EIGENFUND TECHNICAL REPORT EF-2026-01

Version 1.0 — January 2026

Finite-Time Lyapunov Exponent Algorithm Applied to Equity Market Returns: Methodology and Exploratory Analysis

[D. Aliev, S. Nuzhdov]

1. Executive Summary

This report documents the construction, validation, and exploratory application of a Finite-Time Lyapunov Exponent (FTLE) computation algorithm for time series analysis. The algorithm was validated on the Lorenz system under two parameter configurations using both Euler and fourth-order Runge-Kutta integration schemes, recovering values consistent with the known largest Lyapunov exponents ($\lambda_1 \approx 0.9$ for standard parameters; $\lambda_1 \approx 1.5$ for alternative parameters). The method was applied to daily log-returns of three equity ETFs—SPY, QQQ, and IWM—over the period 2000–2024 at two temporal scales: 250 trading days (approximately one trading year) and 500 trading days (approximately two trading years). FTLE time series were computed and characterised, revealing consistently positive median values with occasional negative excursions. Surrogate data testing assessed whether the observed FTLE structure is distinguishable from linear stochastic processes; no configuration rejected the null hypothesis at $p = 0.05$, though marginal rejection at $p = 0.10$ for select configurations suggests that nonlinear structure cannot be ruled out. This report presents methodology and exploratory observations; no predictive claims are made.

2. Introduction

2.1 Beyond Stochastic Models

Standard quantitative models treat financial returns as realisations of stochastic processes [1]. The random walk hypothesis, geometric Brownian motion, and GARCH-family models share a common assumption: price dynamics are fundamentally random, governed by probability distributions that may be stationary or slowly varying [2]. These frameworks have proven useful for option pricing, risk measurement, and portfolio construction [3]. However, several robust empirical features of financial markets are not adequately captured by such models: fat-tailed return distributions exceeding Gaussian predictions, volatility clustering where periods of high and low volatility persist, long-range dependence in absolute returns [4], and abrupt regime changes that standard models treat as exogenous shocks [1]. An alternative perspective, drawn from dynamical systems theory, treats price dynamics as the output of a deterministic nonlinear system subject to observational or dynamical noise [1]. Under this framing, apparent randomness

may arise not from intrinsic stochasticity but from sensitive dependence on initial conditions—the hallmark of deterministic chaos [5]. This motivates the application of tools from nonlinear dynamics to financial time series, not to assert that markets are chaotic, but to characterise their dynamical structure.

2.2 Chaos-Theoretic Approaches in Finance

The application of chaos theory to financial markets gained momentum in the late 1980s and early 1990s. Peters (1991) [6] proposed that markets exhibit fractal structure and long-memory properties inconsistent with efficient market assumptions. Brock, Hsieh, and LeBaron (1991) [7] developed statistical tests for nonlinear dependence and applied them to financial data. These early efforts generated substantial interest but also significant critique. Hsieh (1991) [1] demonstrated that while nonlinear structure exists in financial returns, distinguishing low-dimensional chaos from high-dimensional stochastic processes requires data volumes and precision rarely available in practice. The consensus that emerged was cautious: nonlinear methods can characterise market dynamics, but claims of deterministic chaos in financial markets are unsubstantiated [8]. Contemporary applications focus on regime detection, early warning indicators, and dynamical characterisation rather than deterministic prediction [9]. The Finite-Time Lyapunov Exponent, originally developed for fluid dynamics and Lagrangian coherent structure detection [10], offers a local, time-dependent measure of dynamical instability suitable for non-stationary systems. In this report, FTLE is applied to equity market returns as a diagnostic tool; methodology and observations are documented without claims of predictive power.

2.3 Scope

The documented project has two objectives. The primary is methodological: to implement a rigorous FTLE computation pipeline, validate it against a system with known Lyapunov exponents, and document the computational procedure in sufficient detail for reproduction. The secondary objective is exploratory: to apply the validated method to equity market data and document empirical observations. The given report presents the outcomes of both objectives. Three US equity ETFs (SPY, QQQ, IWM) are examined over a 24-year period. No claims are made that equity markets are chaotic, that FTLE predicts market movements, detects market regime, or that the observations reported here constitute trading signals. The work is presented as a methodological foundation for potential future investigation.

3. Theoretical Framework

3.1 Dynamical Systems

A dynamical system is defined by a state space and an evolution law that determines how the state changes with time [5]. Consider a state vector $x \in \mathbb{R}^n$ representing the complete instantaneous configuration of a system. For continuous-time autonomous systems, the evolution is governed by a system of ordinary differential equations:

$$dx/dt = F(x) \quad (1)$$

where F: $\mathbb{R}^n \to \mathbb{R}^n$ is a vector field [11]. For discrete-time systems, the evolution follows an iterated map:

$$x_{n+1} = F(x_n) \quad (2)$$

The space $\mathbb{R}^n$ in which the state vector evolves is termed the phase space or state space; its dimension n equals the number of independent variables required to specify the state [5]. A trajectory (or orbit) is the path x(t) traced through phase space as the system evolves from an initial condition x_0. The flow map φ_t encapsulates this evolution:

$$\varphi_t: x_0 \mapsto x(t) \quad (3)$$

mapping any initial state to its position after time t [10]. An attractor is an invariant set in phase space toward which trajectories from nearby initial conditions converge asymptotically. Attractors may be fixed point, periodic orbit for continuous time, or strange attractors which have a complicated aspect typical of fractal objects [11].

3.2 State Space Reconstruction

In experimental settings, the full state vector x is typically inaccessible. Often, only a scalar time series s(t) = h(x(t)) is observed, where h is some measurement function [5]. The method of delay coordinate embedding reconstructs a proxy for the original phase space from this scalar observable. Given a time series $\{s_t\}$, delay vectors are constructed:

$$y(t) = [s(t), s(t-\tau), s(t-2\tau), \ldots, s(t-(m-1)\tau)] \quad (4)$$

where τ is the embedding delay and m is the embedding dimension. The Takens embedding theorem [12] provides the theoretical foundation: for a generic smooth observable h and sufficiently large m (specifically, m > 2d where d is the box-counting dimension of the attractor), the delay embedding is a diffeomorphism from the original attractor to its reconstruction. This guarantees that topological and metric properties of the dynamics are preserved. Quantities such as Lyapunov exponents, which are invariants under smooth coordinate transformations, can therefore be computed from the reconstructed space [11].

3.3 Lyapunov Exponents

Lyapunov exponents quantify the average exponential rate of divergence or convergence of infinitesimally close trajectories in phase space. Consider two initial conditions separated by a small perturbation δx(0). As the system evolves, this perturbation grows or shrinks approximately exponentially:

$$|\delta x(t)| \sim |\delta x(0)| \, exp(\lambda t) \quad (5)$$

The Lyapunov exponent λ is defined as the asymptotic rate of this exponential separation [5]:

$$\lambda = \lim_{t\to\infty} \left(\frac{1}{t}\right) ln(|\delta x(t)| \,/\, |\delta x(0)|) \quad (6)$$

An n-dimensional system possesses n Lyapunov exponents, conventionally ordered $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$, corresponding to different directions in the tangent space. The largest exponent λ_1 dominates the long-term behaviour of generic perturbations. A positive largest Lyapunov exponent ($\lambda_1 > 0$)

indicates sensitive dependence on initial conditions: nearby trajectories diverge exponentially, rendering long-term prediction impossible even with arbitrarily precise knowledge of the initial state. A bounded system with at least one positive Lyapunov exponent is termed chaotic. The classical Lyapunov exponent requires an asymptotic limit, which presupposes stationarity—the dynamical law must remain constant over the measurement period [5]. This assumption is violated in financial markets, where regime changes, structural breaks, and evolving market microstructure preclude long-term stationarity [1].

3.4 Finite-Time Lyapunov Exponents

The Finite-Time Lyapunov Exponent (FTLE) addresses non-stationarity by replacing the asymptotic limit with a finite integration window [13]. Rather than measuring trajectory separation as $t \to \infty$, FTLE quantifies separation over a fixed time horizon T:

$$\sigma^T(x_0) = \left(\frac{1}{T}\right) ln(|\delta x(T)| / |\delta x(0)|) \quad (7)$$

The FTLE σ^T is a local quantity: it depends on the starting point x_0 and the time at which measurement begins [10]. This locality renders it suitable for non-stationary systems where the dynamical character may vary over time.

A rigorous formulation employs the deformation gradient tensor. The flow map φ^T takes initial conditions to their evolved positions. Its gradient (Jacobian) $F = \nabla\varphi^T$ describes how infinitesimal displacements transform under the flow. The right Cauchy-Green deformation tensor is defined as:

$$C = F^T F \quad (8)$$

The eigenvalues of C characterise the stretching of material elements [13]. The FTLE is defined via the largest eigenvalue $\lambda_{max}(C)$:

$$\sigma^T = \left(\frac{1}{2T}\right) ln\left(\lambda_{max}(C)\right) \quad (9)$$

This formulation, developed in the context of fluid dynamics and Lagrangian coherent structures, provides a geometrically meaningful measure of local instability. High FTLE values indicate regions where trajectories diverge rapidly—dynamically unstable zones. Low FTLE values indicate coherent regions where nearby trajectories remain close [10].

3.5 FTLE and Volatility: Distinct Quantities

In financial markets, FTLE and statistical volatility are distinct quantities and should not be conflated. Volatility, typically defined as the standard deviation of returns, measures the magnitude of price fluctuations—the dispersion of outcomes [14]. FTLE measures the rate of trajectory divergence in reconstructed phase space—the sensitivity of a system's evolution to its current state [5]. These quantities are not equivalent. A market may exhibit low volatility (small daily returns) while possessing high FTLE (the dynamical state is unstable; small perturbations grow rapidly). Conversely, high volatility (large returns) may coincide with low FTLE (trajectories in phase space converge; the system approaches a coherent regime). FTLE is a geometric property of the reconstructed dynamics; volatility is a statistical property of observed

outcomes. Any empirical relationship between them requires investigation rather than assumption.

4. Methodology

4.1 FTLE Computation Algorithm

The FTLE computation algorithm proceeds in six steps. The method adapts trajectory-based approaches for estimating the largest Lyapunov exponent from time series data. Extending them to finite-time computation via QR-stabilised Jacobian chaining [15] and interpreting the result in the Cauchy-Green deformation framework [10].

Step 1 — Delay Embedding. From a scalar time series $\{s_t\}$, m-dimensional delay vectors are constructed for each time point within the analysis window [12]:

$$y(t) = [s(t), s(t-\tau), s(t-2\tau), \ldots, s(t-(m-1)\tau)] \quad (10)$$

where τ is the embedding delay and m is the embedding dimension.

Step 2 — Neighbour Identification. For each reference point y(t), k nearest neighbours in the embedding space are identified using Euclidean distance. Temporally adjacent points within the Theiler window are excluded to prevent spurious correlations arising from serial dependence [16].

Step 3 — Local Jacobian Estimation. For each time step t, the local Jacobian A_t of the discrete map is estimated via least-squares regression. Initial displacements $\Delta y_j(0) = y(t_j) - y(t)$ and one-step evolved displacements $\Delta y_j(1) = y(t_j + 1) - y(t + 1)$ are computed for each neighbour j. The Jacobian A_t satisfies:

$$\Delta y(1) \approx A_t \cdot \Delta y(0) \quad (11)$$

in the least-squares sense [17].

Step 4 — QR-Stabilised Integration. The flow map Jacobian over integration time T is computed by chaining local Jacobians:

$$D\Phi_t^T = A_{t+T-1} \cdot A_{t+T-2} \cdot \ldots \cdot A_t \quad (12)$$

To prevent numerical overflow from exponentially growing eigenvalues, QR decomposition is applied at each multiplication step. After multiplying by each successive Jacobian, the result is factored as Q·R where Q is orthogonal and R is upper triangular. The orthogonal factor Q is retained for the next multiplication, while the diagonal elements of R—which capture the stretching factors—are stored as logarithms and summed across all T steps. This accumulates the total logarithmic growth along each direction without overflow [15].

Step 5 — FTLE Computation. The sum $\Sigma_k \log|R_{k,ii}|$ gives total logarithmic stretching along direction i over T steps. The FTLE is the maximum rate across all directions - the largest finite-time Lyapunov exponent is extracted as:

$$\sigma^T = \left(\frac{1}{T}\right) max_i\left(\Sigma_k \, log|R_{k,ii}|\right) \quad (13)$$

This is mathematically equivalent to:

$$\sigma^T = \left(\frac{1}{2T}\right) ln\left(\lambda_{max}(C)\right) \quad (14)$$

where $C = (D\Phi_t^T)^T D\Phi_t^T$ is the right Cauchy-Green deformation tensor [10].

Step 6 — Sliding Window. Steps 1–5 are repeated for successive time windows, advancing by one time step, to produce an FTLE time series, tracking how local stability evolves.

Parameters

The algorithm requires specification of six parameters. The embedding delay τ determines the temporal separation between coordinates in the delay vector; values that are too small yield redundant information, while values that are too large lose dynamical connection between coordinates. The embedding dimension m specifies the number of coordinates in the delay vector and must be sufficient to unfold the attractor geometry without self-intersections. The Theiler window excludes temporally adjacent points from the neighbour search, preventing autocorrelation from contaminating divergence estimates. The integration time T defines the horizon over which trajectory separation is measured, balancing locality against sufficient time for exponential divergence to manifest. The neighbour count k determines the number of neighbours used for Jacobian estimation; small k is noise-sensitive while large k blurs local dynamics. The window size W specifies the data segment for each FTLE computation, converting the global measure into a time-varying indicator suitable for non-stationary systems.

4.2 Validation: Lorenz System

The algorithm is validated on the Lorenz system, a three-dimensional autonomous dissipative flow that exhibits chaotic behaviour for certain parameter values [18]. The system is defined by:

$$dx/dt = \sigma(y - x), \; dy/dt = \rho x - y - xz, \; dz/dt = xy - \beta z \quad (15)$$

Two parameter configurations are employed to test algorithm robustness across different attractor geometries. Configuration A uses standard parameters ($\sigma = 10$, $\rho = 28$, $\beta = 8/3$) for which the largest Lyapunov exponent is known to be $\lambda_1 \approx 0.9$ nats per time unit [19]. Configuration B uses alternative parameters ($\sigma = 16$, $\rho = 45.92$, $\beta = 4$) yielding a known largest exponent of $\lambda_1 \approx 1.50$ nats per time unit [20].

To assess the sensitivity of the FTLE algorithm to numerical integration accuracy, both configurations are integrated using two methods: forward Euler and fourth-order Runge-Kutta (RK4). Both integrators use a time step of $dt = 0.01$, generating 10,000 data points after discarding initial transients. If the FTLE estimates are consistent across integration methods, this provides evidence that the Lyapunov computation is robust to the underlying numerical scheme—a necessary condition for confidence in the algorithm's correctness.

For the Lorenz validation, substantially larger algorithm parameters are employed than for market data: $W = 9{,}000$ and $T = 5{,}000$. The rationale is that the Lorenz system is stationary—its attractor geometry remains constant indefinitely—requiring extensive temporal sampling to accurately estimate the true asymptotic Lyapunov exponent [17]. A stricter threshold for the False Nearest Neighbours method (see Section 4.4 for method description) of 1% is also applied,

exploiting the low noise level in synthetic data. In contrast, financial markets are non-stationary, requiring shorter windows to track evolving dynamics rather than averaging over regime changes.

4.3 Market Data

The method is applied to three US equity exchange-traded funds representing different market segments: SPY (SPDR S&P 500 ETF Trust, tracking large-capitalisation US equities), QQQ (Invesco QQQ Trust, tracking the Nasdaq-100 index of technology-weighted equities), and IWM (iShares Russell 2000 ETF, tracking small-capitalisation US equities). Daily adjusted closing prices are obtained from open-source financial data providers for SPY and QQQ from January 2000 through December 2024, and for IWM from May 2000 (inception) through December 2024.

Prices are converted to log-returns:

$$r_t = ln(P_t / P_{t-1}) \tag{16}$$

Log-returns are standard in financial time series analysis due to their additive properties over time and approximate normality for short intervals [4]. The FTLE computation is performed on the log-return series rather than price levels, as returns more directly reflect the underlying dynamics of market movements.

4.4 Parameter Selection

Parameter selection for market data follows established procedures from the nonlinear time series analysis literature. The embedding delay τ is selected as the first minimum of the Average Mutual Information (AMI) function, which identifies the lag at which successive coordinates provide maximally independent information [21]. Unlike linear autocorrelation, AMI captures nonlinear dependencies, making it appropriate for potentially nonlinear dynamics.

The embedding dimension m is determined by the False Nearest Neighbours (FNN) method [22]. The principle is intuitive: if the embedding dimension is too low, the attractor is projected onto a space too small to contain it without self-intersection, causing geometrically distant points to appear close. For each point, the algorithm identifies its nearest neighbour, then checks whether this neighbour remains close when the embedding dimension is increased by one. If the neighbour moves away substantially, it was a 'false' neighbour arising from projection rather than true dynamical proximity. As the dimension increases, the attractor unfolds and false neighbours separate. The embedding dimension is selected as the first value at which the proportion of false neighbours falls below a threshold (5% for market data, 1% for the cleaner Lorenz system).

The Theiler window is set to 2τ, following the recommendation of Theiler (1990) [23] to exclude temporally correlated points from neighbour searches.

The integration time T is selected to minimise FTLE variance across different values of k, providing a stability criterion that identifies the regime where the exponent estimate is least sensitive to the number of neighbours. The neighbour count k is determined by plateau search: the first value of k at which the relative change in FTLE falls below 20% is selected, indicating convergence of the estimate. Two window sizes W are tested: 250 trading days (approximately one calendar year of market activity) and 500 trading days (approximately two calendar years of

market activity). These scales are chosen to span typical market cycle durations while retaining sufficient temporal resolution to track evolving dynamics. The 250-day window provides higher temporal resolution but may be more susceptible to noise; the 500-day window provides greater averaging but may smooth over shorter-term dynamical changes.

4.5 Statistical Framework

To assess whether the observed FTLE structure contains information beyond that present in a linear stochastic process, surrogate data testing is employed [24]. The null hypothesis is that the original time series was generated by a Gaussian linear stochastic process—specifically, Gaussian white noise passed through a linear filter (an autoregressive moving-average process, ARMA). Under this null hypothesis, all temporal structure in the data arises from linear correlations; any apparent nonlinear features are statistical fluctuations. If the null hypothesis cannot be rejected, the observed FTLE values are consistent with having arisen from a linear process, providing no evidence for nonlinear dynamics.

The Amplitude-Adjusted Fourier Transform (AAFT) method generates surrogate time series that preserve both the power spectrum (linear correlation structure) and the marginal amplitude distribution of the original data, while destroying any nonlinear temporal dependencies [25]. The procedure randomises the phases of the Fourier transform while preserving amplitudes, then rescales to match the original distribution. The result is a constrained randomisation: surrogates are statistically indistinguishable from the original under any linear test, but differ if nonlinear structure is present.

For each market data configuration, N = 100 AAFT surrogates are generated. The FTLE algorithm is applied to each surrogate, producing a distribution of FTLE values under the null hypothesis. The test statistic is computed as the rank of the original FTLE value within the surrogate distribution. If the original value is more extreme than 95 of 100 surrogates (i.e., lies in the top or bottom 5%), the null hypothesis is rejected at $p = 0.05$; if more extreme than 90 of 100, rejection at $p = 0.10$ is reported. Rejection indicates that the FTLE structure of the original data is unlikely to have arisen from a linear stochastic process.

5. Results

5.1 Lorenz Validation

The algorithm recovers Lyapunov exponents consistent with known values for both Lorenz configurations across both integration methods. For Configuration A (standard parameters, known $\lambda_1 \approx 0.9$), the Euler-integrated data yields an estimated exponent of $\lambda = 1.01$, while the RK4-integrated data yields $\lambda = 0.97$. For Configuration B (alternative parameters, known $\lambda_1 \approx 1.50$), the Euler-integrated data yields $\lambda = 1.54$, while the RK4-integrated data yields $\lambda = 1.33$.

All four estimates correctly identify the systems as chaotic (positive exponent) and recover the approximate magnitude of the known values. The agreement between Euler and RK4 results for each configuration suggests that the FTLE computation is not unduly sensitive to the numerical integration scheme, though some variation is present. This validation supports correct algorithm implementation and provides a basis for application to empirical data where the true exponent is unknown.

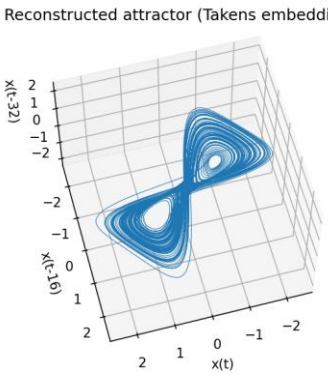

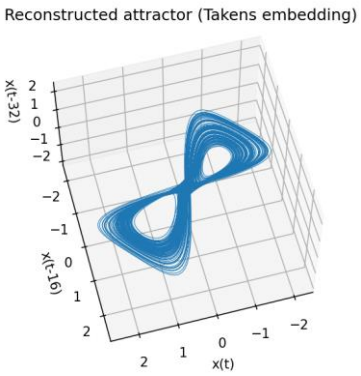

[Figures 1,2: Lorenz attractor reconstructions and FTLE convergence for Configuration A, Euler and RK4 integration]

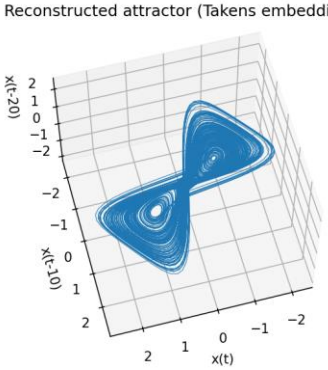

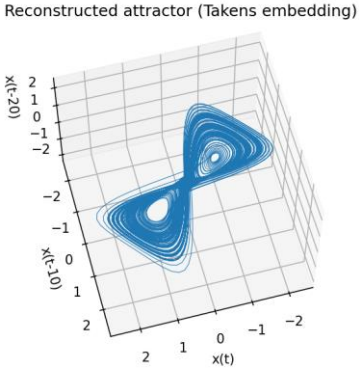

[Figures 3,4: Lorenz attractor reconstructions and FTLE convergence for Configuration B, Euler and RK4 integration]

5.2 Embedding Parameters

The Average Mutual Information and False Nearest Neighbours analyses determine embedding parameters for each asset. The AMI function exhibits a clear first minimum for all three ETFs, providing unambiguous delay selection. The FNN percentage decreases monotonically with embedding dimensions, falling below the 5% threshold at dimensions consistent across assets. Figures 1-6 depict the analyses.

Data set	tau	m
SPY	3	6
QQQ	4	6
IWM	5	6

[Table 1: Embedding parameters (τ, m) for SPY, QQQ, IWM]

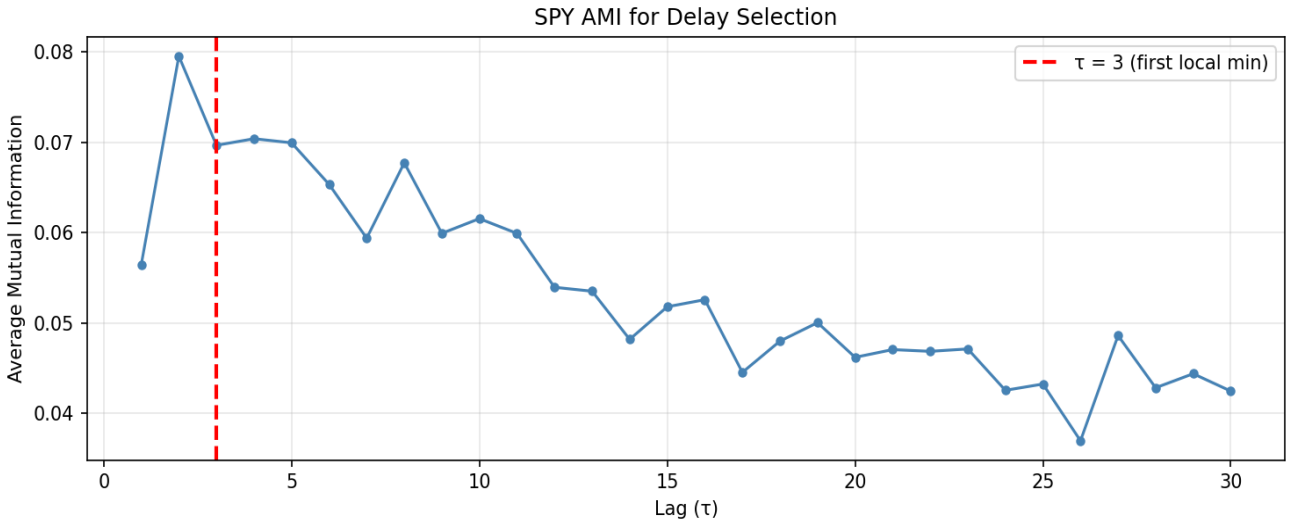

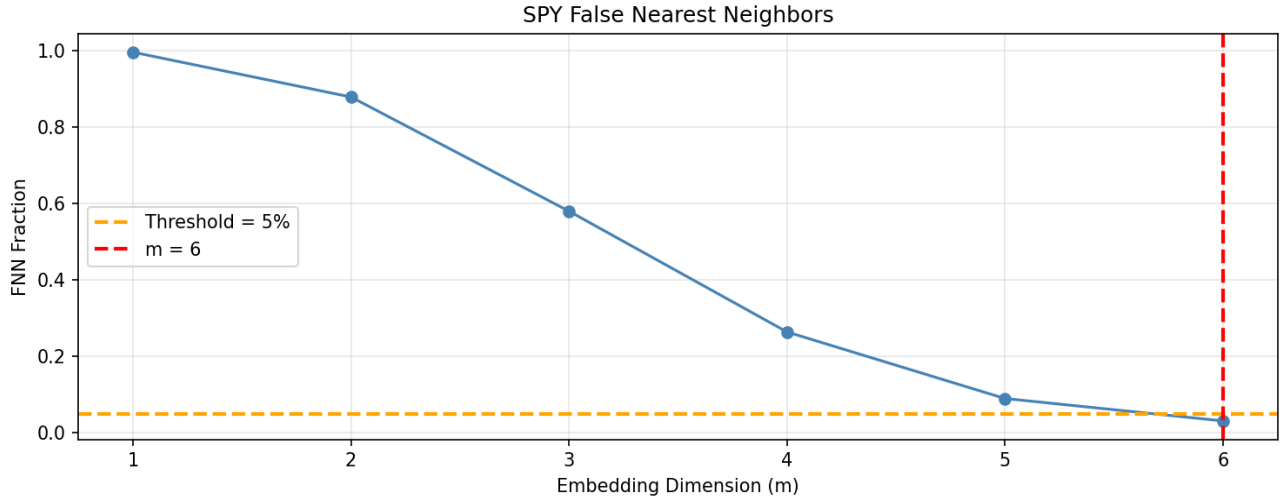

[Figures 5,6: AMI and FNN plots for SPY]

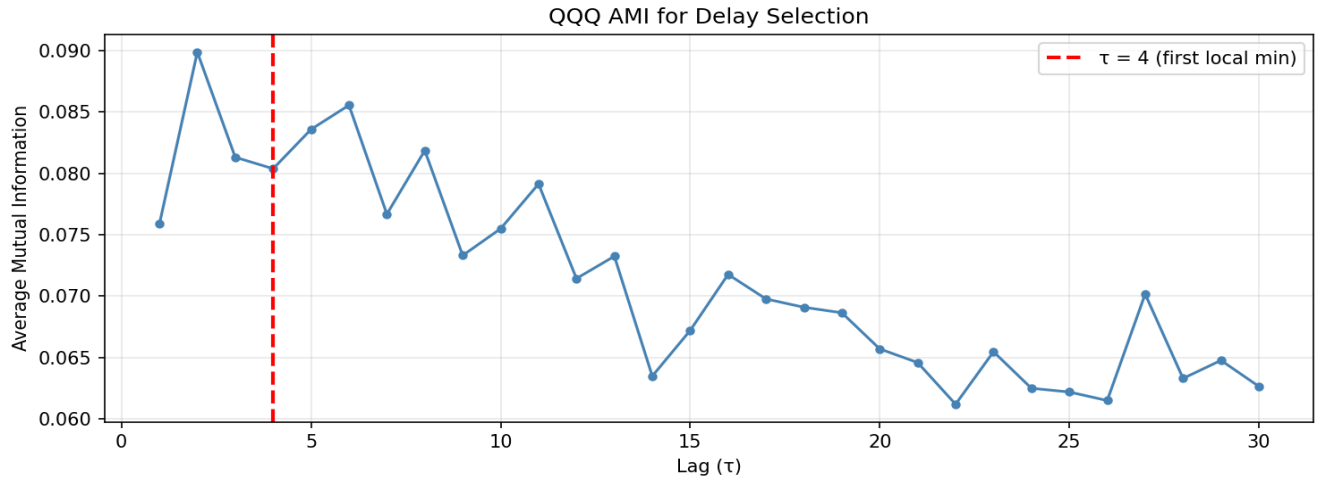

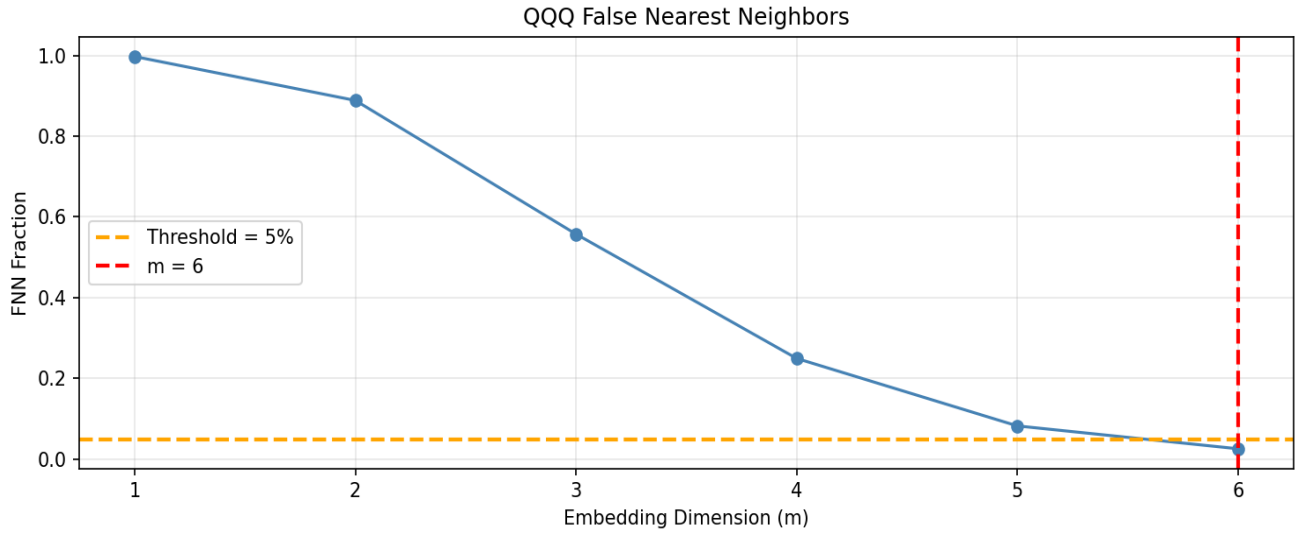

[Figures 7,8: AMI and FNN plots for QQQ]

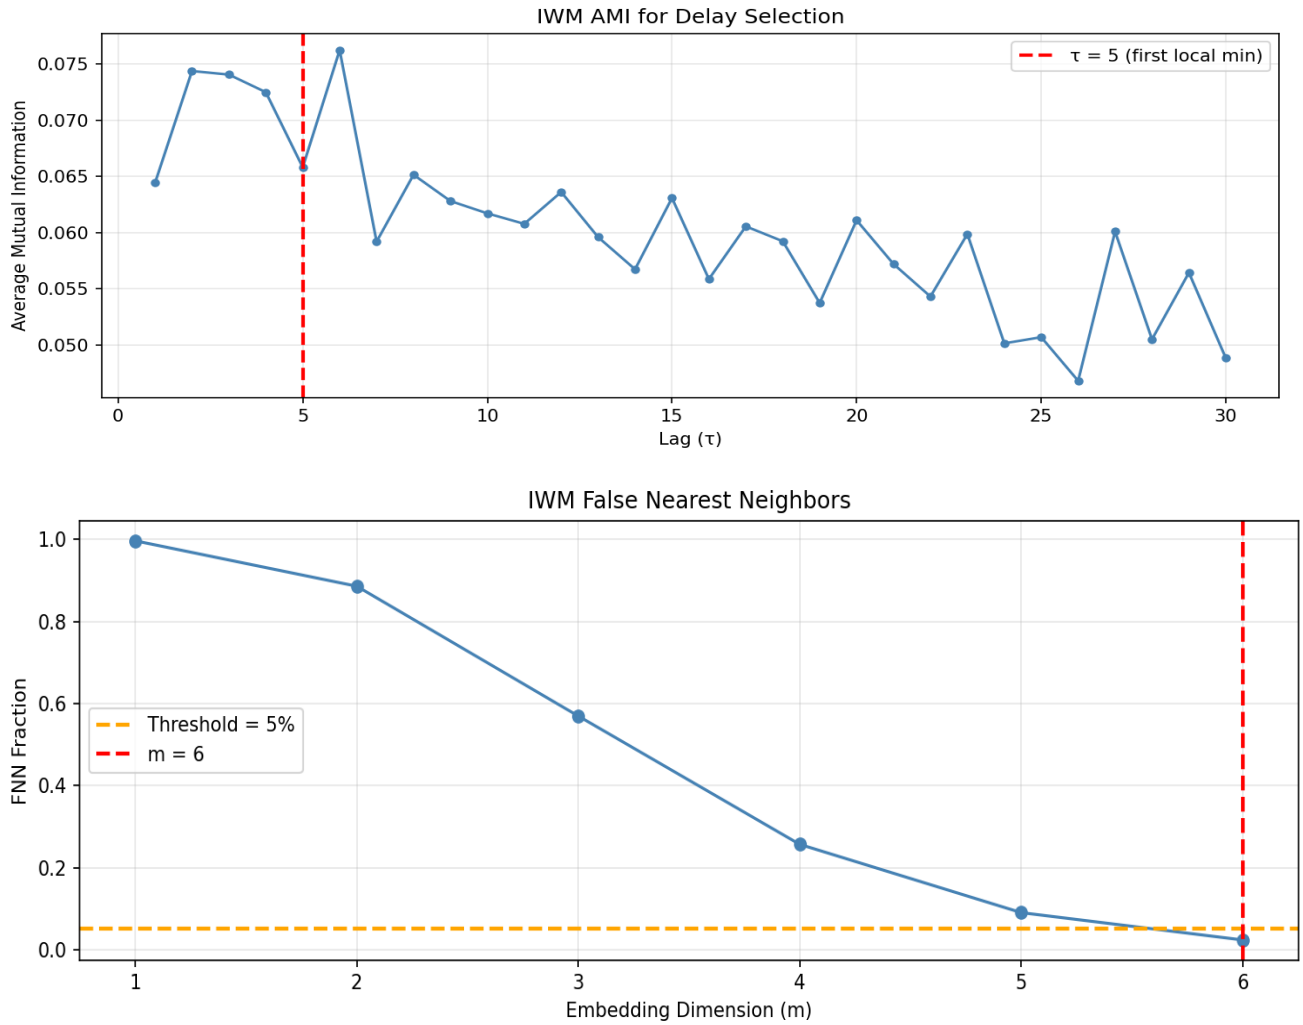

[Figures 9,10: AMI and FNN plots for IWM]

5.3 FTLE Time Series

FTLE time series are computed for all three assets at both window sizes. Median FTLE values are consistently positive across all configurations, indicating net trajectory divergence in the reconstructed phase space. Negative FTLE values, indicating local trajectory convergence, occur occasionally but are less frequent at the 500-day window than at the 250-day window.

The accompanying figures (Figures 11–16) mark the 2008–2009 financial crisis and 2020 market disruption for visual reference. Systematic investigation of FTLE behaviour during these and other periods, including analysis of how parameter choices (τ, m, T, k, W) affect the FTLE time series and potential correlations with volatility indices, is reserved for future work.

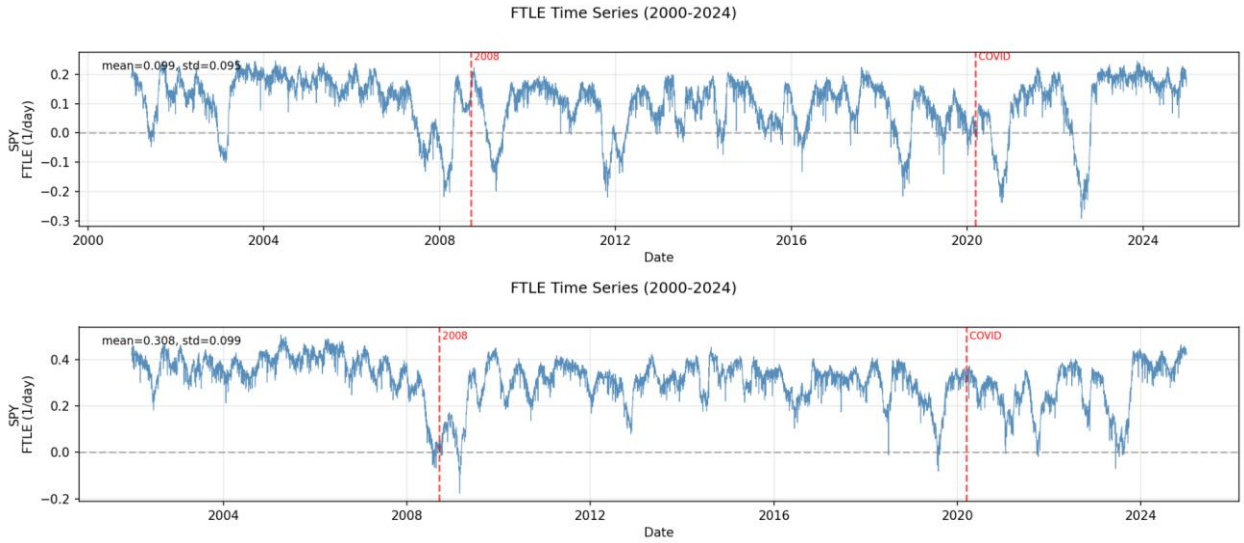

[Figures 11,12: FTLE time series for SPY at W = 250 and W = 500]

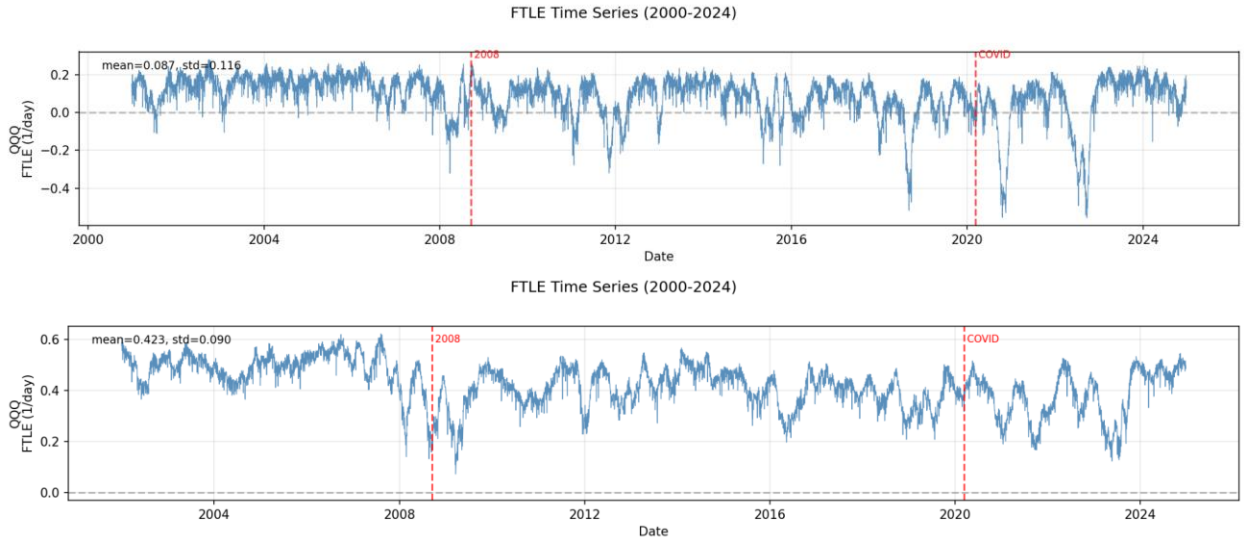

[Figures 13,14: FTLE time series for QQQ at W = 250 and W = 500]

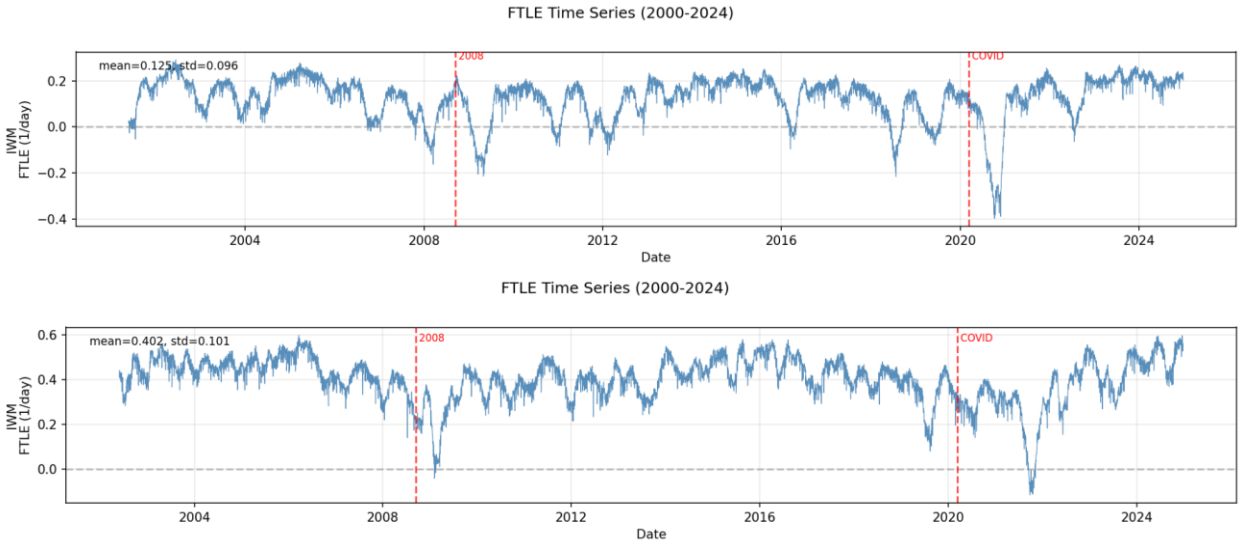

[Figures 15,16: FTLE time series for IWM at W = 250 and W = 500]

5.4 Surrogate Tests

Surrogate data testing yields the following p-values for the null hypothesis that the observed FTLE structure is consistent with a linear stochastic generating process. For SPY: p = 0.152 (250-day window), p = 0.119 (500-day window). For QQQ: p = 0.097 (250-day window), p = 0.091 (500-day window). For IWM: p = 0.156 (250-day window), p = 0.067 (500-day window).

At the conventional significance level of p = 0.05, none of the six configurations rejects the null hypothesis. The observed FTLE values are statistically consistent with those that would arise from linearly filtered Gaussian noise. However, at the relaxed threshold of p = 0.10, three of six configurations—QQQ at both window sizes and IWM at the 500-day window—show marginal rejection. This suggests that while the data are broadly consistent with linear stochastic processes, nonlinear structure cannot be ruled out and may be present in certain assets or parameter configurations.

6. Discussion

6.1 Validation

The algorithm recovers Lyapunov exponents consistent with known values for the Lorenz system across both parameter configurations ($\lambda_1 \approx 0.9$ nats per time unit for Configuration A; $\lambda_1 \approx 1.5$ nats per time unit for Configuration B), confirming correct construction and implementation of the computational pipeline. However, the Lorenz system is an idealised test case: it is deterministic, noise-free, and stationary. Financial market data presents a multitude of additional challenges including measurement noise, non-stationarity, and fundamental uncertainty regarding the nature of the underlying generating process.

6.2 Observations

FTLE computation on equity log-returns is technically feasible across all three ETFs, yielding time series with consistently positive median values. Occasional negative FTLE excursions are observed; systematic investigation of their relationship to market conditions and parameter choices is warranted. The 250-day and 500-day windows yield qualitatively different FTLE behaviour, reflecting the inherent trade-off between temporal resolution and statistical stability in non-stationary systems. Surrogate data testing produced mixed results: no configuration rejected the null hypothesis at $p = 0.05$, though three of six configurations rejected at $p = 0.10$. The data are broadly consistent with linear stochastic processes, though nonlinear structure cannot be ruled out and motivates further investigation.

6.3 Limitations

Several limitations constrain the conclusions that can be drawn from this work. Results may be sensitive to parameter choices (τ, m, T, k, W); systematic sensitivity analysis across the full parameter space was not exhaustively performed. No individual asset rejected the surrogate null hypothesis at the conventional $p = 0.05$ threshold; marginal rejection was observed only at $p = 0.10$. No out-of-sample, walk-forward, or subperiod robustness testing was conducted at this stage. The physical interpretation of FTLE values in financial markets—what trajectory divergence in reconstructed phase space means for market dynamics—remains unclear.

6.4 Open Questions

Several questions emerge from this exploratory work. What drives the occasional negative FTLE excursions—genuine dynamical convergence in the underlying system, or artefacts of the estimation procedure? How does FTLE behaviour vary across the parameter space (τ, m, T, k, W), and are there regions of greater stability? Why do QQQ and IWM show marginal surrogate rejection while SPY does not—does this reflect genuine differences in market microstructure, or statistical fluctuation? Does FTLE exhibit systematic relationships with market observables such as volatility indices, and if so, what is the causal structure? Would alternative embedding methods or longer datasets yield stronger evidence for nonlinear structure?

7. Conclusion

This report documents the construction, validation, and exploratory application of a Finite-Time Lyapunov Exponent computation algorithm. The algorithm was validated on the Lorenz system, recovering values consistent with known Lyapunov exponents across two parameter configurations and two numerical integration methods. The method was then applied to daily log-returns of three US equity ETFs over a 24-year period at two temporal scales. FTLE time series were computed and characterised, revealing consistently positive median values with occasional negative excursions. Surrogate data testing provided no evidence for nonlinear structure at $p = 0.05$, though marginal rejection at $p = 0.10$ for select configurations suggests that nonlinear dynamics cannot be ruled out. This work establishes a methodological foundation; further investigation with systematic parameter analysis, alternative embedding methods, and

extended datasets is required before stronger conclusions can be drawn. The work motivates investigations of correlation between FTLE time series and market observables such as volatility indices.

References

[1] Hsieh, D.A. (1991). Chaos and nonlinear dynamics: Application to financial markets. Journal of Finance, 46(5), 1839–1877.
[2] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327.
[3] Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. The Review of Economics and Statistics, 51(3), 247–257.
[4] Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance, 1(2), 223–236.
[5] Ott, E. (2002). Chaos in Dynamical Systems. Cambridge University Press
[6] Peters, E.E. (1991). Chaos and Order in the Capital Markets. John Wiley & Sons.
[7] Brock, W.A., Hsieh, D.A. & LeBaron, B. (1991). Nonlinear Dynamics, Chaos, and Instability. MIT Press.
[8] Ruelle, D. (1994). Where can one hope to profitably apply the ideas of chaos? (Available via IHES or reprinted collections).
[9] Diks, C. G. H., Hommes, C. H., & Wang, J. (2019). Critical slowing down as early warning signals for financial crises? Empirical Economics, 57(4), 1201–1228.
[10] Shadden, S.C., Lekien, F. & Marsden, J.E. (2005). Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. Physica D, 212(3–4), 271–304.
[11] Eckmann, J.-P. & Ruelle, D. (1985). Ergodic theory of chaos and strange attractors. Reviews of Modern Physics, 57(3), 617–656.
[12] Takens, F. (1981). Detecting strange attractors in turbulence. In D.A. Rand & L.-S. Young (Eds.), Dynamical Systems and Turbulence, Warwick 1980, Lecture Notes in Mathematics, Vol. 898 (pp. 366–381). Springer.
[13] Haller, G. (2015). Lagrangian coherent structures. Annual Review of Fluid Mechanics, 47, 137–162.
[14] Hull, J. C. (2021). Options, futures, and other derivatives. Pearson.
[15] Eckmann, J.-P., Kamphorst, S. O., Ruelle, D., & Ciliberto, S. (1986). Liapunov exponents from time series. Physical Review A, 34(6), 4971–4979.
[16] Theiler, J. (1986). Spurious dimension from correlation algorithms applied to limited time-series data. Physical Review A, 34(3), 2427–2432.
[17] Kantz, H. & Schreiber, T. (2004). Nonlinear Time Series Analysis (2nd ed.). Cambridge University Press.
[18] Lorenz, E.N. (1963). Deterministic nonperiodic flow. Journal of the Atmospheric Sciences, 20(2), 130–141.
[19] Viswanath, D. (1998). Lyapunov exponents from random Fibonacci sequences to the Lorenz equations (Doctoral dissertation). Cornell University.

[20] Wolf, A., Swift, J. B., Swinney, H. L., & Vastano, J. A. (1985). Determining Lyapunov exponents from a time series. Physica D: Nonlinear Phenomena, 16(3), 285–317.
[21] Fraser, A.M. & Swinney, H.L. (1986). Independent coordinates for strange attractors from mutual information. Physical Review A, 33(2), 1134–1140.
[22] Kennel, M.B., Brown, R. & Abarbanel, H.D.I. (1992). Determining embedding dimension for phase-space reconstruction using a geometrical construction. Physical Review A, 45(6), 3403–3411.
[23] Theiler, J. (1990). Estimating fractal dimension. Journal of the Optical Society of America A, 7(6), 1055–1073.
[24] Theiler, J., Eubank, S., Longtin, A., Galdrikian, B. & Farmer, J.D. (1992). Testing for nonlinearity in time series: The method of surrogate data. Physica D, 58(1–4), 77–94.
[25] Schreiber, T. & Schmitz, A. (1996). Improved surrogate data for nonlinearity tests. Physical Review Letters, 77(4), 635–638.

Appendix

A. Glossary

Financial Terms

- ETF: Exchange-traded fund tracking index or asset basket
- Log returns: $r_t = \ln(P_t/P_{t-1})$
- Trading day: Day on which exchanges are open; approximately 252 per calendar year

Dynamical Systems Terms

- Phase space: Space of all possible system states
- Trajectory: Path through phase space
- Attractor: Invariant set toward which trajectories converge
- Lyapunov exponent: Asymptotic rate of trajectory divergence
- FTLE: Finite-time, local measure of trajectory divergence rate
- Embedding: Phase space reconstruction from scalar time series
- Chaos: Bounded dynamics with sensitive dependence on initial conditions ($\lambda_1 > 0$)
- Theiler window: Temporal exclusion preventing spurious neighbour correlations

Statistical Terms

- Surrogate data: Synthetic time series for hypothesis testing
- AAFT: Amplitude-Adjusted Fourier Transform surrogate method

B. Parameter Tables

Final parameters per asset						
	tau	m	W	T	Theiler	k
lorenz_A_Euler	16	3	9000	5000	32	100
lorenz_A_RK4	16	3	9000	5000	32	100
lorenz_B_Euler	10	3	9000	5000	20	100
lorenz_B_RK4	10	3	9000	5000	20	100
SPY_250	3	6	250	60	6	75
SPY_500	3	6	500	50	6	50
QQQ_250	4	6	250	30	8	75
QQQ_500	4	6	500	60	8	30
IWM_250	5	6	250	60	10	75
IWM_500	5	6	500	50	10	40