2.544	0.03	4.113	0.00	0.877	0.50	0.395	0.85
2	2006-01-13 to 2008-01-12	1.619	0.15	2.961	0.01	4.680	0.00	1.316	0.26	0.762	0.58
3	2008-01-13 to 2010-01-12	0.468	0.80	1.516	0.18	5.227	0.00	6.290	0.00	3.530	0.00
4	2010-01-13 to 2012-01-12	0.860	0.51	1.812	0.11	1.405	0.22	0.295	0.92	0.375	0.87
5	2012-01-13 to 2014-01-12	3.011	0.01	3.646	0.00	1.680	0.14	0.777	0.57	1.322	0.25
6	2014-01-13 to 2016-01-12	0.798	0.55	1.027	0.40	0.739	0.59	0.812	0.54	0.586	0.71
7	2016-01-13 to 2018-01-12	1.246	0.29	1.529	0.18	2.556	0.03	1.547	0.17	2.950	0.01
8	2018-01-13 to 2020-01-12	1.194	0.31	1.878	0.10	2.212	0.05	0.623	0.68	0.191	0.97
9	2020-01-13 to 2022-01-12	1.354	0.24	0.753	0.58	0.117	0.99	0.533	0.75	0.881	0.49
10	2022-01-13 to 2024-01-12	0.176	0.97	0.290	0.92	0.895	0.48	1.273	0.27	0.800	0.55
11	2024-01-13 to 2024-11-14	0.431	0.83	0.361	0.87	3.914	0.00	2.001	0.08	1.531	0.18
Train	2004-01-13 to 2020-09-01	0.233	0.95	2.794	0.02	9.619	0.00	6.039	0.00	3.736	0.00
Test	2020-09-01 to 2024-11-14	0.444	0.82	0.589	0.71	1.867	0.10	1.512	0.18	0.498	0.78

Note: This table presents the GRS test results for the extended CAPM model that incorporates the memory-based High-Minus-Low factor. For each time interval (full sample, subsamples, and out-of-sample), the table reports the t-statistics and corresponding p-values for the intercepts (alphas) across the five quintile portfolios (Q1–Q5). The sample period spans from 2004-01-13 to 2024-11-14, with subsample analyses and an 80/20 train/test split also considered.

4.6.4 CAPM Forecasting Performance.

Table 4.7 presents out-of-sample forecasting errors for both the conventional CAPM and its memory augmented HML extension across the five memory-sorted portfolios (Q1–Q5). We evaluate prediction accuracy on the 20 % holdout sample using MSE, MAE and RMSE. For each metric, the extended CAPM delivers uniformly lower loss for the same day forecasts, as evidenced by the negative relative performance figures reported in the RFP section of the table.

TABLE 4.7 : Conventional vs Extended CAPM Out-Of-Sample Forecast Performance.

CAPM	Metric	Q1	Q2	Q3	Q4	Q5
Extended	MSE	0.00003	0.00004	0.00004	0.00004	0.00003
	MAE	0.00402	0.00480	0.00500	0.00501	0.00402
	RMSE	0.00532	0.00633	0.00665	0.00669	0.00532
Standard	MSE	0.00005	0.00004	0.00004	0.00005	0.00005
	MAE	0.00503	0.00483	0.00499	0.00506	0.00507
	RMSE	0.00712	0.00642	0.00663	0.00675	0.00685
RFP	MSE	–44.07 %	–2.90 %	0.86 %	–1.79 %	–39.57 %
	MAE	–20.18 %	–0.73 %	0.29 %	–0.99 %	–20.84 %
	RMSE	–25.21 %	–1.46 %	0.43 %	–0.90 %	–22.26 %

Note: This Table compares out-of-sample forecast accuracy of the extended CAPM (including the memory-based HML factor) against the standard CAPM for five memory-sorted portfolios (Q1–Q5). Forecast errors are measured by mean squared error (MSE), mean absolute error (MAE) and root mean squared error (RMSE) on a 20% hold-out sample. Relative forecast performance (RFP) is computed as the percentage change in each loss metric for the extended model versus the standard model (negative values indicate improvement).

Specifically, the strongest improvements appear at the extremes of the memory spectrum. In Q1 (high antipersistence), the extended model reduces MSE by 44 %, MAE by 20 % and RMSE by 25 %, while in Q5 (strong persistence), the corresponding gains are 40 %, 21 % and 22 % respectively. In contrast, mid-memory portfolios (Q2–Q4) see only marginal changes (±1–2 %). These results corroborate our theoretical prediction that the memory captures the dominant serial dependence signals mean-reversion in antipersistent regimes and persistence in high-memory assets, thereby materially enhancing simultaneous return forecasts where long–memory effects are most pronounced, but adding little in regimes of moderate memory.

4.6.5 Portfolio Performance Analysis

Table 4.8 and Figure 4.2 demonstrate portfolio performance results based on a frictionless daily rebalance strategy. They show that memory-aware strategies systematically reshape the risk-return landscape in ways directly anticipated by our theoretical framework. First, the gap between raw and effective volatility precisely quantifies the drag reducing effect of serial dependence (cf. Eq. (4.13)).

In the frictionless daily rebalancing regime, Q2 (moderate persistence) achieves the largest proportional reduction in volatility (19.23% → 15.90%), translating into Sharpe ratio of 0.792, far above 0.440 in the market yet with lower risk expressed through Beta and tail exposures (VaR, ES). This demonstrates the practical opportunity to capture persistence without incurring excessive compounding drag, as predicted by Eq. (4.13).

TABLE 4.8 : Portfolio Performance Metrics - Daily Rebalanced Frictionless Equally Weighted Portfolios.

Portfolio	Q1	Q2	Q3	Q4	Q5	Market
Return	12.42%	16.03%	12.47%	11.08%	10.39%	8.34%
Volatility	18.82%	19.23%	19.88%	20.74%	21.02%	18.90%
Eff Volatility	16.07%	15.90%	17.61%	18.55%	17.18%	15.35%
Sharpe Ratio	0.638	0.792	0.616	0.539	0.503	0.440
Sortino Ratio	0.635	0.777	0.588	0.524	0.487	0.415
Beta	0.885	0.925	0.964	1.005	0.999	1.000
Alpha	4.86%	8.19%	4.37%	2.71%	2.05%	–
Alpha p-value	0.011	0.000	0.014	0.122	0.260	–
Treynor Ratio	0.136	0.165	0.127	0.111	0.106	0.083
ROI	1,044%	2,109%	1,055%	792%	683%	431%
VaR (5%)	-0.0168	-0.0171	-0.0180	-0.0187	-0.0181	-0.0175
ES (5%)	-0.0272	-0.0280	-0.0298	-0.0304	-0.0314	-0.0292
Min Correlation	0.3275	0.3280	0.3349	0.3331	0.3310	–
Max Correlation	0.4217	0.4144	0.4305	0.4480	0.4450	–
Avg Correlation	0.3783	0.3787	0.3871	0.3933	0.3897	–

Note: This table presents performance metrics for daily-rebalanced, frictionless, equally-weighted portfolios formed by quintile sorting on estimated memory parameters. 'Return' and 'Volatility' are annualised; 'Sharpe', 'Sortino' and 'Treynor' ratios assume a zero risk-free rate; 'Alpha' and its p-value are obtained from the memory-augmented CAPM (Eq. (4.41)); 'ROI' is cumulative return relative to inception; 'VaR' and 'ES' denote the 5% one-day Value-at-Risk and Expected Shortfall, respectively; correlation measures report the minimum, maximum and average pairwise correlations among portfolio constituents. 'Eff Volatility' is the annualised effective standard deviation for a 30-trading-day buy-and-hold horizon. It is computed from the full-sample daily variance σ^2 and the full-sample lag-k autocorrelations $\rho(k)$ by applying Eq. (4.13). All moments are estimated from the entire sample span from 13-01-2004 to 14-11-2024.

In contrast, strongest antipersistence Q1 converts negative autocorrelations into a modest but reliable volatility saving (18.82% $\rightarrow$ 16.07%) and a significant CAPM alpha (4.86%, p=0.011). This suggest that $\rho(k) < 0$ contributes to attenuating drag and enhances expected log growth (see Section 4.3.3) even beyond the memory range of -0.5 and 0.5. High memory quintiles (Q3–Q5) also incur smaller relative volatility reductions (e.g. Q5: 21.02% $\rightarrow$ 17.18%) and their Sharpe ratios remain above market levels, in contrast to the expectation that strong persistence amplifies variance and can overwhelm any tilt in mean (cf. Eq. (4.35)). Perhaps, synchronisation of assets based on their persistence profiles compensates for the increased variance penalty, thus eliminating the drag imposed by positive autocorrelation. However, in light of insignificant p-values for higher persistence portfolios (Q4 & Q5), these results should be treated with caution.

Introducing 30-day rebalancing with a transaction cost of 0.2% per rebalance (Table 4.9, Figure 4.3) attenuates all Sharpe ratios, yet memory ordering persists. Intermediate persistence Q4 now leads with a 0.678 Sharpe and an effective volatility of 18.23%, along with a highly significant alpha (7.14%, p=0.001).

The rebalancing regime exposes a new opportunity: intermediate memory levels strike the optimal trade-off between drag reduction and turnover cost, validating the dual role of d in both the numerator and denominator of our optimiser (Eq. (4.38)). These

FIGURE 4.2 : Cumulative Returns of Frictionless, Equally Weighted Portfolios under Daily Rebalancing.

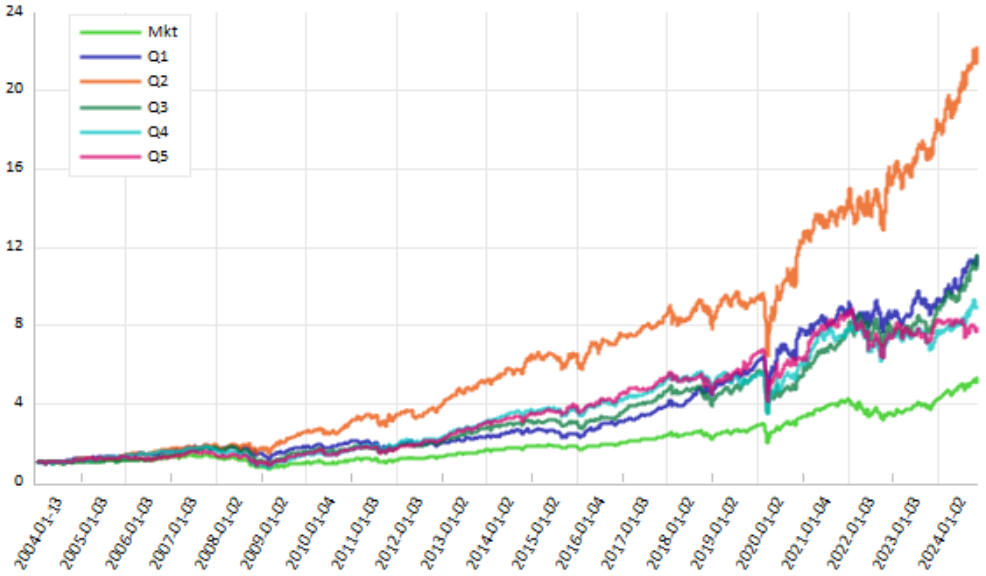

Note: This figure plots the cumulative returns of five quantile portfolios (Q1–Q5) and the market benchmark (Mkt) from 2004-01-13 to 2024-11-14. Each portfolio is equally weighted and rebalanced daily, assuming no transaction costs. Q1 represents the lowest memory assets, while Q5 represents the highest memory assets. The y-axis shows cumulative returns, normalised to 1.0 at the start date, and the x-axis indicates the timeline in daily increments.

findings parallel the dynamic trading rules of Gârleanu and Pedersen (2013), who show that signal autocorrelation, and by extension memory, directly shapes optimal turnover in the presence of transaction costs.

4.6.5.1 Intermediate-Memory Portfolios Dominance

A natural expectation in sorted-portfolio tests is that the extreme portfolios, those containing the most antipersistent (Q1) or most persistent (Q5) assets, should exhibit the strongest performance. In our empirical results, however, for both the US and Japan (for Japan market analysis refer to Section 4.7) the highest risk-adjusted returns arise in the intermediate memory portfolios (Q2 for the US and Q3 for Japan). This pattern is neither anomalous nor inconsistent with the theoretical framework. Instead, it reflects the structure of the estimated memory parameters and the economics of multi-period compounding.

Figure 4.4 plots the average estimated memory parameter d and corresponding average returns across the full US equity universe. Two features stand out. First, the extreme portfolios draw heavily from the tails of the distribution, where d frequently lies outside of stable memory range $|d| < 0.5$. In these ranges, ARFIMA-based estimates are known to exhibit higher sampling variability and less stable finite sample behaviour, making

TABLE 4.9 : Portfolio Performance Metrics - Equally Weighted Portfolio with Transaction Cost & 30 Day Rebalancing Frequency.

Portfolio	Q1	Q2	Q3	Q4	Q5	Market
Return	10.24%	12.13%	8.85%	15.20%	13.26%	8.34%
Volatility	18.76%	18.95%	19.07%	20.22%	20.28%	18.90%
Eff Volatility	16.88%	15.94%	17.26%	18.23%	16.97%	15.35%
Sharpe Ratio	0.535	0.621	0.463	0.728	0.642	0.440
Sortino Ratio	0.527	0.621	0.452	0.730	0.650	0.415
Beta	0.878	0.903	0.918	0.958	0.951	1.000
Alpha	2.73%	4.45%	1.07%	7.14%	5.25%	–
Alpha p-value	0.129	0.015	0.450	0.001	0.012	–
Treynor Ratio	0.114	0.130	0.096	0.154	0.137	0.083
ROI	661%	984%	484%	1,803%	1,237%	431%
VaR (5%)	-0.0167	-0.0171	-0.0169	-0.0176	-0.0176	-0.0175
ES (5%)	-0.0277	-0.0276	-0.0284	-0.0298	-0.0297	-0.0292
Min Correlation	0.3091	0.3116	0.2835	0.3260	0.2932	–
Max Correlation	0.4646	0.4742	0.4719	0.4662	0.4463	–
Avg Correlation	0.3904	0.3721	0.3768	0.3939	0.3722	–

Note: This table presents performance metrics for equally weighted portfolios rebalanced every 30 trading days with a transaction cost of 0.2%. 'Return' and 'Volatility' are annualised; 'Sharpe', 'Sortino' and 'Treynor' ratios assume a zero risk-free rate; 'Alpha' and its p-value are from the memory-augmented CAPM (Eq. (4.41)); 'ROI' is cumulative return since inception; 'VaR' and 'ES' denote the one-day 5% Value-at-Risk and Expected Shortfall; correlation measures report the minimum, maximum and average pairwise correlations among portfolio constituents. 'Eff Volatility' is the annualised effective standard deviation for a 30-trading-day buy-and-hold horizon. It is computed from the full-sample daily variance σ^2 and the full-sample lag-k autocorrelations $\rho(k)$ by applying Eq. (4.13). All moments are estimated from the entire sample span from 13-01-2004 to 14-11-2024.

both inference and portfolio construction more sensitive to noise. Second, the empirical return-memory relationship peaks in the interior of the distribution rather than at the extremes: the assets with moderate values of d (roughly $-0.3 < d < 0.2$) generate the most favourable combination of raw returns and effective variance.

These observations help explain why the intermediate-memory portfolios outperform the extremes. Portfolios Q2 (US) and Q3 (Japan) are composed primarily of assets whose memory estimates fall within this range. As a result, the compounding properties outlined in Section 4.3 reduced volatility drag for antipersistent dynamics, stabilised effective variance, and controlled skewness accumulation operate more consistently. In contrast, Q1 and Q5 contain assets whose extreme d values either magnify estimation uncertainty or induce higher turnover once rankings fluctuate, both of which deteriorate realised Sharpe ratios.

In summary, these findings confirm that the intermediate portfolios do not contradict the theoretical predictions. Rather, they highlight that the economic effects of memory are most reliably realised within the empirically stable region of the distribution. The 'winning' portfolios therefore emerge not at the extremes of d, but in the region where

FIGURE 4.3 : Cumulative Returns of Equally Weighted Portfolios with Transaction Costs 30 Day Rebalancing Frequency.

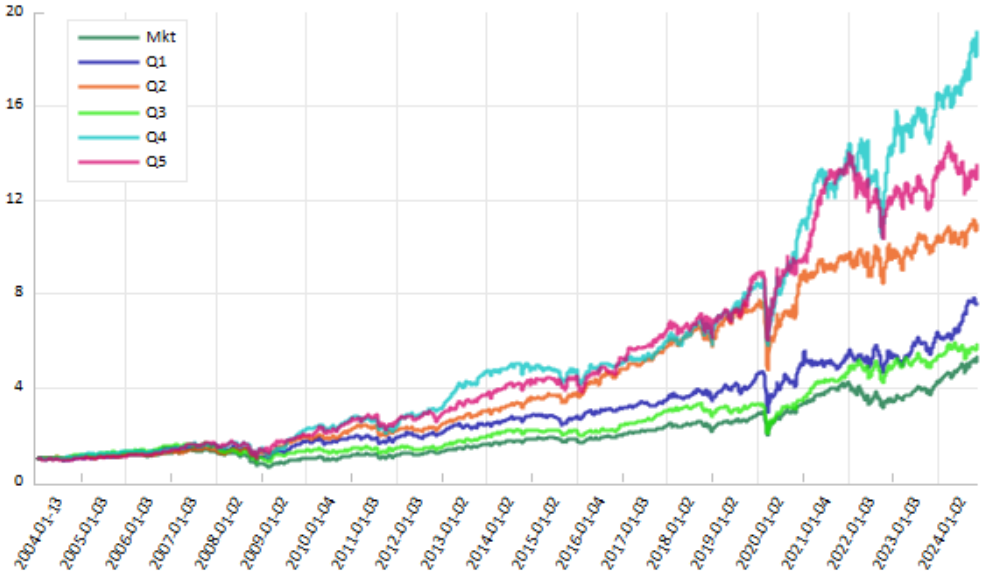

Note: This figure plots the cumulative returns of five quantile portfolios (Q1 Q5) and the market benchmark (Mkt) from 2004-01-13 to 2024-11-14. Each portfolio is equally weighted and rebalanced every 30 observations, incurring a 0.2% transaction cost per rebalancing event. Q1 represents the lowest memory assets, while Q5 represents the highest memory assets. The y-axis shows cumulative returns, normalised to 1.0 at the start date, and the x-axis indicates the timeline in daily increments.

memory estimates are statistically stable and the compounding dynamics predicted by the model operate most consistently.

4.6.6 Memory Portfolios vs Market Relative Performance.

A notable feature of our portfolio results is that each memory-sorted quintile consistently outperforms the market under both daily and monthly rebalancing strategies. Part of the explanation lies in higher-order moment behaviour under serial dependence. When $d > 0$, returns exhibit long-range dependence that propagates and amplifies any existing asymmetry in the innovation distribution, potentially generating a positive third cumulant (skew bonus) over multiple periods (Hoskin, 1981; Beran, 1994).

The skew bonus enters the expected log-return via a positive term $\kappa_{3,\text{eff}}(T)/(3(1+\mu)^3)$ in the Taylor expansion, which can partially offset the increased variance drag $\sigma^2_{\text{eff}}/(2(1+\mu)^2)$. Antipersistence ($d < 0$) reduces multi-period variance, narrowing return distributions, and mitigating drag, but dampens any skew accumulation. Persistence increases the accumulation of variance, but can augment the absolute third cumulant if underlying returns are skewed, offering upside potential. Portfolio strategies

FIGURE 4.4 : Average Estimated Memory vs. Average Daily Returns of US Market.

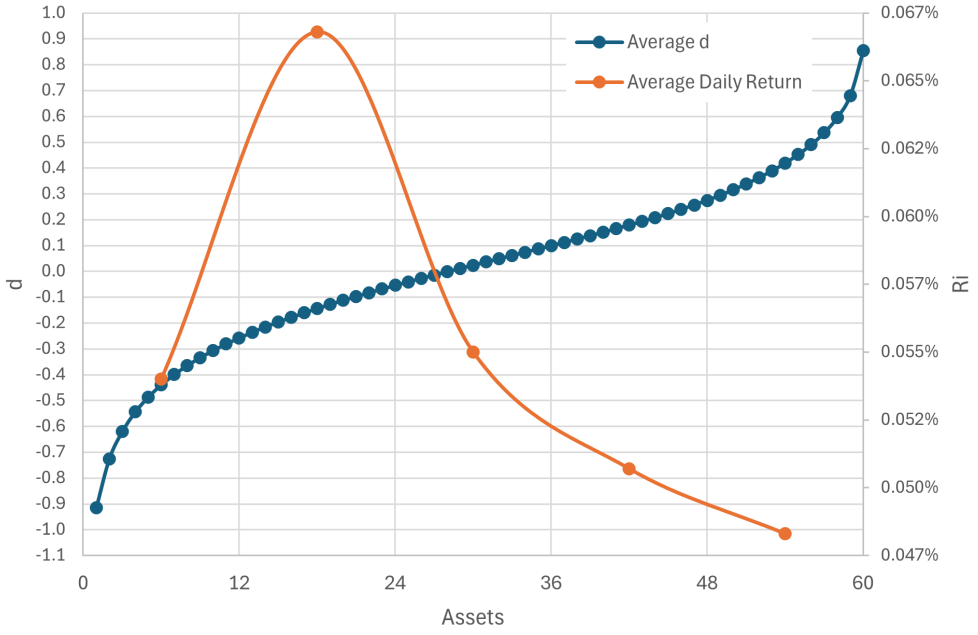

Note: The figure shows that extreme portfolios (Q1 and Q5) contain assets with $|d|$ close to or exceeding 0.5, where the ARFIMA$(0, d, 0)$ model becomes less stable and estimates exhibit higher sampling variability. In contrast, intermediate portfolios (Q2–Q3) lie firmly within the theoretically tractable region ($-0.3 < d < 0.2$), where the relationship between memory and multi-period compounding is most reliable. This pattern explains why the 'winning' portfolio does not always occur at the extremes.

that sort by memory profile can thus exploit mean-reversion benefits in antipersistent assets and skewness-amplified opportunities in persistent assets, balancing higher variance with skew bonus in expected log-growth.

Figure 4.5 provides an intuitive demonstration of how skewness alters expected log-returns. The solid black curve is $\ln(1+r)$, emphasising its concavity. For each distribution, we draw the **chord** joining its 1% and 99% quantile points: blue for the symmetric (normal) benchmark, magenta for the positively skewed alternative which has heavier right-tail mass. Beneath each chord the shaded probability density functions use the matching colour. The vertical dotted lines project the quantile endpoints and the mean return to the x and y-axes, highlighting the Jensen gaps. Because the gradient of the magenta chord is a closer match of the concave log-curve than the blue chord, its Jensen gap is smaller. This reduced shortfall is exactly the 'skew bonus' in $\mathbb{E}[\ln(1+r)]$. This mechanism corresponds exactly to the third-order term in Eq. (4.46), which can partially offset (or inflate, if negative) the variance drag $\sigma_{\text{eff}}^2/(2(1+\mu)^2)$.

Formally, extending Eq.(4.6) to third-order Taylor expansion of the single-period log-return $\ell_t = \ln(1+r_t)$ about $r_t = \mu$:

$$\ell_t = \ln(1+\mu) + \frac{r_t-\mu}{1+\mu} - \frac{(r_t-\mu)^2}{2(1+\mu)^2} + \frac{(r_t-\mu)^3}{3(1+\mu)^3} + O\Big((r_t-\mu)^4\Big).$$

FIGURE 4.5 : Illustration of The Effect of Volatility Drag of Investor's Expected Log-Returns in Presence of Positive Skewness.

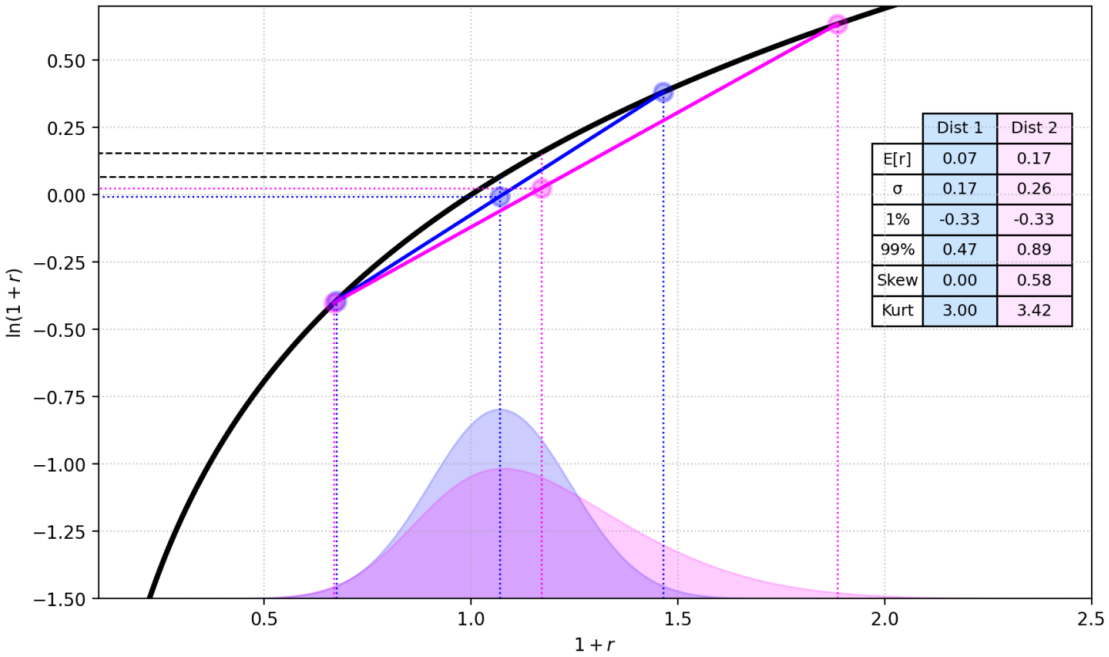

Note: The solid black curve plots $y = \ln(1 + r)$, highlighting its strict concavity. The magenta **chord** joins the two $(1 + r)$-points that bound the central mass of the positively skewed return distribution, whose mean is $r = \mu$. Beneath the curve, the blue shaded area represents a symmetric (normal) return distribution with mean μ and variance σ^2; the magenta shaded area depicts a positively skewed distribution with larger variance and a right–shifted mean due to heavier density in the right tail. Vertical distances between the two chords and the log-curve, projected via horizontal dotted lines to the y-axis, measure the Jensen gap for each density: the blue region shows the baseline shortfall under symmetry, while the additional mass in the right tail of the skewed distribution lies closer to its chord, thereby reducing the shortfall and generating the 'skew bonus'.

Summing over $t = 1, \ldots, T$ and taking expectations yields

$$\mathbb{E}\left[\sum_{t=1}^{T} \ell_t\right] \approx T \ln(1 + \mu) - \frac{\sigma^2_{\text{eff}}(T)}{2(1 + \mu)^2} + \frac{\kappa_{3,\text{eff}}(T)}{3(1 + \mu)^3}, \tag{4.46}$$

where *effective third central moment* is

$$\kappa_{3,\text{eff}}(T) = \mathbb{E}\left[\left(\sum_{t=1}^{T}(r_t - \mu)\right)^3\right] = T\,\kappa_3 + 3\sum_{k=1}^{T-1}(T - k)\ \text{Cov}\Big((r_t - \mu)^2,\ r_{t+k} - \mu\Big). \tag{4.47}$$

Eq. (4.47) relies on two regularity conditions. First, the demeaned return series $x_t = r_t - \mu$ is taken to have a time-invariant third central moment, $\kappa_3 = \mathbb{E}[x_t^3]$, so the single-time marginal skewness does not drift over the sample. Second, the third-order joint cumulant $\text{cum}(x_i, x_j, x_k)$ is set to zero whenever the three time indices are distinct. This eliminates the triple sum term $6\sum_{i<j<k} \mathbb{E}[x_i x_j x_k]$ that would otherwise appear in the full expansion. The latter condition holds exactly for Gaussian (and, more broadly, elliptical) innovations and is a close empirical approximation once heavy-tail property is already captured by the one-period moment κ_3. These simplifications facilitate tractable

expressions for $\kappa_{3,\text{eff}}(T)$, but their empirical adequacy should be assessed in applications.

Define the zero-lag squared–linear covariance

$$C_{21}(0) = \text{Cov}\Big((r_t - \mu)^2,\ r_t - \mu\Big), \tag{4.48}$$

and the lag-k skew-correlation

$$\rho_{21}(k) = \frac{\text{Cov}\Big((r_t - \mu)^2,\ r_{t+k} - \mu\Big)}{C_{21}(0)}. \tag{4.49}$$

Then each lag's contribution factors as $C_{21}(0)\,\rho_{21}(k)$, giving the equivalent form

$$\kappa_{3,\text{eff}}(T) = T\,\kappa_3 + 3\,C_{21}(0)\sum_{k=1}^{T-1}(T-k)\,\rho_{21}(k). \tag{4.50}$$

Under $ARFIMA(0,\ d,\ 0)$, the filter coefficients satisfy $\psi_j \sim C\,j^{d-1}$ (Hoskin, 1981; Beran, 1994), and one shows

$$\rho_{21}(k)\ \sim\ \frac{\Gamma(1-3d)}{\Gamma(d)\,\Gamma(1-2d)}\ k^{3d-2} \quad (k \to \infty). \tag{4.51}$$

Hence persistence ($d > 0$) slows the power-law decay of $\rho_{21}(k)$, amplifying $\kappa_{3,\text{eff}}(T)$, while antipersistence ($d < 0$) accelerates its decay.

Finally, define the *effective skewness*

$$\text{Skew}_{\text{eff}}(T) = \frac{\kappa_{3,\text{eff}}(T)}{\Big[\text{Var}\sum_{t=1}^{T} r_t\Big]^{3/2}} \approx \frac{\kappa_{3,\text{eff}}(T)}{\Big[T\,\sigma^2_{\text{eff}}(T)\Big]^{3/2}}. \tag{4.52}$$

Finally, we note that tail thickness in returns distribution expressed through kurtosis further complicates these dynamics. Leptokurtosis increases the likelihood of extreme returns in either tail, magnifying both the upside skew bonus and the downside volatility drag. As kurtosis increases, higher-order terms beyond the third moment begin to matter, and its effect on cumulative returns becomes implicit in the covariance structures that influence cumulant dynamics over time. Thus, in addition to variance and skewness, kurtosis should be considered when assessing the complete impact of memory-driven distributional shifts on long-term investment outcomes.

By sorting assets into memory-based quintiles, our strategy allows one to simultaneously exploit skewness bonuses and mitigate variance drag. Antipersistent portfolios ($d < 0$) benefit from reduced multi-period variance at the cost of muted skew bonus, while persistent portfolios ($d > 0$) incur greater variance drag but capture a positive third moment uplift when returns are positively skewed. This dual effect, analytically depicted

in Eq. (4.46), underpins the superior risk-adjusted returns of high memory portfolios (Q4 and Q5) observed in our empirical results.

4.6.7 Efficient–Frontier Diagnostics

In addition to our dynamic rebalancing strategies presented in previous section, we conduct a static analysis by examining efficient frontiers of each portfolio. Specifically, we assess how the diagonal and full-memory covariance models reshape the classical Markowitz frontier via three envelope-based diagnostics: the tangency Sharpe ratio, the signed area Δ (Eq. (4.44)), and the point-wise uplift at an arbitrary chosen 20 % volatility target (Eq. (4.45)). Table 4.10 summarises these metrics for each memory-sorted quintile, and Figure 4.6 plots convex envelope frontiers.

TABLE 4.10 : Tangency Portfolio Analysis of Different Covariance Models.

	Covariance	Sharpe	p-value	Area Δ	MaxRet	ΔMaxRet	MinRisk	ΔMinRisk	Uplift@20%
Q1	Conv	0.729			0.1924		0.1468		
	Diag	0.793	0.33	0.00118	0.1962	0.0038	0.1437	-0.0031	1.23%
	Full	0.811	0.54	0.00135	0.1962	0.0038	0.1369	-0.0099	1.34%
Q2	Conv	0.899			0.2331		0.1550		
	Diag	0.987	0.09	0.00150	0.2272	-0.0059	0.1502	-0.0048	1.13%
	Full	1.000	0.30	0.00170	0.2272	-0.0059	0.1470	-0.0080	1.27%
Q3	Conv	1.165			0.3834		0.1500		
	Diag	1.186	0.07	0.00145	0.3950	0.0116	0.1453	-0.0047	0.43%
	Full	1.191	0.02	0.00176	0.3950	0.0116	0.1426	-0.0074	0.64%
Q4	Conv	0.715			0.1769		0.1592		
	Diag	0.767	0.90	0.00617	0.2025	0.0256	0.1538	-0.0054	0.85%
	Full	0.774	0.29	0.00640	0.2025	0.0256	0.1529	-0.0063	1.03%
Q5	Conv	0.587			0.2724		0.1952		
	Diag	0.615	0.90	0.00566	0.3007	0.0256	0.1915	-0.0054	-0.37%
	Full	0.647	0.29	0.00750	0.3007	0.0256	0.1878	-0.0063	1.71%

Note: This table reports, for each memory-sorted quintile (Q1–Q5), the tangency Sharpe ratio under the conventional covariance (Conv), diagonal (Diag) and full memory-covariance (Full) specifications over the period 13-01-2004 to 14-11-2024. The p-value tests the null that the Sharpe of each memory-augmented model equals the Conv Sharpe. Area Δ is the signed area between the memory-augmented and conventional frontiers, computed as the integral of $\text{Return}_{\text{mem}} - \text{Return}_{\text{conv}}$ over the volatility domain. ΔMaxRet and ΔMinRisk are the shifts in the maximum-return and minimum-volatility endpoints of the frontier. Uplift@20% gives the excess return (in %) of the memory-augmented frontier at a fixed 20% volatility level. All return, volatility and Sharpe figures are annualised. See Section 4.6.7 for formal definitions.

In the strongly antipersistent universe (Q1), the diagonal memory covariance model increases the tangency Sharpe from 0.729 to 0.793 (8.8%), with a signed area change of $\Delta = 0.00118$ and a 1.23% uplift at 20% volatility. Incorporating full cross serial covariances raises Sharpe further to 0.811 (11.3%), increases the signed-area to $\Delta = 0.00135$ and delivers a 1.34% uplift. These results confirm that negative autocorrelations reduce multi-period variance (cf. Eq. (4.13)) and that modelling off-diagonal memory effects provides a further incremental benefit.

For Q2 (low-moderate memory) and Q3 (moderate memory), both memory-augmented models continue to outperform the conventional frontier. In Q2, Sharpe increases from 0.899 to 0.987 (9.8%) under diagonal covariance and to 1.000 (+11.3%) under full covariance, with signed area changes of $\Delta = 0.0017$ (Diagonal) and $\Delta = 0.0017$ (Full), and negligible 20% uplifts of 1.13% and 1.27% respectively. In Q3, Sharpe rises from 1.165 to 1.186 (1.8%) and to 1.191 (2.2%), with $\Delta = 0.0015$ and $\Delta = 0.0018$, producing 0.43% (Diagonal) and 0.64% (Full) uplifts compared to the conventional covariance model. These modest improvements mirror the milder persistence dynamics in these quintiles.

In high-memory regimes (Q4–Q5), full-memory covariance is essential to secure positive frontier shifts at the investor relevant 20% risk level. Although both diagonal and full models generate significant signed areas ($\Delta \approx 0.0062$ & 0.0064 in Q4 and $\Delta \approx 0.0057$ & 0.0075 in Q5), the diagonal model's 20% uplifts (0.85% in Q4, –0.37% in Q5) lag behind, whereas the full model delivers 1.03% and 1.71% respectively. This divergence further highlights the value of capturing cross-serial covariance when memory is strongly persistent.

Together, these envelope-based diagnostics validate our theoretical framework (Section 4.3): augmenting variance with serial dependence information systematically mitigates volatility drag and, when fully modelled, rotates the efficient frontier outward, thus reducing risk levels and improving returns.

4.7 Robustness Analysis: Evidence from the Japanese Market

Finally, to ensure that our results are not specific to the primary dataset and extend beyond the original empirical context, we replicate our empirical and analytical framework, including GRS testing, extended CAPM forecasting, and portfolio performance analysis on an alternative asset universe derived from the Japanese equity market. The robustness analysis employs a comprehensive set of equities listed on the Tokyo Stock Exchange, alongside the benchmark Nikkei 225 Index and the 10-Year Japanese government bond yield as the risk-free rate, spanning the period from 04-01-2010 to 03-04-2025 (3,669 observations).

Descriptive statistics reported in Table C.1 and Table C.2 in Appendix C demonstrate properties of returns and memory characteristics for considered assets. Asset-level memory estimates, derived via the Feasible Exact Local Whittle (FELW) method over a rolling 60-day window (calibrated so that the Nikkei 225 mean and median memory are equal zero), show a considerable variation across equities, indicating a wide spectrum of persistence profiles consistent with the original study.

FIGURE 4.6 : Efficient Frontier Graphs of Different Covariance Models.

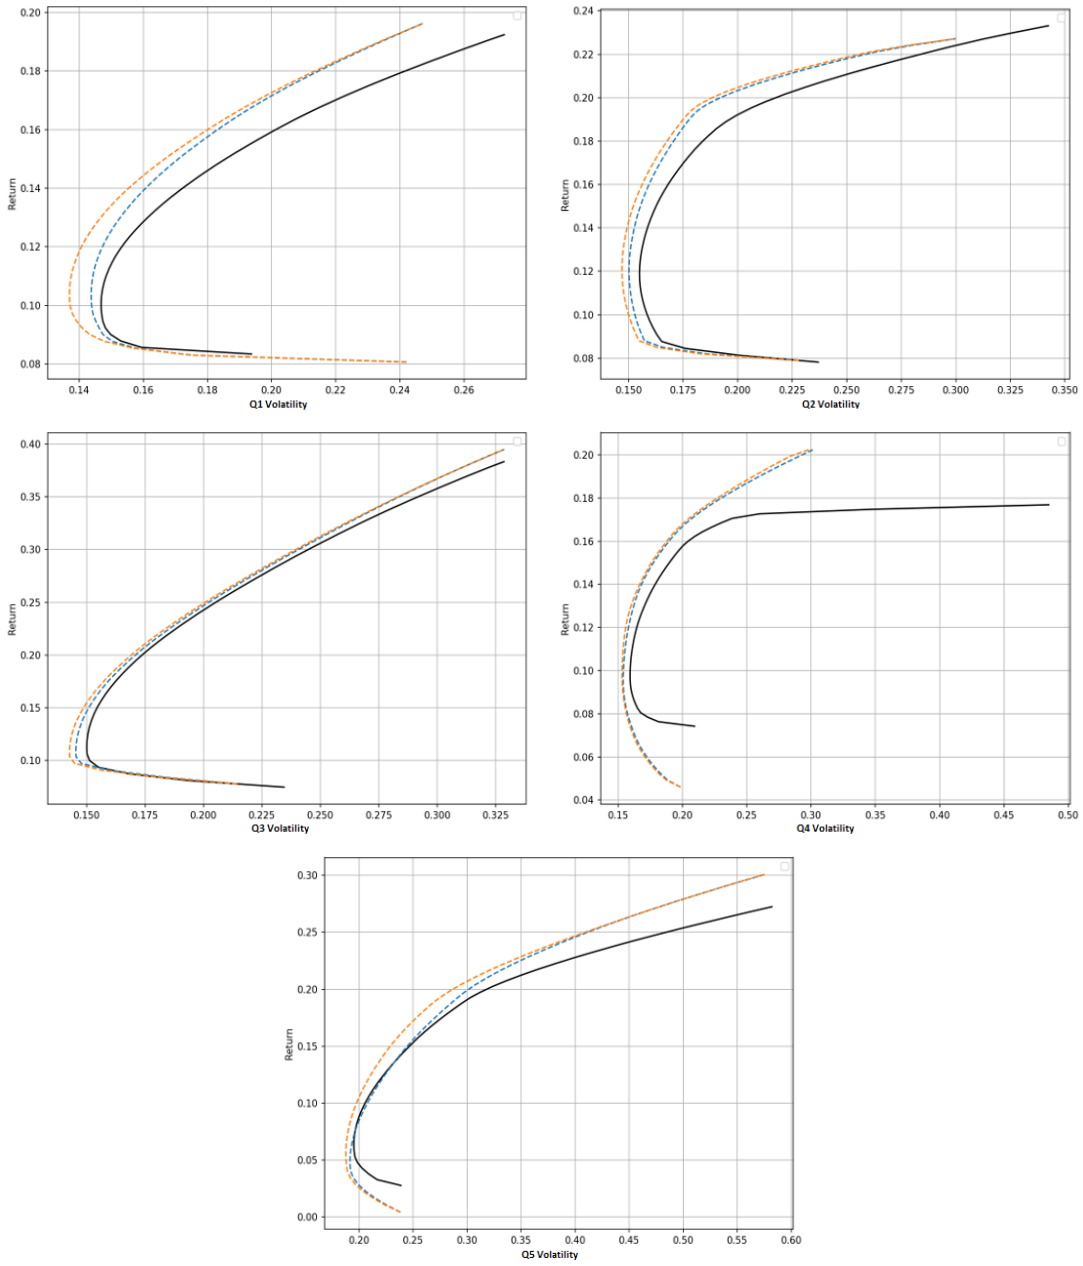

Note: Each panel displays the conventional efficient frontier - black line, the diagonal memory based covariance frontier - dashed blue line, and the full memory based covariance frontier - dashed orange line, for a given memory sorted quintile (Q1–Q5). The tangency portfolios underlying these frontiers maximise expected return per unit of effective variance under each covariance specification. Curves are plotted over annualised volatility ranging from the minimum-variance point to the maximum-return point. Detailed definitions of each frontier and the underlying optimisation appear in Section 4.6.7.

Portfolio construction based on memory-sorted quintiles (Tables C.3 and C.4) reaffirms the core patterns of the primary analysis. Quintile portfolios exhibit significant differentiation in memory and return characteristics. Notably, lower-memory quintiles display relatively higher moments compared to higher-memory quintiles, contrasting somewhat

with initial expectations, and highlighting distinctive return distributional characteristics across different persistence profiles. For instance, the lower-memory (antipersistent) Q2 portfolio demonstrated higher skewness and kurtosis, suggesting more pronounced occurrences of large positive or negative returns and more extreme deviations relative to a normal distribution. In contrast, higher-memory (persistent) portfolios Q4 and Q5 presented relatively lower kurtosis and skewness, implying return distributions closer to normality with fewer extreme events.

We extend our analysis by assessing the augmented CAPM, employing a memory-based HML factor. The results of Gibbons-Ross-Shanken tests presented in Table C.5 confirm the presence of significant intercepts (alphas) across specific quintiles and intervals, notably quintile 1 across the full sample and quintiles 1 and 2 during particular subintervals, providing robust evidence that memory also remains a priced factor in the Japanese market.

The forecasting performance evaluation of the extended CAPM versus the standard model, detailed in Table C.6, aligns closely with the original findings. Notably, memory-based augmentation substantially reduces forecast errors at the extremes of the memory spectrum (Q1 and Q5), reinforcing our initial conclusion that memory characteristics substantially enhance predictive accuracy.

Finally, portfolio performance metrics, reported in Tables C.7 and C.8 under daily and 30-day rebalancing strategies, corroborate our theoretical predictions. However, the notably strong performance of intermediate memory portfolios observed in the Japanese market was traced to anomalies in the price series for Kawasaki Heavy Industries (7012.T), specifically spurious 800-percent single-day returns drop and rise occurring in October 2011 and September 2015 respectively due to adjustment methodology of the data vendor. These artificial extremes inflated volatility estimates for the lowest and highest-memory quintiles, distorted memory parameter estimates towards the centre of the spectrum, and mechanically enhanced the risk-adjusted returns (Sharpe and Sortino ratios) of the intermediate portfolios. Removing 7012.T from the portfolio analysis, the performance hierarchy reverted to align with our main-sample results (persistent portfolios deliver robust cumulative returns performance), confirming that the initially observed intermediate-memory outperformance was a data-driven anomaly rather than a structural market characteristic (Figure C.3).

Overall, the replication on the Japanese market provides strong validation of our theoretical propositions, confirming that incorporating asset memory characteristics systematically enhances forecasting power, risk-adjusted returns, and strategic portfolio allocation efficacy across diverse market contexts.

4.7.1 Link with Existing Studies

The empirical patterns observed here are consistent with, but extend, several strands of the literature. First, the superior risk-adjusted performance of anti-persistent portfolios parallels evidence from volatility and intertemporal risk management strategies (Moreira and Muir, 2017; Barroso and Santa-Clara, 2015), which show that favourable temporal dependence structures enhance geometric returns for a given level of realised volatility. Second, the monotonic link between the memory parameter and effective variance aligns with long memory studies in asset returns (Baillie, 1996; Hoskin, 1981), as well as horizon-specific mispricing and risk amplification documented in fractional integration environments (León and Revéiz, 2010). Third, the strong performance of memory-sorted portfolios echoes findings in the sequence-of-returns and drawdown literature (Grossman and Zhou, 1993; Chekhlov et al., 2005; Clare et al., 2020), which emphasise that long-horizon outcomes are shaped by the ordering of shocks rather than average returns alone.

Overall, these results reinforce the view that temporal dependence constitutes an economically meaningful state variable. By linking fractional memory to compounding efficiency and realised portfolio performance, the chapter extends existing reduced-form evidence by providing a structural explanation of why assets with anti-persistent dynamics generate higher long-run risk-adjusted returns.

4.8 Limitations

While the chapter demonstrates robust theoretical and empirical support for memory-based portfolio construction, several limitations should be acknowledged. First, the asset universe consists of large capitalisation US equities, which may introduce survivorship, liquidity, and sectoral biases, although results are validated on the Japanese market to mitigate these concerns. Second, fractional memory estimates obtained through the FELW estimator are subject to sampling variability, particularly when computed over finite rolling windows, and may shift across regimes. Third, the analysis treats the memory parameter d as a continuous structural input rather than a hypothesis testing coefficient. As such, its economic impact arises through the geometry of multi-period compounding rather than statistical significance from zero. Finally, the modelling choices, including the estimation window, bandwidth parameter m, and innovation assumptions, should be interpreted as practical approximations rather than exhaustive representations of underlying market dynamics. These limitations highlight directions for future work, including time-varying memory estimation, alternative sampling designs, and broader multi-asset validation.

4.9 Conclusion

This study has addressed critical limitations inherent in classical mean-variance optimisation methodologies, specifically their single-period orientation, assumptions of independent and identically distributed returns, and neglect of the multi-period compounding dynamics shaped by serial dependence. While multi-period frameworks have been previously considered in the literature, this thesis advances the field by introducing a comprehensive theoretical approach that explicitly incorporates fractional memory through $ARFIMA(0,\ d,\ 0)$ processes. This enables precise derivations of closed-form solutions for effective variance, volatility drag, and higher-order moment adjustments for serial correlations.

Empirical tests robustly confirmed that memory characteristics represent economically significant and priced dimensions of risk, extending beyond traditional market exposures. Asset pricing analyses conducted on both US and Japanese markets confirmed the existence of a significant memory premium, particularly pronounced at extreme ends of the memory spectrum, improving forecasting accuracy by as much as 44% relative to conventional models. These findings underscore the necessity of explicitly incorporating fractional memory dynamics into modern asset-pricing models and systematic risk management practices.

Portfolio performance analyses provided clear empirical validation of the theoretical predictions regarding volatility drag and higher-order moments such as skewness. Portfolios formed according to memory-sorted quintiles consistently showed marked differences in realised volatility, tail risk, and risk-adjusted returns. Strong antipersistence reliably translated into substantial volatility drag mitigation and enhanced Sharpe ratios, confirming theoretical insights that negative serial correlations systematically reduce effective volatility. Conversely, high-memory portfolios exhibited robust returns despite anticipated volatility amplification, largely attributable to diversification effects arising from synchronised memory profiles and beneficial skewness adjustments.

Introducing realistic transaction costs into the analysis identified crucial trade-offs between drag reduction and turnover-related expenses. Intermediate levels of memory emerged as particularly advantageous, balancing the efficiency gains from reduced volatility drag with practical trading frictions. This finding was robustly reinforced by our empirical validation in the Japanese market, demonstrating that memory-based optimisation is consistently applicable across diverse market contexts. Moreover, portfolio diagnostics, such as efficient frontier analyses and skewness-adjusted returns, further solidified the practical relevance of explicitly modelling memory.

The robustness checks conducted in the Japanese market strongly corroborated the generalisability and applicability of our framework. Results were qualitatively consistent, reaffirming the theoretical predictions across different market structures, and

highlighting the persistent relevance of fractional memory in shaping return dynamics and portfolio performance. The systematic inclusion of memory parameters delivered meaningful performance improvements, validating our framework's efficacy beyond the primary market of analysis.

Future research avenues are promising and should focus on extensions including time-varying fractional memory parameters, developing dynamic, regime-dependent rebalancing rules responsive to shifting market conditions, and integrating forecasts of transaction costs within optimisation procedures. Furthermore, expanding analyses across asset classes, international markets, and alternative investor utility frameworks could provide further validation and broader application contexts. A particularly fruitful direction would leverage our derived expressions for effective variance and skewness to analytically pinpoint optimal rebalancing intervals, effectively balancing volatility drag, skewness benefits, and transaction costs in real-world portfolio management settings.

In summary, this study makes four key contributions. First, it introduces a comprehensive analytic toolkit that explicitly links multi-period log-wealth moments to k-lag autocorrelations and higher-order cumulants, enabling closed-form calculations of effective variance and skewness directly from observed dependence structures. Second, it generalises traditional momentum-based investment approaches by incorporating the persistence parameter d into expected return forecasts, seamlessly unifying positive momentum and negative mean-reversion effects within a single optimisation framework. Third, the proposed methodology advances long-horizon risk management by explicitly quantifying volatility drag and skew bonus as functions of autocovariance patterns, enabling precise risk budgeting and dynamic hedging strategies that improve upon heuristic stop-loss or calendar-based approaches. Fourth, it demonstrates cross-market robustness through empirical validation across both US and Japanese equity markets, highlighting portability and practical relevance of this framework. By integrating rigorous theoretical foundations with extensive empirical evidence, this framework provides investors with powerful, academically robust tools to enhance portfolio returns and manage risk effectively.

Chapter 5

Conclusions

This thesis set out to understand and harness the nature of persistence in financial markets, demonstrating that persistence is not merely a statistical artefact but a practically actionable and economically meaningful property of financial return-generating systems. Across three interconnected chapters, the research progresses from defining economically coherent forecasting systems using spillover-based connectedness, to characterising price formation through bidirectional persistence, and finally to embedding persistence directly within a multi-period asset-pricing framework. This framework is explicitly aligned with long-horizon investor objectives, including the maximisation of compounded wealth and risk-adjusted performance in the presence of serial dependence.

In Chapter 2, the thesis addressed the fundamental challenge of variable selection in financial forecasting. Conventional selection procedures, whether driven by in-sample fit or expert judgement, can under-represent the interconnected structure of financial systems, leading to unstable model specification and fragile out-of-sample performance. The spillover-based selection framework introduced here operationalises system definition by explicitly quantifying directional shock transmission and identifying variables that meaningfully contribute to the internal dynamics of the target market. This has two practical implications. First, it provides a disciplined route for building forecasting models that are less sensitive to arbitrary feature choices, since inclusion is justified by incremental informational role of a variable within the system rather than by predictive coincidence. Second, it produces a more defensible basis for persistence measurement, because memory estimation becomes conditional on a well-defined system boundary, this reduces the risk that estimated persistence reflects omitted interconnected drivers or contamination from loosely related variables. In applied settings, the spillover framework therefore functions not just as a forecasting enhancement but as a model governance tool, such that it produces interpretable evidence for why each variable belongs in the model, how information flows through the system, and which nodes are structurally influential.

Building on this foundation, Chapter 3 examined the behavioural and structural content of persistence through a bidirectional view of memory in price formation, contrasting a fundamentally anchored market (Brent) with a market that is more driven by expectation and sentiment (Bitcoin). The central contribution is to show that persistence should be interpreted as a signature of how information is incorporated across time, and that the relevant 'time direction' differs by market structure. Fundamentally driven markets exhibit stronger backward-looking persistence consistent with slow-moving adjustment to physical constraints, inventories, policy, and geopolitical shocks, whereas speculative markets exhibit stronger forward-looking persistence consistent with anticipatory belief formation and feedback trading. Practically, this result reframes how persistence can be used in forecasting and risk management. It implies that the same persistence estimate can represent different economic mechanisms across markets, and that model design should reflect this: forecasting frameworks for speculative assets should explicitly accommodate expectation channels and regime sensitivity, whereas forecasting frameworks for fundamentally anchored assets should prioritise slow-moving equilibrium adjustment and persistent fundamentals. More broadly, the bidirectional framework provides a mechanism for interpreting why forecast performance varies across assets and horizons: persistence embeds information about whether the market is primarily processing history, anticipating futures, or switching between the two.

Chapter 4 translates persistence into investment decision-making by developing a memory-driven, multi-period asset pricing model aligned with long-horizon compounding objectives. The central theoretical insight is that the economic cost of risk is not fully captured by single-period variance, since volatility drag arising from the concavity of geometric returns accumulates over time and is either amplified or attenuated by serial dependence. By linking fractional memory to closed-form effective variance, and extending the analysis to effective skewness under serial dependence, the chapter provides a tractable mapping from dependence structure to long-horizon wealth dynamics. Empirically, the chapter demonstrates that profiling assets by memory is economically consequential: antipersistence is associated with reduced effective variance and improved risk-adjusted performance, while persistence can incur greater variance drag yet may deliver skewness-related upside when positive asymmetry is present. The memory-based factor construction further shows that persistence behaves as a priced dimension of risk beyond conventional market exposure, and that incorporating this dimension can materially improve predictive accuracy where memory signals are strongest.

Jointly, these results imply several direct practical implications for quantitative investment practice. First, persistence should be treated as an input to risk measurement rather than as a secondary diagnostic. In multi-period contexts, the relevant risk for long-horizon investors is an effective risk that depends on the entire autocorrelation structure, not just contemporaneous volatility. Consequently, risk budgeting, portfolio

optimisation, and performance attribution can be improved by replacing purely one-period variance inputs with effective variance measures that reflect the dependence structure of returns. Second, persistence offers a unifying interpretation of mean-reversion and momentum behaviours within a single analytical framework. In this thesis, negative memory is shown to correspond to economically exploitable mean-reversion effects through volatility drag mitigation, while positive memory aligns with momentum-like persistence that can be beneficial when accompanied by favourable higher order moments. This interpretation provides a coherent bridge between time series dependence and the practical language of trading signals and portfolio tilts. Third, the results emphasise that portfolio design is inseparable from implementation frictions. The transaction cost analysis shows that the gains from persistence-aware allocation are realised through a balance between dependence-driven drag reduction and turnover; in practice, this implies that the optimal use of persistence is likely to be dependent on horizon and cost, and that memory-aware strategies should be evaluated in conjunction with explicit rebalancing rules.

The thesis also has clear limitations that bound interpretation and motivate further work. First, persistence estimates are subject to finite sample uncertainty and potential regime instability, particularly when computed in rolling windows and when assets exhibit structural breaks, microstructure effects, or episodic liquidity constraints. Second, the theoretical mapping in Chapter 4 relies on approximations that are economically transparent but not exhaustive: Taylor expansions and cumulant truncations provide tractable links between dependence and long-horizon utility, yet higher-order effects may become material under extreme volatility, heavy tails, or strongly time-varying skewness. Third, the empirical exercises are necessarily conditional on data quality and measurement choices, including the selected bandwidth for semiparametric memory estimation, the rolling-window length, and the practical treatment of outliers and corporate actions; the Japanese robustness exercise highlights that data vendor adjustments can generate spurious extremes that distort both volatility and memory measures if not detected and corrected. Fourth, the portfolio tests focus on large-cap equities and two representative markets for price formation analysis; while cross-market robustness is supported, generalisation to other asset classes, higher-frequency horizons, and leveraged or short-enabled implementations requires additional validation.

These limitations naturally suggest several promising avenues for future research. A first direction is to develop explicitly time-varying and regime-dependent persistence models, in which memory is treated as a latent state that evolves with market structure, liquidity, and the information environment. This would align persistence estimation more closely with the empirical observation that dependence strength can shift during stress episodes and structural transitions. A second direction is to integrate persistence-aware risk measures into richer multivariate volatility frameworks, allowing effective variance and covariance to be forecast jointly with conditional heteroskedasticity and

cross-asset dependence; this would provide a more complete bridge between the long-memory literature and standard risk engines. A third direction is to formalise the interaction between persistence, transaction costs, and optimal rebalancing frequency by using the derived effective variance and skewness expressions to identify rebalancing intervals that maximise long-horizon expected utility net of costs. A fourth direction is broader validation across asset classes and market structures, including rates, credit, commodities beyond oil, and systematic alternative risk premia, where persistence may arise from different structural constraints and institutional trading behaviours. Finally, there is scope to connect the spillover-based system definition of Chapter 2 with the pricing framework of Chapter 4 more tightly, by using network-derived measures of systemic influence as conditioning variables for time-varying persistence premia and for stress-sensitive portfolio tilts.

In summary, this thesis establishes persistence as a fundamental and actionable characteristic of financial markets. It provides a coherent toolkit for identifying economically relevant systems, interpreting price formation through the time direction of dependence, and translating persistence into investor-relevant risk and return adjustments under multi-period compounding. By combining analytical derivations with empirical validation across markets and by confronting implementation frictions explicitly, the thesis advances both the theoretical understanding of persistence and its practical use in forecasting, asset pricing, and portfolio management.

Appendix A

Supplement to Chapter 2

This appendix provides the formulation and explanation of our spillover analysis approach, following the approach outlined by Diebold and Yılmaz (2014). This framework quantifies the magnitude and direction of spillovers among the variables of interest, offering a comprehensive perspective on how shocks propagate through the system. We derive three key spillover measures: the TSI, Directional Spillovers, and Net Spillover.

A.1 Static (Time-Domain) Spillover Analysis

We begin with a VAR model:

$$y_t = \sum_{i=1}^{p} \mu_i \, y_{t-i} + \epsilon_t, \tag{A.1}$$

which is expressed in its infinite-order moving average (MA) representation:

$$y_t = \sum_{i=0}^{\infty} \delta_i \, \epsilon_{t-i}, \tag{A.2}$$

with δ_0 as the identity matrix and, for $i \geq 1$,

$$\delta_i = \mu_1 \, \delta_{i-1} + \mu_2 \, \delta_{i-2} + \cdots + \mu_p \, \delta_{i-p}.$$

Following Diebold and Yilmaz (2012), we use the generalised variance decomposition (GVD) to compute the h-step-ahead forecast error variance decomposition for variable i due to shocks in variable j:

$$\tau_{ij}^{g}(h) = \frac{\sigma_{ii}^{-1} \sum_{h=0}^{H-1} \left(e_i' \, \delta_h \, \Sigma \, e_j \right)^2}{\sum_{h=0}^{H-1} \left(e_i' \, \delta_h \, \Sigma \, \delta_h' \, e_i \right)}, \tag{A.3}$$

where e_i is a selection vector with unity in the ith position, Σ is the covariance matrix of ϵ_t, and σ_{ii} is the ith diagonal element of Σ.

Because the forecast error contributions do not naturally sum to one, we normalise them:

$$\tilde{\tau}_{ij}^{g}(h) = \frac{\tau_{ij}^{g}(h)}{\sum_{j=1}^{N} \tau_{ij}^{g}(h)}, \tag{A.4}$$

ensuring $\sum_{j=1}^{N} \tilde{\tau}_{ij}^{g}(h) = 1$ for each i, and that the overall sum across the system equals N.

From these normalised values, we derive:

- **Total Spillover Index (TSI):**

$$S.O^{g}(h) = \frac{\sum\limits_{\substack{i,j=1 \\ i \neq j}}^{N} \tilde{\tau}_{ij}^{g}(h)}{N} \times 100, \tag{A.5}$$

  which measures overall cross-variable shock transmission.

- **Directional Spillovers:**

$$S.O_{i.}^{g}(h) = \frac{\sum_{j \neq i} \tilde{\tau}_{ij}^{g}(h)}{\sum_{j=1}^{N} \tilde{\tau}_{ij}^{g}(h)} \times 100, \quad S.O_{.i}^{g}(h) = \frac{\sum_{j \neq i} \tilde{\tau}_{ji}^{g}(h)}{\sum_{j=1}^{N} \tilde{\tau}_{ji}^{g}(h)} \times 100,$$

  indicating the extent to which each variable transmits or receives shocks.

- **Net Spillover:**

$$S.O_{i}^{g}(h) = S.O_{.i}^{g}(h) - S.O_{i.}^{g}(h), \tag{A.6}$$

  which indicates whether a variable is a net transmitter (positive) or receiver (negative).

A.2 Fractional Cointegration Vector Autoregressive Model

One approach for modelling systems of variables that are defined by interdependent dynamic and capture slow convergent memory is FCVAR framework. This coherent framework was first proposed by Johansen (2008) and further extended by Johansen and Nielsen (2012b) as a generalisation of Johansen's cointegrated vector autoregressive model (CVAR) (Johansen, 1995a).

A.2.1 Model Derivation

One approach for modelling systems of variables that are defined by interdependent dynamics and exhibit slow, persistent adjustment is the FCVAR framework. This model was first proposed by Johansen (2008) and subsequently extended by Johansen and Nielsen (2012b) as a generalisation of Johansen's cointegrated vector autoregressive model (CVAR) (Johansen, 1995a).

Before presenting the formal derivation of the FCVAR model, it is useful to clarify the notation used for fractional differencing and lag operators. The fractional integration parameter d governs the long-run persistence of the system, whereas the parameter b, with $0 \leq b \leq d$, determines the fractional degree of the error-correction mechanism. The case $b = 1$ yields the standard VECM as a special case of the FCVAR model.

Fractional differencing of order d is defined as $\Delta^d = (1-L)^d$, with the corresponding fractional lag operator $L_d = 1 - \Delta^d$. Analogously, for the adjustment parameter b, we write $\Delta^b = (1-L)^b$ and $L_b = 1 - \Delta^b$. These operators play complementary roles: d controls long-memory behaviour in levels, while b governs the smoothness of the equilibrium adjustment. This notation ensures that expressions such as $L_b\Delta^{d-b}X_t$ and $\Delta^d L_b^i X_t$ appearing in the FCVAR representation follow directly from the underlying fractional algebra.

In order to derive FCVAR the easiest approach is to consider CVAR model first. The CVAR model is essential for understanding the dynamics of multiple timeseries that are integrated of order one, $I(1)$, and are cointegrated. This model is defined as:

$$\Delta Y_t = \alpha\beta' Y_{t-1} + \sum_{i=1}^{k} \Gamma_i \Delta Y_{t-i} + \varepsilon = \alpha\beta' L^i Y_t + \varepsilon_t. \tag{A.7}$$

To obtain FCVAR model from the above equation, we can replace the difference and lag operators Δ and $L = 1 - \Delta$ by their fractional counterparts, Δ_b and $L_b = 1 - \Delta^b$

$$\Delta^b Y_t = \alpha\beta' L_b Y_t + \sum_{i=0}^{k} \Gamma_i \Delta^b L_b^i Y_t + \varepsilon_t, \tag{A.8}$$

Applying it to $Y_t = \Delta^{d-b} X_t$ it becomes:

$$\Delta^d X_t = \alpha\beta' L_b \Delta^{d-b} X_t + \sum_{i=o}^{k} \Gamma_i \Delta^d L_b^i X_t + \varepsilon_t \tag{A.9}$$

Where ε_t is an independent identically distributed (iid) $p - dimensional$ error term with $E(\varepsilon) = 0$ and $var(\varepsilon_t) = \sigma^2$.

Similarly to the CVAR framework, the FCVAR model treats α and β as $(p \times r)$ matrices with $0 \leq r \leq p$. The matrix α governs the adjustment speed to equilibrium, while β

encodes the cointegrating relationships so that $\beta' X_t$ represents the long-run equilibrium, with Γ_i driving short-run dynamics. To address bias from approximating fractional differences by an infinite series on finite samples, Johansen and Nielsen (Johansen and Nielsen, 2016) suggest adding a level parameter μ that shifts each series by a constant. This modification improves the capture of long memory dynamics. With this adjustment, the equation becomes:

$$\Delta^d(X_t - \mu) = \alpha\beta' L_d(X_t - \mu) + \sum_{i=0}^{k} \Gamma_i \Delta^d L_d^i(X_t - \mu) + \varepsilon_t, \qquad t = 1, 2, \cdots, T \quad \text{(A.10)}$$

Effectively, this addition adjusts the data series for a constant shift μ before applying fractional differencing and truncating it at $k = t - 1$. This allows the FCVAR model to more accurately reflect the characteristics of systems with long memory.

A.2.2 Forecasting Using FCVAR

Due to the autoregressive nature of FCVAR model, it's forecasting methodology is relatively straight forward. To demonstrate this aspect, consider the following model:

$$\Delta^d(X_{t+1} - \mu) = X_{t+1} - \mu - L_d(X_{t+1} - \mu) \quad \text{(A.11)}$$

Rearranging for X_{t+1} gives the following relationship:

$$X_{t+1} = \mu + L_d(X_{t+1} - \mu) + \alpha\beta' L_b \Delta^{d-b}(X_{t+1} - \mu) + \sum_{i=o}^{k} \Gamma_i \Delta^d L_b^i(X_{t+1} - \mu) + \varepsilon_{t+1} \quad \text{(A.12)}$$

Given that $L_b = 1 - \Delta^b$ functions as a lag operator, making $L_b^i X_{t+1}$ known at the time t for all $i \geq 1$, Equation B.6 serves as the basis for FCVAR forecasting. Expressing the relationship using coefficient estimates and conditional expectation notation based on information available at time t, a one-step-ahead forecast, $\widehat{X}_{t+1|t}$ is defined as:

$$\widehat{X}_{t+1|t} = \widehat{\mu} + L_{\hat{d}}(X_{t+1} - \widehat{\mu}) + \widehat{\alpha}\widehat{\beta}' L_{\hat{b}} \Delta^{\hat{d}-\hat{b}}(X_{t+1} - \widehat{\mu}) + \sum_{i=o}^{k} \widehat{\Gamma}_i \Delta^{\hat{d}} L_{\hat{b}}^i(X_{t+1} - \widehat{\mu}) \quad \text{(A.13)}$$

This relationship allows the FCVAR model to generate estimates for every variable within the system at each forecasting iteration. Thus, n-step ahead forecast can be formed recursively based on generalisation of Equation B.7.

$$\widehat{X}_{t+j|t} = \widehat{\mu} + L_{\hat{d}}(X_{t+j|t} - \widehat{\mu}) + \widehat{\alpha}\widehat{\beta}' L_{\hat{b}} \Delta^{\hat{d}-\hat{b}}(X_{t+j|t} - \widehat{\mu}) + \sum_{i=o}^{k} \widehat{\Gamma}_i \Delta^{\hat{d}} L_{\hat{b}}^i(X_{t+j|t} - \widehat{\mu}) \quad \text{(A.14)}$$

Where $\widehat{X}_{s|t} = X_s$ for $s \leq t$. In this case the forecasts are calculated recursively from Equation B.8 for $j = 1, 2, ..., h$ to generate h-step ahead forecast, $\widehat{X}_{t+j|t}$.

A.3 Six Model Spillover Analysis

This section concisely summarises the spillover analysis for six models investigating price discovery in the oil market over a 60-day forecast horizon covering both directional and frequency domain analysis.

A.3.1 Statics Analysis of Intra-Market Dynamism - Isolated Systems.

A.3.1.1 *Model 1 - BSP, BE, BL, BR.*

Directional Spillovers.

TABLE A.1 : Model 1 - Static Directional Connectedness in Brent Market.

Variable	BSP	BE	BL	BR	FROM
BSP	42.20	33.91	0.00	23.89	57.80
BE	41.94	44.07	0.02	13.98	55.93
BL	0.04	0.22	99.58	0.16	0.42
BR	32.99	14.20	0.01	52.81	47.19
TO	74.97	48.32	0.03	38.03	*TSI: 40.33%*
NET	17.17	-7.61	-0.39	-9.16	

Note: Following Diebold and Yilmaz (2012). Numbers are percentages. "TSI" stands for Total Spillover Index. $FROM$ - represents the measure of instrument's sensitivity to uncertainty generated by other variables in the system. TO - represents the measure of instrument's ability to transmit uncertainty to other variables in the system.

Table A.1 indicates a total spillover index of 40.33% and reveals that innovations in BE and BR notably influence BSP, which in turn transmits 74.97% to the system. The network representation in Figure A.1 (see below) confirms the dominance of BSP and the relative isolation of BL.

Frequency Domain Spillovers.

Table A.2 shows that interconnectedness is primarily concentrated over horizons exceeding five days, while short-term efficiency is maintained by high-frequency trading.

FIGURE A.1 : Model 1 - Network analysis of intra-market dynamics in Brent oil market.

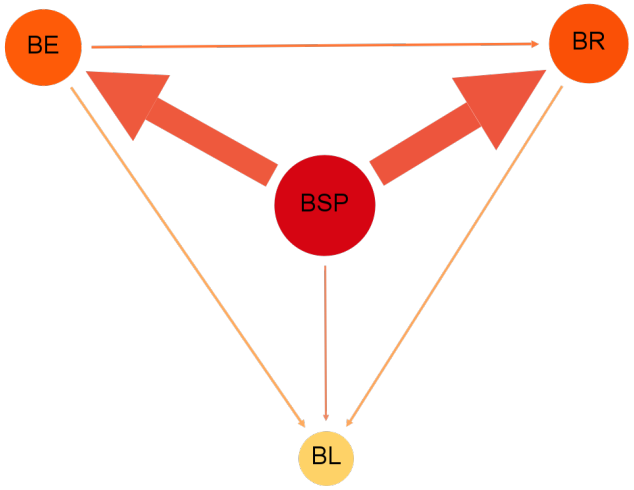

Note: Network analysis intra-market information exchange mechanism. Nodes notation represents: BE - Brent market expansion, BR - Brent market regime, BL - Brent market liquidity, BSP - Brent Spot.

TABLE A.2 : Model 1 - Frequency Decomposition of Static Directional Connectedness in Brent Market.

(a) Short horizon - 1 to 5 days

Variable	BSP	BE	BL	BR	FROM
BSP	0.00	0.00	0.00	0.00	0.01
BE	0.00	0.00	0.00	0.00	0.00
BL	0.00	0.00	0.05	0.00	0.01
BR	0.02	0.01	0.00	0.03	0.03
TO	0.02	0.01	0.00	0.01	*TSI: 33.71%*

(b) Long horizon - 5 days to infinity

Variable	BSP	BE	BL	BR	FROM
BSP	42.20	33.90	0.00	23.89	57.79
BE	41.93	44.07	0.02	13.97	55.93
BL	0.04	0.21	99.53	0.16	0.41
BR	32.97	14.19	0.00	52.78	47.16
TO	74.95	48.30	0.02	38.02	*TSI: 40.34%*

Note: Frequency domain analysis following Baruník and Krehlík (2018) methodology. Short and Long horizons refer to '1 to 5 days' and '5 days to infinity'. Numbers report percentages.

A.3.1.2 *Model 2 - WSP, WE, WL, WR.*

Directional Spillovers.

Table A.3 reports a total spillover index of 35.88%, with WE emerging as the main transmitter and WL contributing negligibly. The network graph in Figure A.2 (see

TABLE A.3 : Model 2 - Static Directional Connectedness in WTI market.

Variable	WSP	WE	WL	WR	FROM
WSP	29.04	40.59	0.02	30.35	70.96
WE	20.04	65.32	0.60	14.05%	34.68
WL	1.34	1.18	96.65	0.83 %	3.35
WR	22.26	11.52	0.74	65.47	34.53
TO	43.64	53.29	1.35	45.23	*TSI: 35.88%*
NET	-27.32	18.61	-2.00	10.70	

Note: Following Diebold and Yilmaz (2012). Numbers are percentages. "TSI" stands for Total Spillover Index. *FROM* - represents the measure of instrument's sensitivity to uncertainty generated by other variables in the system. *TO* - represents the measure of instrument's ability to transmit uncertainty to other variables in the system.

below) illustrates these dynamics.

FIGURE A.2 : Model 2 - Network analysis of intra-market dynamics in WTI oil market.

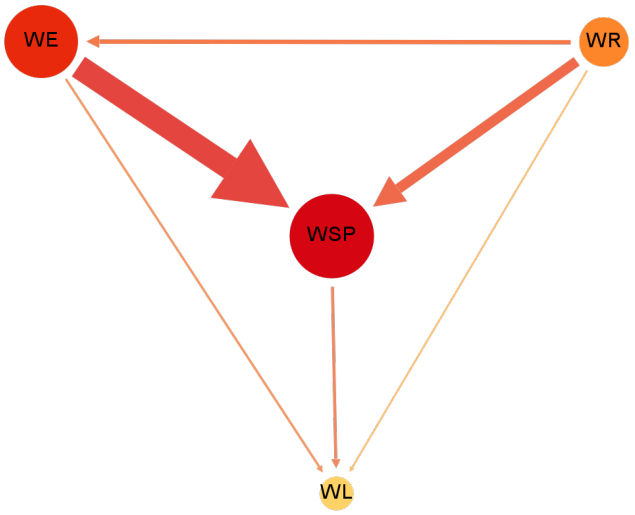

Note: Network analysis intra-market information exchange mechanism. Nodes notation represents: WE - WTI market expansion, WR - WTI market regime, WL - WTI market liquidity, WSP - WTI Spot.

Frequency Domain Spillovers.

TABLE A.4 : Model 1 - Frequency Decomposition of Static Directional Connectedness in WTI Market.

(a) Short horizon - 1 to 5 days

Variable	WSP	WE	WL	WR	FROM
WSP	0.10	0.14	0.00	0.00	0.15
WE	0.02	0.00	0.01	0.02	0.05
WL	0.02	0.01	2.33	0.04	0.07
WR	0.03	0.01	0.01	0.02	0.05
TO	0.08	0.16	0.02	0.05	*TSI: 11.29%*

(b) Long horizon - 5 days to infinity

Variable	WSP	WE	WL	WR	FROM
WSP	28.94	40.45	0.01	30.35	70.81
WE	20.01	65.32	0.59	14.03	34.63
WL	1.31	1.17	94.32	0.80	3.28
WR	22.23	11.51	0.73	65.45	34.47
TO	43.56	53.13	1.33	45.18	*TSI: 36.05%*

Note: Frequency domain analysis following Baruník and Krehlík (2018) methodology. Short and Long horizons refer to '1 to 5 days' and '5 days to infinity'. Numbers report percentages.

The frequency domain decomposition in Table A.4 aligns with Model 1, with long-run spillovers indicating exploitable inefficiencies and short-run stability maintained by speculative activity.

A.3.1.3 *Model 3 - BSP, BE, BL, BR, WSP, WE, WL, WR.*

Directional Spillovers.

TABLE A.5 : Model 2 - Static Directional Connectedness between Brent & WTI markets.

Variable	BSP	BE	BL	BR	WSP	WE	WL	WR	FROM
BSP	22.82	14.85	0.14	14.88	10.47	23.58	0.76	12.51	77.18
BE	20.41	17.76	0.05	9.16	10.05	33.78	1.34	7.46	82.24
BL	0.10	1.09	87.20	0.60	0.31	2.62	7.82	0.25	12.80
BR	21.30	8.77	0.36	28.60	8.22	8.08	0.16	24.50	71.40
WSP	17.51	11.33	0.35	10.18	15.99	25.54	0.70	18.41	84.01
WE	15.25	11.98	0.44	6.79	10.90	45.14	1.69	7.81	54.86
WL	0.40	0.78	9.71	0.29	0.12	4.32	83.63	0.74	16.37
WR	16.25	8.29	0.21	16.68	13.67	6.12	0.03	38.75	61.25
TO	91.22	57.10	11.26	58.58	53.74	104.03	12.51	71.68	*TSI: 57.51%*
NET	14.04	-25.14	-1.54	-12.82	-30.27	49.17	-3.86	10.43	

Note: Following Diebold and Yilmaz (2012). Numbers are percentages. "TSI" stands for Total Spillover Index. *FROM* - represents the measure of instrument's sensitivity to uncertainty generated by other variables in the system. *TO* - represents the measure of instrument's ability to transmit uncertainty to other variables in the system.

Table A.5 yields a total spillover index of 57.51%, reflecting enhanced bidirectional information flow when Brent and WTI markets are jointly considered. Dominant transmitters include WE, BSP, and WR. The corresponding network in Figure A.3 (see below) highlights the pre-eminence of WE.

FIGURE A.3 : Model 3 - Network analysis of inter-market dynamics between Brent & WTI markets.

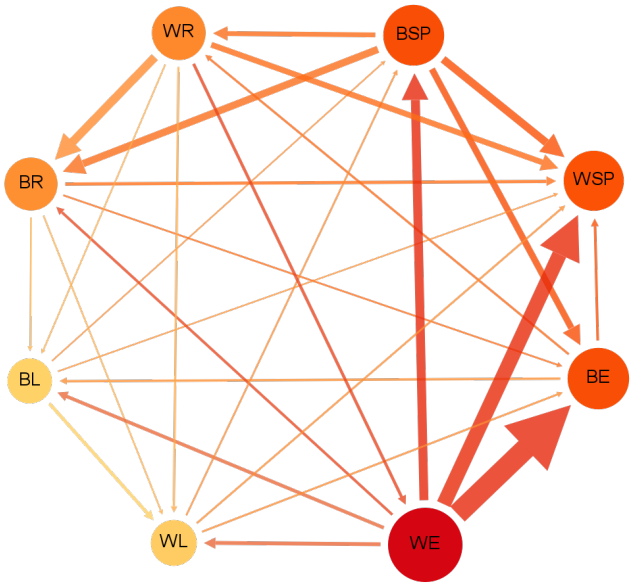

Note: Network analysis intra-market information exchange mechanism. Nodes notation represents: BE - Brent market expansion, BR - Brent market regime, BL - Brent market liquidity, BSP - Brent Spot; WE - WTI market expansion, WR - WTI market regime, WL - WTI market liquidity, WSP - WTI Spot.

{Frequency Domain Spillovers.

Frequency analysis (Table A.6) confirms that significant spillovers occur over long horizons, with short-run effects largely neutralised by market corrections.

TABLE A.6 : Model 3 - Frequency Decomposition of Static Directional Connectedness between Brent & WTI Markets.

(a) Short horizon - 1 to 5 days

Variable	BSP	BE	BL	BR	WSP	WE	WL	WR	FROM
BSP	0.00	0.01	0.01	0.00	0.00	0.09	0.01	0.00	0.13
BE	0.00	0.01	0.00	0.00	0.00	0.17	0.02	0.00	0.20
BL	0.01	0.01	0.63	0.02	0.01	0.06	0.12	0.01	0.24
BR	0.01	0.01	0.01	0.02	0.02	0.02	0.00	0.02	0.09
WSP	0.00	0.01	0.01	0.00	0.08	0.09	0.01	0.00	0.14
WE	0.01	0.00	0.01	0.00	0.00	0.00	0.02	0.00	0.06
WL	0.01	0.01	0.15	0.01	0.01	0.02	2.93	0.02	0.23
WR	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.01	0.05
TO	0.04	0.06	0.19	0.05	0.06	0.46	0.20	0.06	*TSI: 23.35%*

(b) Long horizon - 5 days to infinity

Variable	BSP	BE	BL	BR	WSP	WE	WL	WR	FROM
BSP	22.82	14.84	0.13	14.88	10.46	23.48	0.75	12.51	77.05
BE	20.41	17.75	0.05	9.15	10.05	33.61	1.31	7.45	82.04
BL	0.10	1.09	86.57	0.58	0.30	2.56	7.70	0.25	12.56
BR	21.29	8.76	0.35	28.58	8.20	8.07	0.16	24.48	71.31
WSP	17.51	11.32	0.34	10.18	15.91	25.44	0.69	18.41	83.87
WE	15.24	11.98	0.43	6.78	10.89	45.14	1.67	7.81	54.80
WL	0.39	0.77	9.56	0.29	0.12	4.30	80.70	0.72	16.14
WR	16.24	8.28	0.21	16.67	13.66	6.11	0.03	38.74	61.20
TO	91.17	57.05	11.07	58.53	53.68	103.57	12.31	71.62	*TSI: 57.72%*

Note: Frequency domain analysis following Baruník and Krehlík (2018) methodology. Short and Long horizons refer to '1 to 5 days' and '5 days to infinity'. Numbers report percentages.

A.3.2 Statics Analysis of Intra-Market Dynamism - Global Systems.

A.3.2.1 *Model 4 - BSP, E, L, R.*

Directional Spillovers.

TABLE A.7 : Model 4 - Static Directional Connectedness in Brent Market represented through global oil market variables.

Variable	BSP	E	L	R	FROM
BSP	42.72	45.70	4.20	7.38	57.28
E	36.96	46.41	3.73	12.90	53.59
L	14.87	14.08	70.73	0.31	29.27
R	11.40	18.04	0.66	69.91	30.09
TO	63.23	77.82	8.59	20.60	*TSI: 42.56%*
NET	5.95	24.22	-20.68	-9.50	

Note: Following Diebold and Yilmaz (2012). Numbers are percentages. "TSI" stands for Total Spillover Index. *FROM* - represents the measure of instrument's sensitivity to uncertainty generated by other variables in the system. *TO* - represents the measure of instrument's ability to transmit uncertainty to other variables in the system.

Table A.7 reports a total spillover index of 42.56%, with global market expansion (E) and BSP as the primary transmitters. Figure A.4 (see below) graphically demonstrates that E dominantly transmits information across all edges.

FIGURE A.4 : Model 4 - Network analysis of dynamics between Brent spot and global oil market variables.

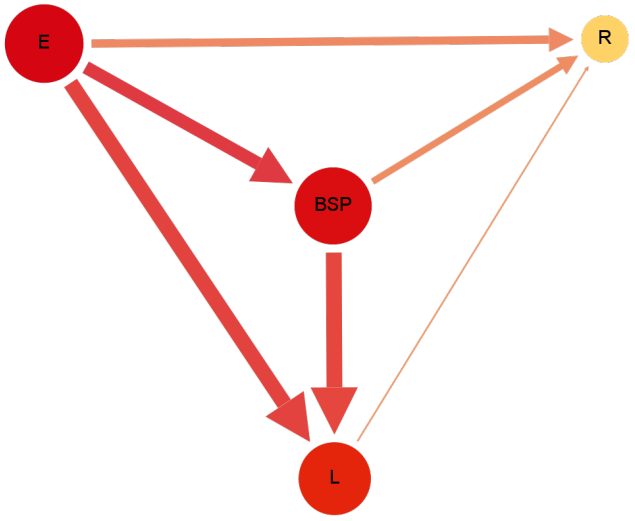

Note: Network analysis intra-market information exchange mechanism. Nodes notation represents: BE - Brent market expansion, BR - Brent market regime, BL - Brent market liquidity, BSP - Brent Spot.

Frequency Domain Spillovers.

TABLE A.8 : Model 4 - Frequency Decomposition of Static Directional Connectedness Between Brent and Global Oil Market.

(a) Short horizon - 1 to 5 days

Variable	BSP	E	L	R	FROM
BSP	0.00	0.01	0.00	0.00	0.01
E	0.00	0.00	0.00	0.00	0.01
L	0.02	0.06	0.42	0.12	0.20
R	0.02	0.05	0.04	0.04	0.11
TO	0.05	0.11	0.05	0.13	*TSI: 41.30%*

(b) Long horizon - 5 days to infinity

Variable	BSP	E	L	R	FROM
BSP	42.72	45.69	4.20	7.38	57.27
E	36.96	46.40	3.73	12.90	53.59
L	14.85	14.03	70.31	0.19	29.07
R	11.38	17.99	0.62	69.86	29.99
TO	63.19	77.71	8.55	20.47	*TSI: 42.56%*

Note: Frequency domain analysis following Baruník and Krehlík (2018) methodology. Short and Long horizons refer to '1 to 5 days' and '5 days to infinity'. Numbers report percentages.

The frequency domain results (Table A.8) indicate that long-run spillovers prevail, despite short-run market efficiency.

A.3.2.2 *Model 5 - WSP, E, L, R.*

Directional Spillovers.

TABLE A.9 : Model 5 - Static Directional Connectedness in WTI oil market represented through global oil market variables.

Variable	WSP	E	L	R	FROM
WSP	37.47	48.84	3.49	10.20	62.53
E	29.01	51.63	4.01	15.35	48.37
L	4.70	11.61	83.15	0.54	16.85
R	6.06	16.68	1.00	76.26	23.74
TO	39.78	77.13	8.50	26.09	*TSI: 37.88%*
NET	-22.76	28.76	-8.35	2.35	

Note: Following Diebold and Yilmaz (2012). Numbers are percentages. "TSI" stands for Total Spillover Index. $FROM$ - represents the measure of instrument's sensitivity to uncertainty generated by other variables in the system. TO - represents the measure of instrument's ability to transmit uncertainty to other variables in the system.

Table A.9 shows a total spillover index of 37.88%, with E as the predominant transmitter and WSP and L as net receivers. The network in Figure A.5 (see below) highlights

FIGURE A.5 : Model 5 - Network analysis of dynamics between WTI spot and global oil market variables.

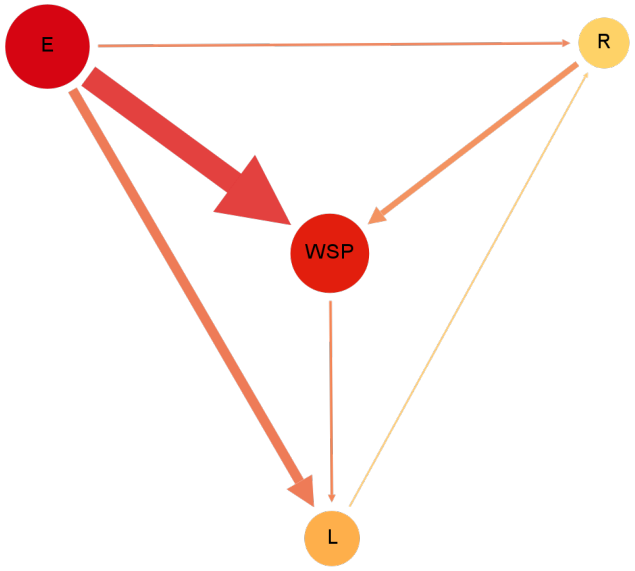

Note: Network analysis intra-market information exchange mechanism. Nodes notation represents: E - Global market expansion, R - Global market regime, L - Global market liquidity, WSP - WTI Spot.

the central role of E.

Frequency Domain Spillovers.

TABLE A.10 : Model 5 - Frequency Decomposition of Static Directional Connectedness Between WTI and Global Oil Market.

(a) Short horizon - 1 to 5 days

Variable	WSP	E	L	R	FROM
WSP	0.14	0.02	0.01	0.01	0.03
E	0.03	0.01	0.01	0.00	0.04
L	0.21	0.15	0.59	0.14	0.51
R	0.10	0.09	0.04	0.05	0.24
TO	0.34	0.26	0.06	0.15	*TSI: 51.03%*

(b) Long horizon - 5 days to infinity

Variable	WSP	E	L	R	FROM
WSP	37.33	48.82	3.49	10.20	62.50
E	28.98	51.62	4.00	15.35	48.33
L	4.49	11.46	82.56	0.39	16.35
R	5.96	16.59	0.95	76.20	23.51
TO	39.43	76.87	8.44	25.94	*TSI: 37.82%*

Note: Frequency domain analysis following Baruník and Krehlík (2018) methodology. Short and Long horizons refer to '1 to 5 days' and '5 days to infinity'. Numbers report percentages.

The frequency domain breakdown in Table A.10 mirrors previous models, with significant long-run spillovers.

A.3.2.3 *Model 6 - BSP, WSP, E, L, R.*

{*Directional Spillovers.*

TABLE A.11 : Model 6 - Static Directional Connectedness in Brent and WTI oil markets represented through global oil market variables.

Variable	BSP	WSP	E	L	R	FROM
BSP	33.93	18.64	37.63	3.63	6.18	66.07
WSP	24.48	28.46	37.06	2.70	7.30	71.54
E	28.06	20.30	37.64	3.09	10.91	62.36
L	12.69	4.35	11.78	70.89	0.30	29.11
R	9.81	5.26	15.83	0.64	68.47	31.53
TO	75.03	48.54	102.30	10.05	24.69	*TSI: 52.12%*
NET	8.97	-23.00	39.94	-19.06	-6.85	

Note: Following Diebold and Yilmaz (2012). Numbers are percentages. "TSI" stands for Total Spillover Index. *FROM* - represents the measure of instrument's sensitivity to uncertainty generated by other variables in the system. *TO* - represents the measure of instrument's ability to transmit uncertainty to other variables in the system.

FIGURE A.6 : Model 6 - Network analysis of dynamics between Brent and WTI spot and global oil market variables.

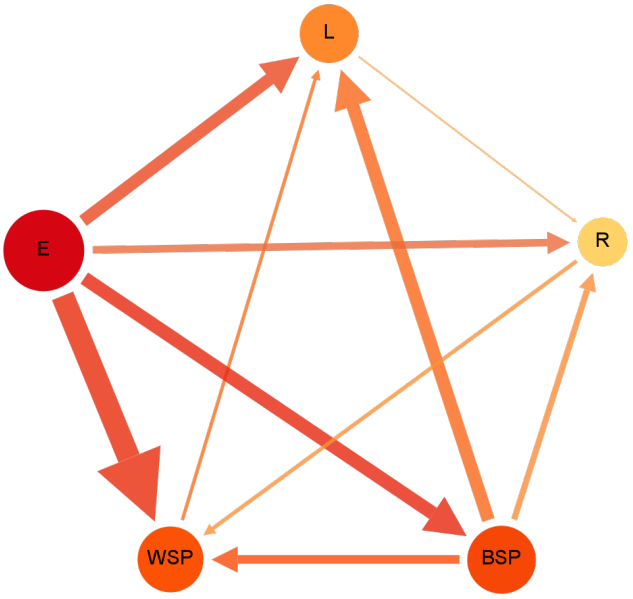

Note: Network analysis intra-market information exchange mechanism. Nodes notation represents: E - Global market expansion, R - Global market regime, L - Global market liquidity, BSP - Brent Spot, WSP - WTI Spot.

Table A.11 presents a total spillover index of 52.12%, slightly lower than Model 3, supporting a parsimonious approach. Dominant contributors are E and BSP, while market liquidity remains marginal. The network graph in Figure A.6 (see below) confirms the dominant role of E.

Frequency Domain Spillovers.

TABLE A.12 : Model 6 - Frequency Decomposition of Static Directional Connectedness Between Brent, WTI and Global Oil Market.

(a) Short horizon - 1 to 5 days

Variable	BSP	WSP	E	L	R	FROM
BSP	0.01	0.01	0.01	0.00	0.00	0.03
WSP	0.02	0.12	0.02	0.01	0.01	0.05
E	0.01	0.03	0.01	0.01	0.00	0.05
L	0.05	0.19	0.11	0.43	0.13	0.48
R	0.03	0.09	0.06	0.04	0.04	0.22
TO	0.11	0.32	0.20	0.06	0.14	*TSI: 57.80%*

(b) Long horizon - 5 days to infinity

Variable	BSP	WSP	E	L	R	FROM
BSP	33.92	18.63	37.62	3.62	6.17	66.04
WSP	24.46	28.34	37.04	2.69	7.29	71.49
E	28.05	20.27	37.64	3.08	10.91	62.31
L	12.63	4.16	11.66	70.46	0.17	28.63
R	9.78	5.17	15.77	0.60	68.43	31.31
TO	74.92	48.23	102.10	9.99	24.55	*TSI: 52.11%*

Note: Frequency domain analysis following Baruník and Krehlík (2018) methodology. Short and Long horizons refer to '1 to 5 days' and '5 days to infinity'. Numbers report percentages.

Table A.12 shows that the majority of spillover effects are concentrated over long horizons, with short-run market efficiency maintained by the corrective action of speculative traders.

A.4 Spillover Analysis for Model 6 vs Guo et al. dataset.

This section details the spillover analysis for the LSTM-based comparison between Model 6 and the dataset from Guo et al. (2023b). The tables break down the Total Spillover Index, directional, and net spillovers, highlighting differences in dynamic interdependencies.

Table A.13 presents spillover measures for Model 6, which integrates Brent and WTI spot prices with global market indicators, while Table A.14 shows the corresponding results for Guo et al. These findings underscore the superior interconnectedness of Model 6 and its implications for forecast accuracy.

TABLE A.13 : Static Spillover Analysis - Model 6 Variables.

	Variable	Brent	WTI	E	L	R	From Others
	Brent	31.09	20.18	27.80	5.34	15.58	68.91
	WTI	23.79	29.78	25.15	4.26	17.01	70.22
	E	26.09	20.15	29.53	3.54	20.69	70.47
	L	15.25	9.61	12.69	58.49	3.96	41.51
	R	18.88	18.13	26.42	1.17	35.41	64.59
To Others	Excluding Own	84.01	68.07	92.06	14.32	57.25	315.71
	Including Own	115.10	97.85	121.59	72.81	92.66	
NET	Excluding Own	15.10	-2.15	21.59	-27.19	-7.34	TSI: 63.14%
	Including Own	46.19	27.62	51.12	31.30	28.07	

Note: Static spillover analysis of Model 6 variables. Data span: from 26-03-2018 to 12-01-2023 of daily frequency.

TABLE A.14 : Static Spillover Analysis - Guo et al. (2023b) Variables.

	Variable	Brent	CF	EI	EPU	ER	INE	IR	SHCI	TR	WTI	From Others
	Brent	37.93	7.59	3.76	0.04	0.03	8.63	0.03	2.13	0.01	39.84	62.07
	CF	15.79	40.08	8.64	0.01	0.36	13.92	0.06	4.62	0.01	16.51	59.92
	EI	6.07	14.18	46.14	0.12	0.35	6.92	0.04	19.33	0.06	6.78	53.86
	EPU	4.47	0.96	0.87	82.36	0.06	1.46	1.73	3.71	0.71	3.66	17.64
	ER	1.99	2.84	2.00	0.06	77.63	1.18	0.03	12.64	0.33	1.30	22.37
	INE	29.24	10.50	4.47	0.02	0.32	22.08	0.09	1.60	0.11	31.57	77.92
	IR	0.59	0.47	0.09	3.08	0.48	0.06	94.18	0.27	0.31	0.47	5.82
	SHCI	1.89	5.27	18.86	0.37	1.44	1.03	0.17	68.77	0.02	2.18	31.23
	TR	0.17	0.71	0.27	0.22	0.02	2.43	20.26	0.12	75.27	0.53	24.73
	WTI	35.57	7.60	3.24	0.07	0.05	8.40	0.09	1.73	0.04	43.22	56.78
To Others	Excluding Own	95.79	50.12	42.21	3.99	3.12	44.03	22.51	46.13	1.60	102.84	412.34
	Including Own	133.72	90.20	88.35	86.35	80.74	66.11	116.69	114.91	76.86	146.06	
NET	Excluding Own	33.72	-9.80	-11.65	-13.65	-19.26	-33.89	16.69	14.91	-23.14	46.06	TSI: 41.23%
	Including Own	71.64	30.28	34.50	68.71	58.37	-11.81	110.87	83.68	52.13	89.27	

Note: Static spillover analysis of variables proposed by Guo et al. (2023b). Data span: from 26-03-2018 to 12-01-2023 of daily frequency.

Appendix B

Supplement to Chapter 3

The key to our methodology is measuring long memory—or dynamic persistence—in financial time series, which reflects how shocks decay over time. A slow, hyperbolic decay of the autocorrelation function (ACF) as the fractional differencing parameter d approaches 1 signals long memory. The operator

$$\Delta^d = (1 - L)^d$$

admits the binomial expansion

$$\Delta^d y_t \;=\; \sum_{k=0}^{\infty} \binom{d}{k} (-L)^k y_t, \qquad \binom{d}{k} = \frac{d(d-1)\cdots(d-k+1)}{k!}.$$

If d is integer, this series truncates and the ACF decays exponentially; for $0 \leq d \leq 1$ it decays hyperbolically, indicating long memory and affecting forecastability. In particular, $0 \leq d \leq 0.5$ yields covariance stationarity; $0.5 < d \leq 1$ implies mean reversion with weakening forecasts; and $d > 1$ implies explosiveness. Such fractionally integrated ($I(0)$–$I(1)$) processes are modelled by

$$\Delta^d y_t = (1 - L)^d y_t + \varepsilon_t, \quad \varepsilon_t \sim \text{iid}(0, \sigma^2).$$

To capture these dynamics fully, we turn to the FCVAR framework in the next section.

B.1 Fractional Cointegrated Vector Auto Regression Model

One approach for modelling systems of variables that are defined by interdependent dynamic and capture slow convergent memory is FCVAR framework. This framework was first proposed by Johansen (2008) and further extended by Johansen and Nielsen (2012b) as a generalisation of Johansen's cointegrated vector autoregressive model (CVAR) (Johansen, 1995a).

Before presenting the formal derivation of the FCVAR model, it is useful to clarify the notation used for fractional differencing and lag operators. The fractional integration parameter d governs the long-run persistence of the system, whereas the parameter b, with $0 \leq b \leq d$, determines the fractional degree of the error-correction mechanism. The case $b = 1$ yields the standard VECM as a special case of the FCVAR model.

Fractional differencing of order d is defined as $\Delta^d = (1 - L)^d$, with the corresponding fractional lag operator $L_d = 1 - \Delta^d$. Analogously, for the adjustment parameter b, we write $\Delta^b = (1 - L)^b$ and $L_b = 1 - \Delta^b$. These operators play complementary roles: d controls long-memory behaviour in levels, while b governs the smoothness of the equilibrium adjustment. This notation ensures that expressions such as $L_b \Delta^{d-b} X_t$ and $\Delta^d L_b^i X_t$ appearing in the FCVAR representation follow directly from the underlying fractional algebra.

In order to derive FCVAR the easiest approach is to consider CVAR model first. The CVAR model is essential for understanding the dynamics of multiple timeseries that are integrated of order one, $I(1)$, and are cointegrated. This model is defined as:

$$\Delta Y_t = \alpha \beta' Y_{t-1} + \sum_{i=1}^{k} \Gamma_i \Delta Y_{t-i} + \varepsilon = \alpha \beta' L^i Y_t + \varepsilon_t. \tag{B.1}$$

To obtain FCVAR model from the above equation, we can replace the difference and lag operators Δ and $L = 1 - \Delta$ by their fractional counterparts, Δ_b and $L_b = 1 - \Delta^b$

$$\Delta^b Y_t = \alpha \beta' L_b Y_t + \sum_{i=0}^{k} \Gamma_i \Delta^b L_b^i Y_t + \varepsilon_t, \tag{B.2}$$

Applying it to $Y_t = \Delta^{d-b} X_t$ it becomes:

$$\Delta^d X_t = \alpha \beta' L_b \Delta^{d-b} X_t + \sum_{i=o}^{k} \Gamma_i \Delta^d L_b^i X_t + \varepsilon_t \tag{B.3}$$

Where ε_t is an independent identically distributed (iid) $p - dimensional$ error term with $E(\varepsilon) = 0$ and $var(\varepsilon_t) = \sigma^2$.

Similarly to CVAR framework, parameters of FCVAR model have the following interpretations: α and β are $(p \times r)$ matrices with $0 \leq r \leq p$. Coefficients in α represent the rate of adjustment of each element in the system to the equilibria, β is a column vector of cointegrating relationships in the system. Stationary combinations or the long run equilibria is represented by the elements of $\beta' X_t$ and short run behaviour is governed by parameter Γ_i.

To mitigate the bias introduced by assuming fractional differences are defined through

an infinite series expansion, while actual data samples comprise a finite number of observations, an additional modification is required. Johansen and Nielsen suggested incorporating a level parameter μ that shifts each series by a constant (Johansen and Nielsen, 2016). This adjustment allows to account for the finite sample size and improve accuracy in capturing the dynamics of long memory processes. Applying this modification, equation becomes:

$$\Delta^d(X_t-\mu)=\alpha\beta' L_d(X_t-\mu)+\sum_{i=0}^{k}\Gamma_i\Delta^d L_d^i(X_t-\mu)+\varepsilon_t,\qquad t=1,2,\cdots,T \tag{B.4}$$

Effectively, this addition adjusts the data series for a constant shift μ before applying fractional differencing and truncating it at $k=t-1$.

B.1.1 Forecasting Using FCVAR

Due to the autoregressive nature of FCVAR model, it's forecasting methodology is relatively straight forward. To demonstrate this aspect, consider the following model:

$$\Delta^d(X_{t+1}-\mu)=X_{t+1}-\mu-L_d(X_{t+1}-\mu) \tag{B.5}$$

Rearranging for X_{t+1} gives the following relationship:

$$X_{t+1}=\mu+L_d(X_{t+1}-\mu)+\alpha\beta' L_b\Delta^{d-b}(X_{t+1}-\mu)+\sum_{i=o}^{k}\Gamma_i\Delta^d L_b^i(X_{t+1}-\mu)+\varepsilon_{t+1} \tag{B.6}$$

Given that $L_b=1-\Delta^b$ functions as a lag operator, making $L_b^i X_{t+1}$ known at the time t for all $i\geq 1$, Equation B.6 serves as the basis for FCVAR forecasting. Expressing the relationship using coefficient estimates and conditional expectation notation based on information available at time t, a one-step-ahead forecast, $\widehat{X}_{t+1|t}$ is defined as:

$$\widehat{X}_{t+1|t}=\widehat{\mu}+L_{\hat{d}}(X_{t+1}-\widehat{\mu})+\widehat{\alpha}\widehat{\beta}' L_{\hat{b}}\Delta^{\hat{d}-\hat{b}}(X_{t+1}-\widehat{\mu})+\sum_{i=o}^{k}\widehat{\Gamma}_i\Delta^{\hat{d}} L_{\hat{b}}^i(X_{t+1}-\widehat{\mu}) \tag{B.7}$$

This relationship allows the FCVAR model to generate estimates for every variable within the system at each forecasting iteration. Thus, n-step ahead forecast can be formed recursively based on generalisation of Equation B.7.

$$\widehat{X}_{t+j|t}=\widehat{\mu}+L_{\hat{d}}(X_{t+j|t}-\widehat{\mu})+\widehat{\alpha}\widehat{\beta}' L_{\hat{b}}\Delta^{\hat{d}-\hat{b}}(X_{t+j|t}-\widehat{\mu})+\sum_{i=o}^{k}\widehat{\Gamma}_i\Delta^{\hat{d}} L_{\hat{b}}^i(X_{t+j|t}-\widehat{\mu}) \tag{B.8}$$

Where $\widehat{X}_{s|t}=X_s$ for $s\leq t$. In this case the forecasts are calculated recursively from Equation B.8 for $j=1,2,...,h$ to generate h-step ahead forecast, $\widehat{X}_{t+j|t}$.

B.2 Rolling Window Size Estimation

Accurately estimating the optimal rolling window size is crucial for improving forecast accuracy and ensuring robust model performance. The academic literature provides a variety of approaches for rolling window estimation, including techniques such as Exponential Smoothing, Moving Average, GARCH-based models, and Dynamic Conditional Correlation (DCC) models. These methods primarily focus on improving forecast precision by accounting for structural breaks, volatility dynamics, and correlations in multivariate series. However, a key concern arises regarding the potential misalignment between the model used for window size selection and the model used in actual analysis. If the rolling window selection model and the actual analysis model diverge in their handling of time series processes, this misalignment can negatively impact forecasting accuracy and lead to misleading conclusions.

The issue that can arise from such a misalignment is the inaccurate representation of market dynamics. For example, if a window estimation model targets volatility clustering or shifts in correlations, but the actual analysis model is more focused on capturing equilibrium relationships, the selected window might not capture the critical dynamics relevant to the forecasting model, and, thus, introduce unintended bias. To address this concern, we develop a custom approach to estimate the optimal rolling-window size. Implemented in R, this approach dynamically identifies the most suitable window by evaluating the forecast accuracy across a range of window sizes within the context of a Vector Error Correction Model (VECM), which aligns well with the principles of the FCVAR methodology.

A natural concern is whether a VECM-based optimisation is appropriate given that the primary analysis relies on the FCVAR framework. This choice is deliberate and follows the theoretical nesting relationship between the two models. As shown in Johansen and Nielsen (2012b) and Johansen and Nielsen (2016), the standard VECM is obtained as the special case of the FCVAR model when the fractional integration parameter d equals one. Both models therefore share the same cointegrating structure, adjustment coefficients, and short-run dynamics, differing only in the degree of fractional persistence. Because FCVAR estimation is computationally intensive, especially when repeated thousands of times in rolling window optimisations, the VECM provides a tractable and statistically coherent preliminary filter for determining window size. This approach ensures that the selected window preserves stable long-run relations and coherent short-run dynamics, which in turn supports reliable FCVAR estimation.

The VECM-based algorithm aims to minimise the average Mean Squared Error (MSE) of one-day-ahead out-of-sample forecasts across various window sizes, ranging from 50 to 890 days. For each window size, the algorithm identifies the optimal lag using the Akaike Information Criterion (AIC) and determines the optimal rank by applying the Johansen cointegration test (Johansen, 1988) to fit models and analyse residuals. After

establishing the optimal lag and rank, the algorithm estimates the VECM for each chosen window size and generates a one-day-ahead out-of-sample forecast, recording the MSE for each window. This process is repeated for each window size up to the specified maximum, moving sequentially along the time series to the last observation. By cycling through all possible windows within the available data, the algorithm aims to find the window size that yields the most accurate predictions. Then it calculates the minimum MSE across all tested windows and reports the window size associated with this minimal value.

B.3 Variable Interdependence Evaluation Methods

To validate the results observed from examining the error profiles, we perform Robustness checks using VAR Granger Causality/Wald Block Exogeneity Tests. This approach allows to identify and cross validate the interaction mechanisms between forward and backward memory parameters and their impact on spot price movements.

B.3.1 VAR Granger Causality & Wald Block Exogeneity Tests

These tests arise from the collective developments in econometric methodologies developed by Wald (1943); Granger (1969); Sims (1980). This framework extends the Wald test within the VAR framework to assess whether a group of variables does not Granger-cause another variable or block of variables in the model.

In this context, a VAR model is central to understanding the linear interdependencies among multiple timeseries. For a system of k variables, the VAR model of order p is described as:

$$Y_t = C + A_1 Y_{t-1} + A_2 Y_{t-2} + \ldots + A_p Y_{t-p} + E_t, \tag{B.9}$$

Where Y_t is a $k \times 1$ vector of endogenous variables at time t, C represents a $k \times 1$ vector of intercepts, $A_1, A_2, ..., A_p$ are $k \times k$ coefficient matrices for the lagged values, and E_t is a $k \times 1$ vector of error terms, representing white noise.

Granger causality tests if data generating process reflected through one timeseries can predict another. Formally, a timeseries X is said to Granger cause Y if past values of X contain predictive information for Y above the level of the information contained in past values of Y alone. This is assessed by comparing a restricted model, excluding X, with an unrestricted model that includes X. For example, mathematical formulation for two timeseries Y_t and X_t, first requires construction of restricted model without X_t component:

$$Y_t = C_1 + A_{11} Y_{t-1} + \ldots + A_{1p} Y_{t-p} + E_{1t}, \tag{B.10}$$

Followed by unrestricted model that includes X_t component:

$$Y_t = C_2 + A_{21}Y_{t-1} + \ldots + A_{2p}Y_{t-p} + B_{21}X_{t-1} + \ldots + B_{2p}X_{t-p} + E_{2t}, \tag{B.11}$$

An F-test is then used to compare both models under the null hypothesis that coefficients of all X_t lags in the unrestricted model are zero, implying X_t does not Granger cause Y_t. Therefore, rejection of null hypothesis indicates Granger causality.

The VAR method is specifically designed for analysing stationary processes that are not cointegrated. However, applying this approach directly to non stationary and cointegrated timeseries can lead to spurious results. To address this issue, we use the VECM, which can be expressed in terms of VAR representation. This approach preserves the full spectrum of the VECM properties, including the long-term equilibrium relationships among cointegrated variables, allowing for their analysis within the VAR framework.

The connection between VECM and VAR models can be mathematically expressed by integrating the error correction term into the VAR model, transforming it into a VECM. To demonstrate this aspect, consider a VECM of order $p-1$:

$$\Delta Y_t = \alpha\beta' Y_{t-1} + \sum_{i=1}^{p-1} \Gamma_i \Delta Y_{t-i} + \epsilon_t, \tag{B.12}$$

Where ΔY_t represents the difference of Y_t, α and β' represent the adjustment coefficients and cointegration matrix which jointly represent the long-term relationship, Γ_i are short-term dynamics coefficients, and ϵ_t is the error term.

Transforming this VECM into a VAR model of order p involves expressing the cointegrated relationships within the VAR framework:

$$Y_t = A_1 Y_{t-1} + \cdots + A_p Y_{t-p} + \mu + \epsilon_t, \tag{B.13}$$

Where $A_1, \ldots, A_p$ are the coefficients matrices, capturing both short-term dynamics and long-term equilibria derived from the VECM, and μ is a constant term.

This integration closes the gap between short-term fluctuations and long-term equilibrium, ensuring that the VAR model accounts for both stationary and non-stationary cointegrated timeseries.

B.3.2 Wald Block Exogeneity Tests

Associated within the VAR framework, these tests aim to evaluate the hypothesis that one or more variables do not Granger cause other variables in the system. This approach checks if all lagged coefficients of the exogenous variables are jointly zero. The Wald

statistic is then used to establish block exogeneity. Wald statistic is formulated as:

$$W = (R\hat{\beta} - r)'[R(\hat{V}_\beta)R']^{-1}(R\hat{\beta} - r), \tag{B.14}$$

where R is a matrix that selects the coefficients being tested from $\hat{\beta}$ which represents the vector of estimated coefficients, $\hat{V}_\beta$ is the covariance matrix of $\hat{\beta}$, and r is the vector under the null hypothesis. Wald statistic (W) follows a χ^2 distribution with degrees of freedom equal to the number of restrictions.

B.4 Loss Functions

The forecasting accuracy in this study quantified using three distinct loss functions: (1) the Linear Exponential Loss (LinExp Loss); (2) Root Mean Squared Error (RMSE); and (3) Tracking Signal loss function.

(1) The LinExp Loss function is calibrated to highlight the sensitivity of the model to overestimations, reinforcing the importance of adaptable modelling techniques in financial analyses. This loss function is defined as:

$$L(y, f) = e^{a(y-f)} - a(y - f) - 1 \tag{B.15}$$

Where $L(y, f)$ denotes the LinExp loss for an actual value y and forecasted value f, with the asymmetry parameter a. We set parameter a = -1 to emphasise the loss function sensitivity to overestimations. Specifically, for overestimations where $f > y$, the term $e^{-1\cdot(y-f)}$ increases exponentially and dominates liner decrease in $-1(y - f) - 1$, penalising these errors more severely. In contrast, in the case of underestimations where $y > f$, the loss grows, but the exponential term decreases, resulting in a more moderated increase due to the linear component.

This asymmetry in the loss function is particularly useful in volatile markets, such as oil and cryptocurrency markets, where the cost of overestimating values can lead to adverse financial impacts, such as overpaying for assets or holding onto investments where expected returns may not materialise.

(2) Forecasts are also assessed using RMSE, which provides a more interpretative measure of the average prediction error between the actual and forecasted values. The formula for RMSE is given as:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - f_i)^2} \tag{B.16}$$

Where n represents the number of forecasts, y_i denotes the actual value, and f_i is the forecasted value at instance i. This metric is widely used due to its ability to quantify the magnitude of forecast error, providing a straightforward interpretation of model performance in terms of average error per forecast. Unlike the LinExp loss function, which can be adjusted to particularly emphasise sensitivity to either underpredictions or overpredictions through its asymmetry parameter a, the RMSE treats all percentage errors symmetrically. This means that it does not automatically penalise underpredictions or overpredictions. However, its sensitivity to large errors due to its quadratic nature means that it effectively highlights significant deviations between forecasted and actual values.

(3) The Tracking Signal (TS) loss function is employed as a diagnostic tool to identify systematic forecast bias over time. It is calculated using the formula:

$$TS = \frac{\sum_{i=1}^{n}(y_i - f_i)}{MAE} \tag{B.17}$$

Where MAE (Mean Absolute Error) is the average of the absolute differences between the forecasted values and the actual values, calculated as:

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - f_i| \tag{B.18}$$

The Tracking Signal measures the ratio of the cumulative forecast errors to the MAE. With bound of acceptance laying between values of -4 to 4, this measure provides insight into whether the forecast bias is drifting over time. A TS close to zero indicates that any forecast errors are random rather than systematic. However, a persistently positive or negative TS suggests a systematic bias in the forecasts, either consistently overestimating or underestimating the actual values, respectively. The practical application of this metric is in the calibration of forecast models to reduce bias and improve accuracy over time.

Appendix C

Supplement to Chapter 4

C.1 Japan's market analysis

TABLE C.1 : Asset Universe Returns – Descriptive Statistics (Japan Market).

Ticker	Mean	Median	Max	Min	Std. Dev.	Skew	Kurt	Jarque–Bera
N225	0.040%	0.072%	0.1023	-0.1240	0.0133	-0.3813	9.93	7,423.0
RFR	0.002%	0.001%	0.0001	-0.0000	0.0000	0.4439	1.98	278.8
1605.T	0.038%	0.000%	0.1662	-0.1542	0.0220	0.0366	7.30	2,833.6
2502.T	0.050%	0.017%	0.0973	-0.1240	0.0162	-0.0780	7.52	3,128.3
2503.T	0.033%	0.043%	0.1185	-0.1068	0.0155	-0.0459	8.60	4,797.3
2801.T	0.072%	0.052%	0.2078	-0.1750	0.0190	0.0767	13.39	16,505.8
2802.T	0.065%	0.038%	0.1135	-0.1399	0.0162	0.1011	9.58	6,625.6
2914.T	0.053%	0.000%	0.1248	-0.1683	0.0154	0.1842	11.47	10,990.8
3401.T	0.020%	0.000%	0.1701	-0.1180	0.0181	0.1413	8.84	5,230.4
3402.T	0.039%	0.000%	0.1684	-0.1325	0.0177	0.1164	10.88	9,512.5
4005.T	0.024%	0.000%	0.1040	-0.1000	0.0203	-0.0326	5.40	878.5
4452.T	0.047%	0.003%	0.0785	-0.1052	0.0148	-0.0525	6.85	2,267.9
4502.T	0.029%	0.025%	0.1040	-0.0895	0.0137	-0.1356	8.79	5,142.4
4503.T	0.043%	0.024%	0.1005	-0.0881	0.0157	0.0844	6.22	1,592.8
4519.T	0.090%	0.000%	0.1538	-0.1396	0.0185	0.4935	9.02	5,693.2
4523.T	0.038%	0.000%	0.2070	-0.2364	0.0210	0.8134	26.26	83,111.0
4568.T	0.078%	0.022%	0.1452	-0.1481	0.0197	0.4173	8.64	4,972.9
4689.T	0.041%	0.000%	0.1795	-0.1444	0.0219	0.6411	11.74	11,934.5
4755.T	0.033%	0.000%	0.2410	-0.1585	0.0230	0.4402	9.50	6,583.1
5019.T	0.059%	0.000%	0.1829	-0.1117	0.0201	0.3184	7.93	3,771.1
5020.T	0.044%	0.000%	0.1247	-0.1511	0.0187	0.0453	8.59	4,781.3
5201.T	0.027%	0.000%	0.1337	-0.1322	0.0187	-0.0663	8.39	4,440.7
5202.T	-0.016%	0.000%	0.1800	-0.1860	0.0265	0.3619	8.64	4,949.4
5411.T	0.018%	0.000%	0.1451	-0.1372	0.0235	0.2082	6.14	1,535.3
6301.T	0.052%	0.041%	0.1152	-0.1060	0.0197	0.1181	5.39	879.7
6367.T	0.063%	0.042%	0.0880	-0.1467	0.0196	-0.1244	5.85	1,255.3
6501.T	0.089%	0.000%	0.1662	-0.1619	0.0206	0.1150	8.03	3,869.3
6701.T	0.071%	0.000%	0.1838	-0.1741	0.0206	0.0874	12.98	15,242.4
6758.T	0.068%	0.036%	0.1241	-0.1279	0.0214	0.2381	6.76	2,197.0
6861.T	0.084%	0.017%	0.1766	-0.1210	0.0188	0.3860	7.66	3,406.4
6869.T	0.074%	0.000%	0.1543	-0.1535	0.0209	-0.0841	10.16	7,847.3
6954.T	0.046%	0.024%	0.1321	-0.1272	0.0202	0.0502	6.09	1,459.5
7011.T	0.081%	0.000%	0.1984	-0.1496	0.0205	0.4213	8.90	5,430.8
7012.T	0.274%	0.000%	8.6514	-0.9041	0.1456	57.1128	3,397.96	1,763,998,540.9
7203.T	0.058%	0.000%	0.1281	-0.1366	0.0168	0.1681	7.68	3,363.0
7261.T	0.027%	0.000%	0.1852	-0.1679	0.0249	0.2171	6.60	2,015.0
7267.T	0.035%	0.000%	0.1470	-0.1777	0.0186	0.0907	8.50	4,621.6
7270.T	0.079%	0.018%	0.1291	-0.1834	0.0218	0.0905	6.97	2,410.8
7751.T	0.029%	0.051%	0.0982	-0.1346	0.0157	-0.1232	8.48	4,605.3
7974.T	0.067%	0.073%	0.4038	-0.1772	0.0236	1.6025	32.83	137,587.1
8001.T	0.085%	0.074%	0.1219	-0.1453	0.0164	-0.0183	8.06	3,911.7
8031.T	0.064%	0.050%	0.1043	-0.1987	0.0171	-0.3522	10.65	9,017.0
8035.T	0.100%	0.076%	0.1803	-0.1848	0.0237	0.0735	7.78	3,502.5
8058.T	0.061%	0.023%	0.1043	-0.1411	0.0167	0.1303	7.70	3,383.7
8306.T	0.067%	0.000%	0.1119	-0.1784	0.0184	-0.0957	8.81	5,167.4
8316.T	0.056%	0.000%	0.1037	-0.1552	0.0173	-0.0262	8.18	4,109.9
8411.T	0.049%	0.000%	0.0929	-0.1971	0.0166	-0.4683	13.28	16,299.3
8473.T	0.063%	0.000%	0.2806	-0.1959	0.0253	0.7280	14.19	19,466.7
8604.T	0.041%	-0.015%	0.1254	-0.1859	0.0212	-0.1420	9.52	6,507.5
8628.T	0.034%	0.000%	0.1074	-0.1595	0.0168	-0.1976	9.14	5,783.4
9101.T	0.077%	0.000%	0.1343	-0.1301	0.0223	0.1160	5.78	1,191.4
9432.T	0.058%	0.000%	0.0727	-0.1474	0.0137	-0.0533	9.54	6,537.4
9433.T	0.065%	0.023%	0.1045	-0.1615	0.0159	-0.1948	10.34	8,257.1
9501.T	0.006%	-0.083%	0.3581	-0.2762	0.0333	1.2590	23.04	62,346.9
9502.T	0.017%	0.000%	0.1345	-0.1107	0.0178	0.0968	7.93	3,726.1
9503.T	0.027%	0.000%	0.1765	-0.1846	0.0219	0.1693	11.02	9,860.6
9531.T	0.044%	0.018%	0.1296	-0.1165	0.0157	0.0244	9.16	5,797.8
9532.T	0.038%	0.000%	0.1002	-0.1033	0.0142	0.1821	7.38	2,947.1
9537.T	0.026%	0.000%	0.0968	-0.0695	0.0105	0.5304	11.26	10,610.8
9613.T	0.064%	0.036%	0.1717	-0.1743	0.0194	-0.0051	11.26	10,424.7
9766.T	0.092%	0.010%	0.1439	-0.1888	0.0217	-0.0058	8.72	4,997.5
9984.T	0.083%	0.037%	0.1895	-0.1866	0.0248	0.1360	9.92	7,339.7

Note: This table presents the descriptive statistics of the asset universe in the Japan market, including the Nikkei 225 index (N225) and the RFR (10-Year Bond Yield). Period 04-01-2010 to 03-04-2025 (3,669 obs.).

TABLE C.2 : Asset Universe Memory - Descriptive Statistics (Japan Market).

Ticker	Mean	Median	Max.	Min.	Std. Dev.	Skew.	Kurt.	Jarque-Bera
N225	-0.001	-0.003	1.4052	-0.7227	0.2119	0.1053	3.95	145.1
1605.T	-0.028	-0.037	1.1460	-0.9972	0.2338	0.2941	3.80	150.1
2502.T	-0.041	-0.033	0.9540	-0.9060	0.2249	-0.2073	3.39	49.1
2503.T	0.019	0.024	0.8046	-1.1034	0.2446	-0.3379	3.81	170.8
2801.T	-0.037	-0.046	0.9945	-1.2800	0.2447	0.0821	3.49	40.2
2802.T	-0.066	-0.044	0.8758	-1.0707	0.2448	-0.2268	3.13	34.1
2914.T	-0.026	-0.009	0.8279	-1.0383	0.2168	-0.0753	3.53	46.6
3401.T	-0.027	-0.022	0.9612	-1.2359	0.2352	-0.1681	3.54	62.2
3402.T	-0.083	-0.087	1.0576	-1.4381	0.2825	-0.1310	3.70	85.7
4005.T	-0.011	-0.002	1.0099	-1.2255	0.2627	-0.2450	4.02	196.8
4452.T	-0.043	-0.020	0.7432	-1.1257	0.2380	-0.4092	3.89	222.7
4502.T	-0.012	0.007	0.9740	-0.8864	0.2307	-0.1090	3.07	7.9
4503.T	-0.029	0.001	0.9840	-1.2131	0.2407	-0.6895	4.34	566.8
4519.T	0.005	0.007	0.9541	-0.7874	0.2224	0.0243	3.19	5.9
4523.T	-0.017	-0.020	1.7323	-1.1250	0.2444	-0.0968	5.81	1,213.8
4568.T	0.044	0.035	1.5739	-1.3098	0.2637	-0.1993	6.07	1,467.5
4689.T	-0.066	-0.049	0.7149	-0.9952	0.2275	-0.2665	3.43	71.6
4755.T	-0.040	-0.040	1.0425	-1.0480	0.2507	0.2852	3.57	99.4
5019.T	-0.043	-0.029	1.1258	-0.9694	0.2424	-0.1593	3.90	139.4
5020.T	-0.035	-0.035	0.8834	-1.6036	0.2476	0.0874	3.68	75.1
5201.T	-0.008	0.015	1.9688	-1.0263	0.2420	-0.3013	5.04	692.3
5202.T	-0.007	-0.006	0.8878	-0.9579	0.2134	-0.0778	3.86	117.7
5411.T	0.059	0.048	0.8025	-0.6147	0.2201	0.1717	3.35	37.1
6301.T	0.015	0.029	1.5383	-1.2064	0.2720	-0.1770	3.72	98.6
6367.T	-0.004	-0.013	0.7747	-0.9918	0.2352	-0.0250	3.08	1.3
6501.T	-0.028	-0.029	0.7218	-1.0225	0.2168	-0.2696	4.07	219.2
6701.T	-0.023	-0.022	0.8860	-1.3158	0.2374	-0.1317	3.99	161.4
6758.T	-0.063	-0.055	0.6961	-1.1108	0.2209	-0.3180	3.90	185.8
6861.T	-0.019	-0.008	0.8830	-1.3242	0.2507	-0.2585	4.24	275.2
6869.T	-0.043	-0.028	0.8610	-1.1386	0.2381	-0.3705	4.07	259.4
6954.T	-0.003	-0.001	1.1846	-1.1075	0.2118	-0.0575	3.84	109.0
7011.T	0.021	0.012	1.1287	-1.1106	0.2521	-0.0815	4.07	178.1
7012.T	0.047	0.040	1.1628	-1.0042	0.2250	0.1774	4.02	178.3
7203.T	-0.011	0.008	0.9259	-1.4171	0.2411	-0.3977	3.80	195.2
7261.T	0.025	0.020	1.0920	-1.3212	0.2050	0.1611	4.29	270.3
7267.T	0.003	0.008	1.0000	-1.4215	0.2421	-0.1609	3.76	104.2
7270.T	-0.011	-0.007	0.9215	-1.4227	0.2189	-0.3419	4.63	479.6
7751.T	-0.047	-0.031	1.1724	-1.0304	0.2612	-0.1773	3.37	40.5
7974.T	0.023	0.016	1.0692	-0.9517	0.2368	0.5013	4.69	588.9
8001.T	-0.003	0.005	1.2788	-1.2309	0.2432	-0.2713	4.50	387.9
8031.T	0.009	0.022	1.1932	-0.8493	0.2273	-0.1100	3.64	70.0
8035.T	0.013	0.026	0.9613	-0.9412	0.2333	-0.4179	3.78	200.4
8058.T	-0.016	-0.018	0.9615	-0.8448	0.2255	0.0935	3.42	31.7
8306.T	-0.014	-0.018	1.0360	-1.2053	0.2260	-0.1457	3.39	36.3
8316.T	-0.007	0.000	1.2146	-0.9916	0.2291	-0.2270	3.76	119.7
8411.T	-0.000	0.003	0.7901	-1.0643	0.2446	-0.3012	3.38	77.6
8473.T	-0.001	-0.024	1.1387	-1.0073	0.2456	0.4425	3.48	155.6
8604.T	0.026	0.039	0.9131	-1.1224	0.2116	-0.2710	3.40	68.8
8628.T	-0.023	-0.000	1.1737	-1.3712	0.2843	-0.6503	4.70	702.3
9101.T	0.051	0.027	0.9557	-1.0223	0.2463	0.2269	3.22	39.0
9432.T	-0.036	-0.046	0.9165	-1.3319	0.2320	-0.1792	3.88	137.9
9433.T	-0.034	-0.038	1.1220	-1.3197	0.2475	-0.2533	4.44	357.4
9501.T	0.044	0.053	1.2153	-1.0682	0.2513	-0.0241	3.86	114.6
9502.T	-0.024	-0.009	1.4546	-1.1449	0.2715	-0.1404	3.39	35.5
9503.T	-0.018	-0.017	1.2079	-1.0098	0.2584	0.0353	3.60	55.2
9531.T	-0.072	-0.072	1.0861	-1.2020	0.2524	-0.2492	4.40	335.6
9532.T	-0.088	-0.066	0.7207	-1.1427	0.2526	-0.3583	3.30	92.4
9537.T	-0.199	-0.185	0.6714	-1.3480	0.2909	-0.2716	3.21	51.9
9613.T	0.004	0.019	1.6291	-1.0016	0.2258	-0.3272	4.81	567.7
9766.T	-0.025	-0.024	1.0046	-0.9852	0.2340	-0.1607	3.51	55.1
9984.T	0.004	0.022	1.0717	-1.1704	0.2370	-0.2924	3.73	133.2

Note: This table presents the descriptive statistics of the memory of individual assets in the asset universe, estimated from their return series. Memory estimated using FELW and rolling window of 60 observations. The table includes the memory of Nikkei 225 index (N225) which represents Japan Market in this study. Period: 04-01-2010 to 03-04-2025 (3,669 obs.).

TABLE C.3 : Quintile and Market Portfolio Returns – Descriptive Statistics (Japan Market).

Portfolio	Mean	Median	Max.	Min.	Std. Dev.	Skew.	Kurt.	Jarque–Bera
Q1	0.057%	0.064%	8.2%	-13.0%	0.012	-0.381	11.50	11 125.3
Q2	0.057%	0.065%	69.3%	-12.6%	0.017	16.929	685.5	71,394,502
Q3	0.061%	0.077%	8.2%	-11.8%	0.013	-0.345	8.82	5 245.8
Q4	0.045%	0.073%	9.7%	-9.4%	0.013	-0.178	8.32	4 348.9
Q5	0.057%	0.078%	9.1%	-12.3%	0.013	-0.334	9.34	6 204.9
MKT	0.040%	0.072%	10.2%	-12.4%	0.013	-0.381	9.93	7 423.0

Note: This table reports the daily return summary for quintile portfolios (Q1–Q5) sorted by the cross-sectional memory parameter, and for the benchmark market portfolio (MKT), in the Japan market. Memory parameters were estimated using a 60-day rolling window on each asset's return series; Q1 contains assets with the lowest memory values, Q5 the highest. The sample runs from 04-01-2010 to 03-04-2025, encompassing 3,669 trading days. Skewness, kurtosis, and Jarque–Bera statistics test for deviations from normality. Higher JB values indicate stronger non-normality in the return distribution.

TABLE C.4 : Quintile and Market Portfolio Memory – Descriptive Statistics (Japan Market).

Portfolio	Mean	Median	Max.	Min.	Std. Dev.	Skew.	Kurt.	Jarque–Bera
Q1	-0.335	-0.339	0.0307	-0.714	0.106	0.194	3.30	37.2
Q2	-0.120	-0.122	0.306	-0.460	0.096	0.412	3.56	151.4
Q3	-0.006	-0.008	0.441	-0.365	0.095	0.361	3.30	93.2
Q4	0.099	0.097	0.551	-0.203	0.097	0.390	3.28	105.1
Q5	0.277	0.276	1.087	-0.067	0.1082	0.565	4.30	452.6
MKT	-0.001	-0.003	1.405	-0.723	0.212	0.105	3.95	145.1

Note: This table reports the daily summary statistics of the estimated memory parameter for quintile-sorted portfolios (Q1–Q5) and the market portfolio (MKT) in the Japan market. Memory is estimated via a 30-day rolling window on each asset's return series; Q1 consists of assets with the lowest memory values, Q5 the highest. The sample period runs from 04-01-2010 to 03-04-2025 (3,669 obs.). Skewness, kurtosis, and the Jarque–Bera test assess deviations from normality in the memory distribution.

TABLE C.5 : HML CAPM Extension – GRS Test Results (Japan Market).

Factor	**HML**	**Q1**		**Q2**		**Q3**		**Q4**		**Q5**	
Period	**Interval**	**tStat**	**Pval**	**tStat**	**Pval**	**tStat**	**Pval**	**tStat**	**Pval**	**tStat**	**Pval**
Full	2010-04-01 to 2025-04-03	2.986	0.01	3.016	0.01	2.621	0.02	1.501	0.19	1.235	0.29
1	2010-04-01 to 2012-03-31	1.640	0.15	1.380	0.23	1.123	0.35	0.162	0.98	1.344	0.24
2	2012-04-01 to 2014-03-31	0.309	0.91	0.297	0.91	0.354	0.88	0.641	0.67	0.894	0.49
3	2014-04-01 to 2016-03-31	0.762	0.58	1.373	0.23	1.753	0.12	1.409	0.21	1.038	0.39
4	2016-04-01 to 2018-03-31	2.062	0.07	0.676	0.64	0.234	0.95	0.429	0.83	0.309	0.91
5	2018-04-01 to 2020-03-31	3.358	0.01	4.037	0.00	0.750	0.59	1.227	0.30	1.011	0.41
6	2020-04-01 to 2022-03-31	1.411	0.22	0.595	0.70	1.013	0.41	2.443	0.03	2.168	0.06
7	2022-04-01 to 2024-03-31	0.334	0.89	0.536	0.75	1.325	0.25	0.572	0.72	0.462	0.81
8	2024-04-01 to 2025-04-03	0.920	0.47	0.507	0.77	1.168	0.33	0.776	0.57	1.493	0.19
Train	2010-04-01 to 2022-03-23	2.495	0.03	2.822	0.02	2.210	0.05	1.281	0.27	1.307	0.26
Test	2022-03-23 to 2025-04-03	0.689	0.63	0.968	0.44	2.325	0.04	1.028	0.40	0.599	0.70

Note: This table reports GRS joint tests of CAPM intercepts augmented with the memory HML factor for quintile portfolios Q1–Q5. "Full" spans 04-01-2010 to 03-04-2025 (3,662 observations). Rows 1–8 cover consecutive two-year intervals, while out-of-sample test section split into "Train" (80%) part which runs 04-01-2010 to 23-03-2022 and "Test" (20%) part runs 23-03-2022 to 03-04-2025. Low p-values signal significant non-zero alphas.

TABLE C.6 : Conventional vs Extended CAPM Out-Of-Sample Forecast Performance (Japan Market).

Model	**Metric**	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**
Extended	MSE	0.00003	0.00004	0.00004	0.00004	0.00003
	MAE	0.00400	0.00487	0.00504	0.00507	0.00400
	RMSE	0.00522	0.00640	0.00657	0.00667	0.00522
Standard	MSE	0.00004	0.00004	0.00004	0.00005	0.00005
	MAE	0.00477	0.00485	0.00504	0.00509	0.00499
	RMSE	0.00617	0.00635	0.00657	0.00672	0.00673
RFP	MSE	–28.31%	1.49%	0.09%	–1.35%	–39.64%
	MAE	–16.02%	0.40%	–0.03%	–0.45%	–19.83%
	RMSE	–15.33%	0.74%	0.04%	–0.68%	–22.31%

Note: This table compares out-of-sample forecast accuracy of the extended CAPM (including the memory-based HML factor) against the standard CAPM for five memory-sorted portfolios (Q1–Q5) in the Japan market. Forecast errors are measured by mean squared error (MSE), mean absolute error (MAE), and root mean squared error (RMSE) on a 20% hold-out sample (04-01-2010 and 03-04-2025). "Relative" shows the percentage change in each metric for the extended model versus the standard model (negative values indicate improvement).

FIGURE C.1 : Cumulative Returns of Frictionless, Equally Weighted Portfolios under Daily Rebalancing (Japan Market).

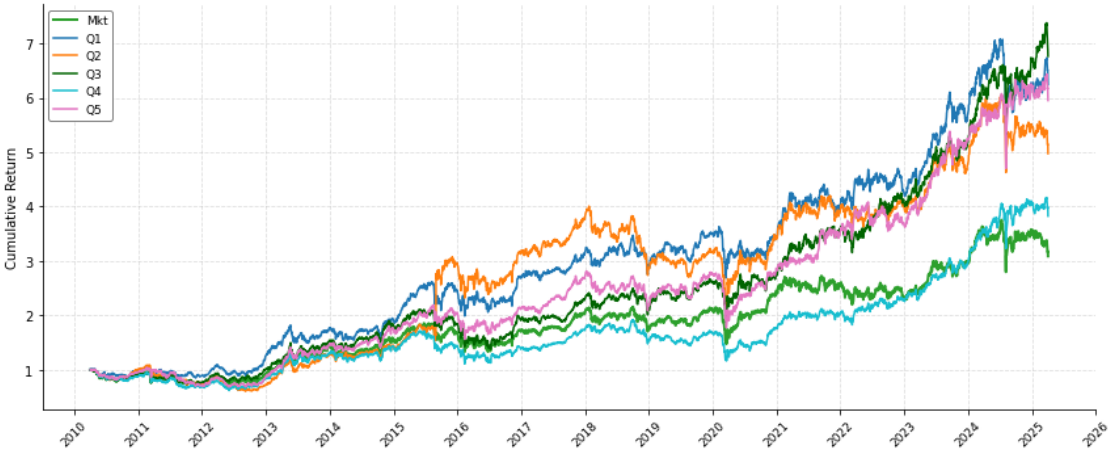

Note: This figure plots the cumulative returns of five quantile portfolios (Q1–Q5) and the market benchmark (Mkt) from 04-01-2010 to 03-04-2025. Each portfolio is equally weighted and rebalanced daily, assuming no transaction costs. Q1 represents the lowest memory assets, while Q5 represents the highest memory assets. The y-axis shows cumulative returns, normalised to 1.0 at the start date, and the x-axis indicates the timeline in daily increments.

TABLE C.7 : Portfolio Performance Metrics – Daily Rebalanced Frictionless Equally Weighted Portfolios (Japan Market).

Portfolio	Q1	Q2	Q3	Q4	Q5	Market
Return	13.41%	11.65%	14.03%	9.66%	13.03%	8.05%
Volatility	19.02%	27.62%	20.64%	20.77%	20.79%	21.06%
Eff Volatility	18.24%	26.15%	17.82%	19.19%	19.42%	18.75%
Sharpe Ratio	0.593	0.708	0.630	0.447	0.379	0.514
Sortino Ratio	0.607	0.726	0.642	0.456	0.374	0.500
Beta	0.783	0.833	0.877	0.879	0.871	1.000
Alpha	7.00%	4.87%	6.92%	2.52%	5.96%	0.00%
Treynor Ratio	0.166	0.135	0.155	0.105	0.145	0.076
ROI	524%	398%	576%	283%	495%	209%
VaR (5%)	-0.0178	-0.0201	-0.0194	-0.0193	-0.0199	-0.0208
ES (5%)	-0.0270	-0.0303	-0.0299	-0.0299	-0.0302	-0.0309
Min Correlation	0.3275	0.3280	0.3349	0.3331	0.3310	–
Max Correlation	0.4217	0.4144	0.4305	0.4480	0.4450	–
Avg Correlation	0.3783	0.3787	0.3871	0.3933	0.3897	–

Note: This table presents performance metrics for daily-rebalanced, frictionless, equally-weighted portfolios formed by quintile sorting on estimated memory parameters. 'Return' and 'Volatility' are annualised; 'Eff Volatility' is the annualised effective standard deviation over a 30-trading-day horizon; 'Sharpe,' 'Sortino,' and 'Treynor' ratios assume a zero risk-free rate; 'Alpha' is the intercept from the memory-augmented CAPM; 'ROI' is cumulative return since inception; 'VaR' and 'ES' denote the one-day 5% Value-at-Risk and Expected Shortfall, respectively; correlation measures report the minimum, maximum, and average pairwise correlations among portfolio constituents. All statistics are computed over 04-01-2010 to 03-04-2025 (3,669 obs.).

FIGURE C.2 : Cumulative Returns of Equally Weighted Portfolios with Transaction Costs and 30-Day Rebalancing Frequency (Japan Market).

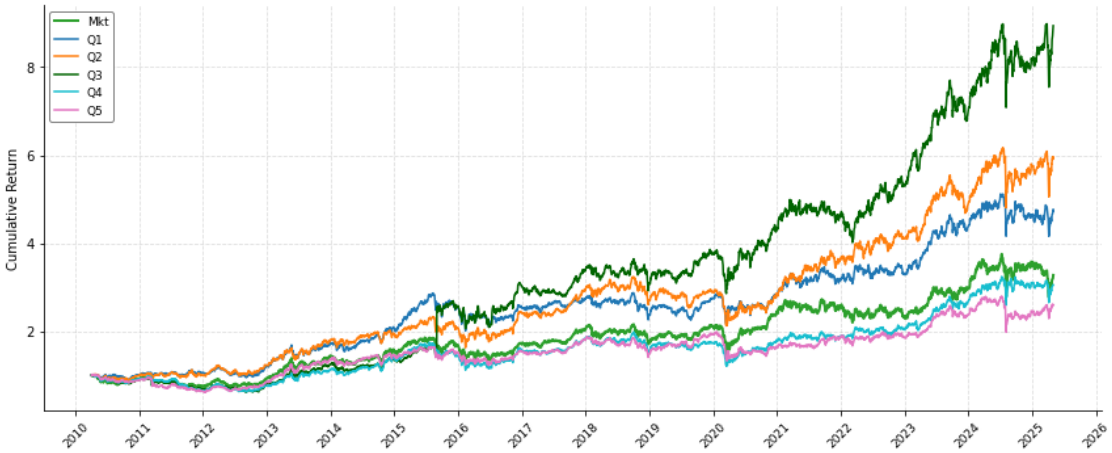

Note: This figure plots the cumulative returns of five quantile portfolios (Q1–Q5) and the market benchmark (Mkt) from 04-01-2010 to 03-04-2025. Each portfolio is equally weighted and rebalanced every 30 days under 0.2% transaction costs per rebalance. Q1 represents the lowest memory assets, while Q5 represents the highest memory assets. The y-axis shows cumulative returns, normalised to 1.0 at the start date, and the x-axis indicates the timeline in daily increments.

TABLE C.8 : Portfolio Performance Metrics – Equally Weighted Portfolios with 30-Day Rebalancing Transaction Costs (Japan Market).

Portfolio	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Market**
Return	11.26%	12.97%	16.14%	7.92%	6.72%	8.44%
Volatility	18.65%	20.12%	27.31%	20.75%	20.48%	21.37%
Eff Volatility	15.81%	17.10%	23.97%	18.47%	18.30%	18.79%
Sharpe Ratio	0.593	0.708	0.630	0.447	0.379	0.514
Sortino Ratio	0.607	0.726	0.642	0.456	0.374	0.500
Beta	0.747	0.821	0.788	0.848	0.829	1.000
Alpha	4.84%	5.96%	9.40%	0.70%	-0.36%	0.00%
Treynor Ratio	0.145	0.153	0.199	0.088	0.076	0.080
ROI	377%	496%	794%	205%	159%	228%
VaR (5%)	-0.0176	-0.0182	-0.0192	-0.0189	-0.0193	-0.0209
ES (5%)	-0.0269	-0.0289	-0.0294	-0.0306	-0.0302	-0.0313
Min Correlation	0.3091	0.3116	0.2835	0.3260	0.2932	–
Max Correlation	0.4646	0.4742	0.4719	0.4662	0.4463	–
Avg Correlation	0.3904	0.3721	0.3768	0.3939	0.3722	–

Note: This table presents performance metrics for equally weighted portfolios rebalanced every 30 trading days with a 0.2% transaction cost per trade. 'Return' and 'Volatility' are annualised; 'Eff Volatility' is the annualised effective standard deviation over the 30-day hold period; 'Sharpe,' 'Sortino,' and 'Treynor' ratios assume a zero risk-free rate; 'Alpha' is the intercept from the memory-augmented CAPM; 'ROI' is cumulative return since inception; 'VaR' and 'ES' denote the one-day 5% Value-at-Risk and Expected Shortfall; correlation measures report the minimum, maximum, and average pairwise correlations among portfolio constituents. All statistics are computed over 04-01-2010 to 03-04-2025 (3,669 obs.).

FIGURE C.3 : Cumulative Returns of Equally Weighted Portfolios with Transaction Costs and 30-Day Rebalancing Frequency (Japan Market) - Excluding 7012.T.

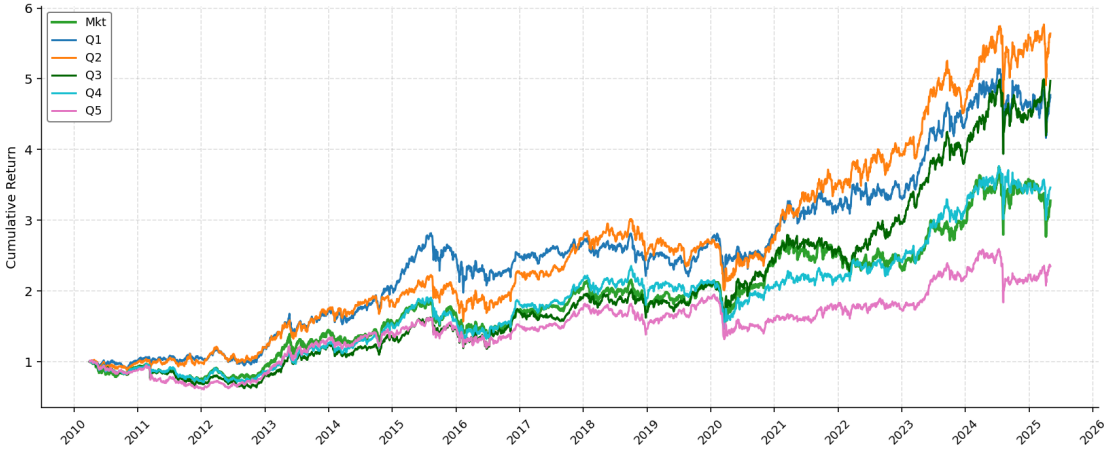

Note: This figure plots the cumulative returns of five quantile portfolios (Q1–Q5) and the market benchmark (Mkt) from 04-01-2010 to 03-04-2025. Portfolio's rebalancing schedules exclude 7012.T asset due to irregularities in its data. Each portfolio is equally weighted and rebalanced every 30 days under 0.2% transaction costs per rebalance. Q1 represents the lowest memory assets, while Q5 represents the highest memory assets. The y-axis shows cumulative returns, normalised to 1.0 at the start date, and the x-axis indicates the timeline in daily increments.

References

Franklin Allen and Douglas Gale. Financial contagion. *Journal of Political Economy*, 108(1):1–33, 2000.

Ron Alquist and Olivier Gervais. The role of financial speculation in driving the price of crude oil. *The Energy Journal*, 34(3):35–54, 2013.

Jose Alvarez-Ramirez, Myriam Cisneros, Carlos Ibarra-Valdez, and Angel Soriano. Multifractal hurst analysis of crude oil prices. *Physica A: Statistical Mechanics and its Applications*, 313(3-4):651–670, 2002.

Alin Andries, D. Ursu, B. Ihnatov, and A. Tiwari. Measuring financial contagion and spillover effects with copulas. *Journal of Financial Stability*, 42:105–124, 2019.

Andrew Ang, Joseph Chen, and Yuhang Xing. Downside risk. *The Review of Financial Studies*, 19(4):1191–1239, 2006.

G. Angeletos and C. Lian. Forward guidance without common knowledge. *American Economic Review*, 2018. .

Brandon Anthonisz, Claire Giziou, and Huy Pham. Asset pricing with downside liquidity risk. *Management Science*, 63(10):3326–3344, 2017.

Nikolaos Antonakakis and David Gabauer. Refined measures of dynamic connectedness based on tvp-var. Technical report, Munich Personal RePEc Archive (MPRA) Paper No. 78282, 2017.

Nikolaos Antonakakis and David Gabauer. Oil market spillovers and time-varying correlation dynamics. *The European Journal of Finance*, 24(13):1015–1031, 2018.

Jushan Bai and Serena Ng. Large dimensional factor analysis. *Foundations and Trends in Econometrics*, 3(2):89–163, 2008.

Richard T Baillie. Long memory processes and fractional integration in econometrics. *Journal of econometrics*, 73(1):5–59, 1996.

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004.

Pedro Barroso and Pedro Santa-Clara. Momentum has its moments. *Journal of Financial Economics*, 116(1):111–120, 2015. .

Jozef Baruník and Tomáš Křehlík. Measuring the frequency dynamics of financial connectedness and systemic risk. *Journal of Financial Econometrics*, 16(2):271–296, 2018.

Christiane Baumeister and Gert Peersman. Time-varying effects of oil supply shocks. *American Economic Journal: Macroeconomics*, 5:1–29, 2013.

Jan Beran. *Statistics for Long-Memory Processes*. Chapman & Hall, London, 1994.

Monica Billio, Mila Getmansky, Andrew W. Lo, and Loriana Pelizzon. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559, 2012.

Pierre Blanc, Jonathan Donier, and Jean-Philippe Bouchaud. Quadratic hawkes processes for financial prices. *Quantitative Finance*, 17(2):171–188, 2017. .

Michael W. Brandt, Pedro Santa-Clara, and Rossen Valkanov. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies*, 22(9):3411–3447, 2009.

Peter J Brockwell and Richard A Davis. *Time series: theory and methods*. Springer science & business media, 1991.

John Y Campbell and Luis M Viceira. *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economic, 2002.

John Y. Campbell, Andrew W. Lo, and A. Craig MacKinlay. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ, 1997.

U Celikyurt and Süleyman Özekici. Multiperiod portfolio optimization models in stochastic markets using the mean–variance approach. *European Journal of Operational Research*, 179(1):186–202, 2007.

Alexei Chekhlov, Stanislav Uryasev, and Michael Zabarankin. Drawdown measure in portfolio optimization. *International Journal of Theoretical and Applied Finance*, 8 (1):13–58, 2005.

L. Chikhi, F. Djaafri, R. Guebli, and A. Tlidji. Does predictive ability of an asset price rest in 'memory'? evidence from arfima-aegas models. *Physica A: Statistical Mechanics and its Applications*, 523:709–728, 2019.

Y. Y. Chong and David F. Hendry. Econometric evaluation of linear macroeconomic models. *Review of Economic Studies*, 53(4):671–690, 1986.

Andrew Clare, James Seaton, Peter Smith, and Stephen Thomas. The sequence-of-returns risk: A new way to think about retiring. *Journal of Pension Economics and Finance*, 19(1):1–21, 2020. .

John H. Cochrane. *Asset Pricing*. Princeton University Press, Princeton, NJ, revised edition edition, 2005. ISBN 978-0691121376.

Lin William Cong, Ke Tang, Yang Wang, and David Zhang. Deep learning of dynamic risk premia from sequential returns, 2021. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3925195. Working paper.

Lin William Cong, Tengyuan Liang, Xiao Zhang, and Wu Zhu. Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information. *Management Science*, 2024. Forthcoming.

Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.

Min Dai, Shuqing Jin, Steven Kou, and Min Yan. A dynamic mean-variance analysis for log returns. *Management Science*, 67(4):2451–2470, 2021.

Stavros Degiannakis and G. Filis. Forecasting oil prices: High-frequency financial data are indeed useful. *Energy Economics*, 2018. .

Danilo Delpini, Giacomo Bormetti, Stefano Marmi, and Fabrizio Lillo. Stochastic volatility with heterogeneous time scales. *The European Physical Journal B*, 85(6):1–9, 2012.

Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The review of Financial studies*, 22(5):1915–1953, 2009.

Shangkun Deng, Takashi Mitsubuchi, Kei Shioda, Tatsuro Shimada, and Akito Sakurai. Combining technical analysis with sentiment analysis for stock price prediction. In *2011 IEEE ninth international conference on dependable, autonomic and secure computing*, pages 800–807. IEEE, 2011.

Francis X. Diebold and Roberto S. Mariano. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263, 1995.

Francis X. Diebold and Kamil Yilmaz. Measuring financial asset return and volatility spillovers, with application to global equity markets. *Journal of Econometrics*, 161 (2):246–266, 2009.

Francis X Diebold and Kamil Yilmaz. Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28 (1):57–66, 2012.

Francis X Diebold and Kamil Yılmaz. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182 (1):119–134, 2014.

Zhuanxin Ding, Clive WJ Granger, and Robert F Engle. A long memory property of stock market returns and a new model. *Journal of empirical finance*, 1(1):83–106, 1993.

Jonathan Donier, Julius Bonart, Iacopo Mastromatteo, and Jean-Philippe Bouchaud. A fully consistent, minimal model for non-linear market impact. *Quantitative Finance*, 15(7):1109–1121, 2015. .

James Doran, David Peterson, and Charles Wright. Investment performance and the sequence of returns. *Financial Services Review*, 21(4):341–356, 2012.

Graham Elliott, Thomas J Rothenberg, and James H Stock. Efficient tests for an autoregressive unit root. *Econometrica*, 64(4):813–836, 1996.

Robert Engle. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3):339–350, 2002.

Larry G Epstein and Stanley E Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 1989.

Eugene F Fama and Kenneth R French. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, 60(1): 55–73, 1987.

Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1):3–56, 1993.

Eugene F. Fama and Kenneth R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, 2015.

J. Doyne Farmer, Austin Gerig, Fabrizio Lillo, and Szabolcs Mike. Market efficiency and the long-memory of supply and demand: Is price impact variable and permanent or fixed and temporary? *Quantitative Finance*, 6(2):107–112, 2006. .

E Robert Fernholz and E Robert Fernholz. *Stochastic portfolio theory*. Springer, 2002.

T. Fischer and C. Krauss. Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2):654–669, 2018.

Mario Forni and Luca Gambetti. Sufficient information in structural vars. *Journal of Monetary Economics*, 66:124–136, 2014.

Raffaella Giacomini and Barbara Rossi. Forecasting in economics and finance. *Annual Review of Economics*, 9:1–27, 2017.

Raffaella Giacomini and Halbert White. Tests of conditional predictive ability. *Econometrica*, 74(6):1545–1578, 2006.

Michael R Gibbons, Stephen A Ross, and Jay Shanken. A test of the efficiency of a given portfolio. *Econometrica: Journal of the Econometric Society*, pages 1121–1152, 1989.

C. Gomes and H. Waelbroeck. Is market impact a measure of the information value of trades? market response to liquidity vs. informed metaorders. *Quantitative Finance*, 15(5):773–793, 2015. .

M. Gonzalez, L. Lima, and R. Marinho. Business cycle durations: An empirical investigation using bayesian structural time series models. *Economic Modelling*, 47:105–116, 2015.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969.

Clive WJ Granger and Roselyne Joyeux. An introduction to long-memory time series models and fractional differencing. *Journal of time series analysis*, 1(1):15–29, 1980.

Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

Sanford J Grossman and Zhongquan Zhou. Optimal investment strategies for controlling drawdowns. *Mathematical Finance*, 3(3):241–276, 1993.

Jian Guo, Xia Wang, and Ming Zhou. Forecasting brent and wti oil prices with hybrid machine learning models. *Energy Economics*, 122:106707, 2023a.

Lili Guo, Xinya Huang, Yanjiao Li, and Houjian Li. Forecasting crude oil futures price using machine learning methods: Evidence from china. *Energy Economics*, 127:107089, 2023b.

Constantin Gurdgiev and D. O'Loughlin. Herding and anchoring in cryptocurrency markets: Investor reaction to fear and uncertainty. *Microeconomics: General Equilibrium Disequilibrium Models of Financial Markets eJournal*, 2020. .

Nicolae Gârleanu and Lasse Heje Pedersen. Dynamic trading with predictable returns and transaction costs. *The Journal of Finance*, 68(6):2309–2340, 2013.

James D Hamilton. *Time series analysis*. Princeton university press, 1994.

James D. Hamilton. Causes and consequences of the oil shock of 2007–08. *Brookings Papers on Economic Activity*, 40(1):215–261, 2009.

Peter Reinhard Hansen, Asger Lunde, and James M. Nason. The model confidence set. *Econometrica*, 79(2):453–497, 2011.

S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9 (8):1735–1780, 1997a.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997b.

Harrison G. Hong and Jeremy C. Stein. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance*, 54(6):2143–2184, 1999.

J. R. M. Hoskin. Fractional differencing. *Biometrika*, 68(1):165–176, 04 1981. ISSN 0006-3444.

Atsushi Inoue and Lutz Kilian. On the selection of forecasting models. *Journal of Econometrics*, 130(2):273–306, 2006.

Johan Ludwig William Valdemar Jensen. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta mathematica*, 30(1):175–193, 1906.

Daniel J. Jobson and Burgess M. Korkie. Putting markowitz theory to work. *Journal of Portfolio Management*, 7:70–74, 1981.

Søren Johansen. Statistical analysis of cointegration vectors. *Journal of economic dynamics and control*, 12(2-3):231–254, 1988.

Søren Johansen. *Likelihood-based inference in cointegrated vector autoregressive models*. Oxford University Press, 1995a.

Søren Johansen. A representation theory for a class of vector autoregressive models for fractional processes. *Econometric Theory*, 24(3):651–676, 2008.

Søren Johansen and Morten Ø. Nielsen. Likelihood inference for a nonstationary fractional autoregressive model. *Journal of Econometrics*, 167(2):382–399, 2012a.

Søren Johansen and Morten Ørregaard Nielsen. Likelihood inference for a fractionally cointegrated vector autoregressive model. *Econometrica*, 80(6):2667–2732, 2012b.

Soren Johansen and Morten Ørregaard Nielsen. The role of initial values in nonstationary fractional time series models. *Univ. of Copenhagen Dept. of Economics Discussion Paper*, (12-18), 2012c.

Søren Johansen. *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford University Press, 1995b.

Søren Johansen and Morten Ørregaard Nielsen. The role of initial values in conditional sum-of-squares estimation of nonstationary fractional time series models. *Econometric Theory*, 32(5):1095, 2016.

Maggie EC Jones, Morten Ørregaard Nielsen, and Michał Ksawery Popiel. A fractionally cointegrated var analysis of economic voting and political support. *Canadian Journal of Economics/Revue canadienne d'économique*, 47(4):1078–1130, 2014.

Philippe Jorion. Bayes–stein estimation for portfolio analysis. *Journal of Financial and Quantitative Analysis*, 21(3):277–292, 1986. .

John L Kelly. A new interpretation of information rate. *the bell system technical journal*, 35(4):917–926, 1956.

Lutz Kilian. Not all oil price shocks are alike. *American Economic Review*, 99(3): 1053–1069, 2009.

Lutz Kilian and Daniel P. Murphy. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478, 2014.

T. King and Dimitrios Koutmos. Herding and feedback trading in cryptocurrency markets. *Annals of Operations Research*, 300:79 – 96, 2021. .

E. Kohlscheen. A cross-country analysis of friedman's plucking theory: Business cycle durations and global recessions. *BIS Working Paper*, (1076), 2023.

Denis Kwiatkowski, Peter CB Phillips, Peter Schmidt, and Yongcheol Shin. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of econometrics*, 54(1-3):159–178, 1992.

Carlos León and Eduardo Revéiz. Portfolio optimization and long-term dependence. *Cuadernos de Economía*, 29(53):35–63, 2010.

Andrew W. Lo. Long-term memory in stock market prices. *Econometrica*, 59(5):1279–1313, 1991.

Andrew W Lo. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management, Forthcoming*, 2004.

David G Luenberger. Products of trees for investment analysis. *Journal of Economic Dynamics and Control*, 22(8-9):1403–1417, 1998.

A. G. Malliaris and M. E. Malliaris. Forecasting energy product prices: Linear versus neural network models. *Energy Economics*, 30(3):939–956, 2008.

Massimiliano Marcellino, James H. Stock, and Mark W. Watson. A comparison of direct and iterated multistep ar methods for forecasting macroeconomic time series. *Journal of Econometrics*, 135(1-2):499–526, 2006. .

Harry Markowitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952.

Robert C. Merton. An intertemporal capital asset pricing model. *Econometrica*, 41(5): 867–887, 1973.

Robert C Merton. Optimum consumption and portfolio rules in a continuous-time model. In *Stochastic optimization models in finance*, pages 621–661. Elsevier, 1975.

Tom Messmore. Variance drain. *Journal of Portfolio Management*, 21(4):104, 1995.

Claudio Morana. Oil price dynamics, macro-finance interactions and the role of financial speculation. *Journal of banking & finance*, 37(1):206–226, 2013.

Alan Moreira and Tyler Muir. Volatility-managed portfolios. *The Journal of Finance*, 72(4):1611–1644, 2017. .

Dmitri Mustanen, Ahmad Maaitah, Tapas Mishra, and Mamata Parhi. The power of investors' optimism and pessimism in oil market forecasting. *Energy Economics*, page 106273, 2022.

Serena Ng and Pierre Perron. Lag length selection and the construction of unit root tests with good size and power. *Econometrica*, 69(6):1519–1554, 2001.

Irena Pavlova, Maria E. de Boyrie, and Alireza M. Parhizgari. Long memory in reit volatility and changes in the unconditional mean. *Journal of Real Estate Finance and Economics*, 49:449–472, 2014.

T. M. Pedersen. The length and the phases of the business cycle: A study of 19 oecd countries, 1948–1994. *Economics Letters*, 58(1):63–69, 1998.

Pierre Perron. The great crash, the oil price shock, and the unit root hypothesis. *Econometrica*, 57(6):1361–1401, 1989.

Wade D Pfau. Safe savings rates: A new approach to retirement planning over the life cycle. *Journal of Financial Planning*, 24(5):42–50, 2011.

Dimitris N. Politis and Joseph P. Romano. The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313, 1994.

John W Pratt. Risk aversion in the small and in the large. In *Uncertainty in economics*, pages 59–79. Elsevier, 1978.

Adrian E. Raftery. Bayesian model averaging in social research. *Sociological Methods & Research*, 25(2):221–239, 1997.

Juan C. Reboredo and Alberto Ugolini. Downside and upside spillovers between oil and financial markets. *Energy Economics*, 54:42–57, 2016.

Peter M Robinson. Gaussian semiparametric estimation of long range dependence. *The Annals of statistics*, pages 1630–1661, 1995.

R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

Carlos Rodríguez. Validez del supuesto de neutralidad del horizonte de tiempo en el capm. *Revista de Métodos Cuantitativos para la Economía y la Empresa*, 12:81–98, 2011.

Afees A. Salisu and Xuan Vinh Vo. Predicting energy market volatility using asymmetric and long memory garch models with exogenous variables. *Energy Economics*, 86: 104628, 2020a.

Afees A Salisu and Xuan Vinh Vo. Predicting stock returns in the presence of covid-19 pandemic: The role of health news. *International Review of Financial Analysis*, 71: 101546, 2020b.

Longyu Shi, Yunyun Wang, Wenyue Li, and Zhimin Zhang. Multi-period mean–variance portfolio optimization with capital injections. *Mathematics and Computers in Simulation*, 233:400–412, 2025.

Robert J Shiller. From efficient markets theory to behavioral finance. *Journal of economic perspectives*, 17(1):83–104, 2003.

Katsumi Shimotsu. Exact local Whittle estimation of fractional integration with unknown mean and time trend. *Econometric Theory*, 26(2):501–540, 2010.

Katsumi Shimotsu and Peter CB Phillips. Local Whittle estimation of fractional integration and some of its variants. *Journal of Econometrics*, 130(2):209–233, 2006.

Herbert A. Simon. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482, 1962.

Christopher A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980.

Kenneth J. Singleton. Investor flows and the 2008 boom/bust in oil prices. *Management Science*, 60(2):300–318, 2014.

Fallaw Sowell. Maximum likelihood estimation of stationary univariate fractionally integrated time series models. *Journal of econometrics*, 53(1-3):165–188, 1992.

J. H. Stock and M. W. Watson. Business cycle fluctuations in u.s. macroeconomic time series. *Handbook of Macroeconomics*, 1A:3–64, 1999.

James H. Stock and Mark W. Watson. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97(460): 1167–1179, 2002.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

Douglas Eduardo Turatti, Fernando Henrique PS Mendes, and João Frois Caldeira. Testing for mean reversion in bitcoin returns with gibbs-sampling-augmented randomization. *Finance Research Letters*, 34:101252, 2020.

Abraham Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of the American Mathematical society*, 54(3):426–482, 1943.

Dayong Zhang et al. Energy-financial market connectedness. *Energy Economics*, 2019.

Guofu Zhou, Haimeng Zhang, and Yuzhao Zhu. Long-run risks with multiple macroeconomic volatilities. *Management Science*, 61(12):2962–2979, 2015.

Xiaowen Zhou and Wei Duan. Dynamic connectedness among commodity and financial markets: A frequency-domain perspective. *The European Journal of Finance*, 27(12): 1113–1135, 2021.

Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2): 301–320, 2005.