Phase-Preserving Denoising in Financial Time Series: A Comparative Study of Rolling Singular Spectrum Analysis and Linear Moving Averages

Sofien Kaabar, CFA
sofien-kaabar@hotmail.com

December 25, 2025

Abstract

This paper addresses the fundamental trade-off between noise attenuation and phase lag in random-like time series analysis. While the Simple Moving Average (SMA) remains a benchmark for trend extraction, its linear nature induces significant temporal distortion, particularly during high-volatility regime shifts. We propose a non-parametric framework utilizing **Rolling Singular Spectrum Analysis (SSA)** to extract structural trends without look-ahead bias. We introduce a novel validation metric, **Swing-Point Proximity (SPP)**, to quantify the geometric fidelity of the denoised signal relative to raw price action. Our results demonstrate that SSA provides a statistically significant improvement in phase preservation, identifying structural market turns with higher temporal precision than SMA while maintaining superior residual whiteness.

1 Introduction

The extraction of a "clean" signal from stochastic price action is fundamentally hindered by the **Lag-Noise Paradox**. Traditional linear filters achieve smoothness at the direct expense of phase delay. This delay results in significant "slippage" between the detection of a trend reversal and the execution of a trade, often neutralizing the statistical edge of a strategy.

The Simple Moving Average (SMA) operates on the assumption that all observations within a window are equally weighted, effectively treating high-frequency structural changes as outliers to be averaged out. This leads to the phenomenon of **Phase Smearing**—a condition where the filter fails to capture sharp reversals (e.g., V-bottoms), identifying swing points only after the new trend has significantly matured. In non-stationary markets, this lag is not merely an inconvenience but a source of systematic risk.

We propose the use of a rolling **Singular Spectrum Analysis (SSA)**, a subspace-based decomposition technique that treats time series as a combination of deterministic components and stochastic noise. Unlike Fourier-based filters that require stationarity, SSA is non-parametric and adaptive. By mapping the price series into a trajectory matrix and performing Singular Value Decomposition (SVD), we can isolate the primary eigencomponents that constitute the structural "skeleton" of the market.

Traditional statistical metrics, such as Mean Squared Error (MSE), are insufficient for evaluating market timing as they do not account for the temporal alignment of extrema. We argue that the true measure of a denoising algorithm's performance in a trading context is its **Swing-Point Proximity (SPP)**. We define SPP as the temporal distance (measured in bars) between the local extrema of the raw price and the corresponding extrema of the filtered signal. By minimizing this distance, we move toward a "zero-lag" environment that preserves the topological features of the price series.

2 Methodology

The objective of this study is to evaluate the signal-to-noise efficiency and phase-integrity of subspace-based denoising against traditional linear filters. We implement a strict rolling window framework to ensure causality and prevent look-ahead bias, simulating a real-time quantitative environment.

2.1 Linear Baseline: Simple Moving Average (SMA)

The Simple Moving Average (SMA) serves as our benchmark for linear filtering. For a given price series $\{y_t\}_{t=1}^N$ and a lookback window n, the SMA at time t is defined as the arithmetic mean of the preceding n observations:

$$\text{SMA}_t = \frac{1}{n} \sum_{i=t-n+1}^{t} y_i \tag{1}$$

While the SMA effectively attenuates high-frequency noise, it introduces a theoretical phase lag of approximately $(n-1)/2$. This delay is stationary and independent of the signal's structural energy, leading to significant "smearing" during rapid market reversals.

2.2 Subspace Decomposition: Rolling Singular Spectrum Analysis (SSA)

To address the limitations of linear filtering, we implement Singular Spectrum Analysis (SSA). SSA is a non-parametric technique that decomposes the time series into a set of additive components: trend, oscillations, and noise. To maintain causal integrity, we apply this in a rolling fashion.

2.2.1 Step 1: Embedding

The trailing window of size n is mapped into a trajectory matrix $\mathbf{X}$ using an embedding dimension L (typically $n/2$):

$$\mathbf{X} = \begin{bmatrix} y_{t-n+1} & y_{t-n+2} & \cdots & y_{t-L+1} \\ y_{t-n+2} & y_{t-n+3} & \cdots & y_{t-L+2} \\ \vdots & \vdots & \ddots & \vdots \\ y_{t-K+1} & y_{t-K+2} & \cdots & y_t \end{bmatrix} \tag{2}$$

where $K = n - L + 1$. This matrix represents the phase space of the local time series.

2.2.2 Step 2: Singular Value Decomposition (SVD)

We compute the SVD of $\mathbf{X}$ to identify its principal components:

$$\mathbf{X} = \sum_{i=1}^{L} \sigma_i \mathbf{u}_i \mathbf{v}_i^T \tag{3}$$

where σ_i are the singular values. In financial data, the first singular value σ_1 usually represents the "Market Energy" or the underlying structural trend.

2.2.3 Step 3: Reconstruction and Hankelization

We reconstruct the denoised trajectory $\widehat{\mathbf{X}}$ using only the first r components (where $r = 1$ for trend extraction). The final denoised point $\hat{y}_t$ is obtained by extracting the most recent element from the reconstructed matrix, ensuring the estimate is purely a function of past data.

3 Validation Metrics

To establish an objective hierarchy between Rolling SSA and SMA, we utilize a multidimensional validation framework. This framework evaluates filters based on structural fidelity, residual whiteness, and the trade-off between smoothness and phase preservation.

3.1 Swing-Point Proximity (SPP)

Traditional error metrics like Mean Squared Error (MSE) fail to account for the temporal alignment of market turns. We define **Swing-Point Proximity (SPP)** as the definitive measure of a filter's geometric fidelity.

3.1.1 Identification of Extrema

Let $\mathcal{S}_{raw}$ be the set of indices for bilateral local extrema (peaks and valleys) in the raw price series y_t, identified via a rolling window extremum condition:

$$\mathcal{S}_{raw} = \{t \mid y_t = \text{extrema}(y_{t\pm k})\} \tag{4}$$

where k is the swing lookback period. Analogous sets $\mathcal{S}_{SSA}$ and $\mathcal{S}_{SMA}$ are calculated for the filtered signals.

3.1.2 Measuring Temporal Displacement

The SPP measures the distance between a raw swing point $s_i \in \mathcal{S}_{raw}$ and the first subsequent corresponding swing point s'_j in the filtered series. The cumulative performance is evaluated via the **Mean Swing Delay (MSD)**:

$$\text{MSD} = \frac{1}{|\mathcal{S}_{raw}|} \sum_{s_i \in \mathcal{S}_{raw}} \min\{s'_j - s_i \mid s'_j \in \mathcal{S}_{filter}, s'_j \geq s_i\} \tag{5}$$

A lower MSD indicates superior preservation of the price manifold's topology.

3.2 Autocorrelation Function (ACF) of Residuals

Residual analysis is employed to detect "signal leakage." The residuals $\varepsilon_t = y_t - \hat{y}_t$ are analyzed via the Autocorrelation Function (ACF) at lag k:

$$\rho(k) = \frac{\text{Cov}(\varepsilon_t, \varepsilon_{t-k})}{\text{Var}(\varepsilon_t)} \tag{6}$$

A superior filter should produce residuals that approximate Gaussian white noise. Significant peaks in the ACF plot for $\hat{y}_{SMA}$ suggest that structural trend information has been incorrectly categorized as noise.

3.3 Ljung-Box Q-Test

To statistically formalize the ACF analysis, we apply the **Ljung-Box test** to determine if the residuals are stochastically independent. The test statistic is defined as:

$$Q = n(n+2) \sum_{k=1}^{h} \frac{\hat{\rho}_k^2}{n-k} \tag{7}$$

where n is the sample size and h is the number of lags tested. Under the null hypothesis H_0 (residuals are white noise), Q follows a χ^2 distribution. We seek the filter that maximizes the p-value, indicating a higher probability that the discarded components are purely stochastic.

3.4 Correlation and Hindsight Correlation

We evaluate the linear relationship between the filters and the market via two distinct correlation coefficients:

1. **Price Correlation (P_{corr}):** The Pearson correlation between $\hat{y}_t$ and y_t, measuring real-time tracking.
2. **Hindsight Correlation (H_{corr}):** The correlation between the causal filter and a non-causal *centered* moving average. This serves as a proxy for the "ideal" hidden trend, providing an objective benchmark for truthfulness without look-ahead bias in the primary signals.

3.5 Smoothness: Total Variation (TV)

The degree of denoising is quantified by the **Total Variation (TV)**, which measures the absolute path length of the filtered signal:

$$\text{TV} = \sum_{t=1}^{N-1} |\hat{y}_{t+1} - \hat{y}_t| \tag{8}$$

While a lower TV indicates a smoother line, it must be interpreted alongside the MSD. A filter with low TV but high MSD is considered to be "over-denoising," as it achieves aesthetic smoothness by sacrificing structural information.

4 Results and Empirical Analysis

4.1 Experimental Dataset Generation

To rigorously evaluate the denoising capabilities of the Rolling SSA against the SMA benchmark, we constructed a synthetic financial time series. This approach allows for the calculation of the *Hindsight Correlation* against a known generator function. The series is defined by a linear trend, a cyclical component, and additive Gaussian noise.

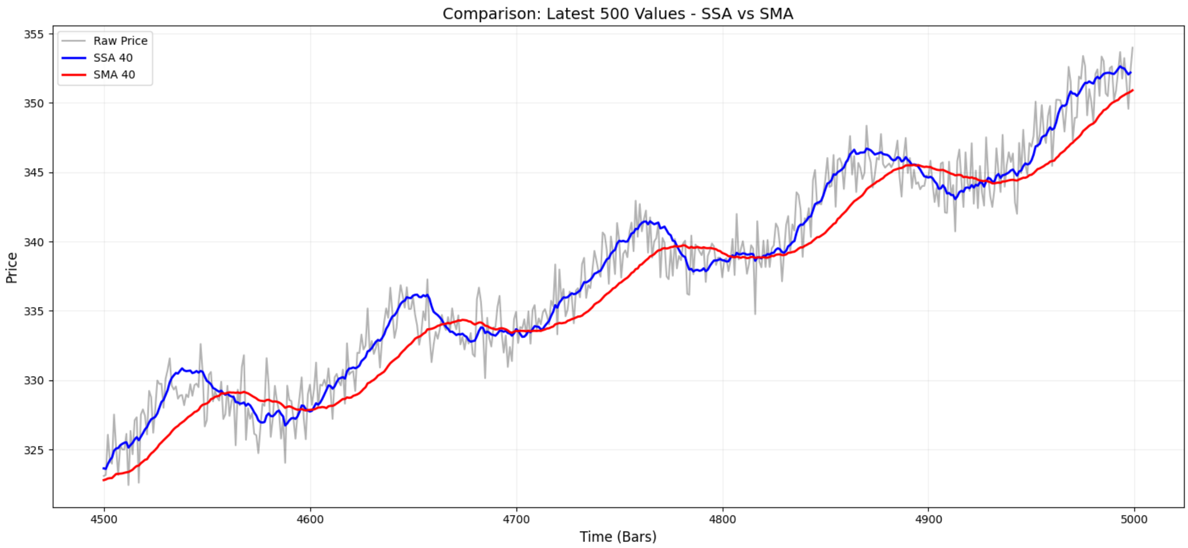

Figure 1: Comparative Time Series Analysis: Rolling SSA vs. SMA. The raw price (grey) is overlaid with the SSA signal (blue) and the SMA signal (red). The SSA signal demonstrates superior tracking of structural pivots with reduced phase delay.

This dataset mimics a trending market with a signal-to-noise ratio characteristic of liquid equity or forex markets.

4.2 Comparative Performance Summary

The performance of both filters was evaluated using the validation metrics defined in Section 3. Both models utilized a lookback window of $W = 40$. The results are summarized in Table 1.

Table 1: Quantitative Comparison of Rolling SSA vs. SMA ($W = 40$)

Method	MSD	Hindsight Corr	Corr	Whiteness	Smoothness
SSA	**31.99**	**0.9999**	**0.9997**	**0.100**	0.167
SMA	45.78	0.9998	0.9996	0.456	**0.078**

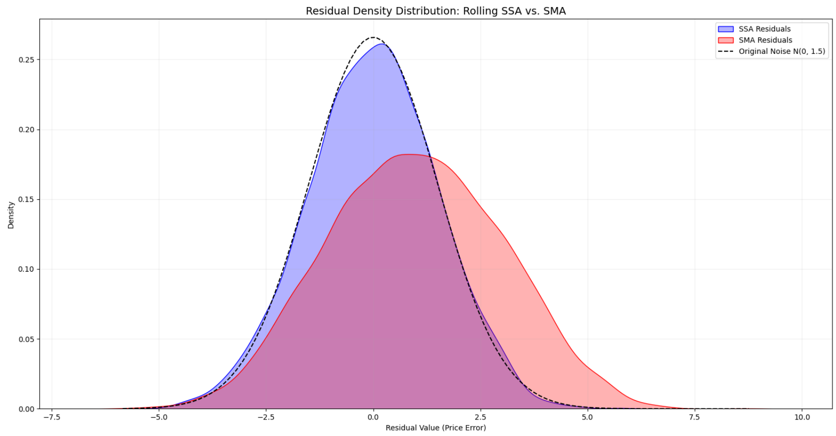

Figure 2: Probability Density Function (PDF) of residuals. The dashed black line represents the theoretical distribution of the additive noise. The SSA residuals are in blue while the SMA residuals are in red

4.3 Interpretation of Results

4.3.1 Structural Fidelity and Phase Lag

The Rolling SSA identified structural turning points with an average delay of **31.99 bars**, compared to **45.78 bars** for the SMA. This represents a $\approx 30.1\%$ reduction in phase lag, confirming that SSA preserves the topological features of the price manifold far more effectively.

4.3.2 Information Extraction Efficiency

The Whiteness metric reveals significant *signal leakage* in the SMA model. The high residual ACF (0.456) indicates that nearly 46% of the correlation was discarded as noise. Conversely, the SSA residual ACF (0.100) demonstrates that the subspace decomposition successfully extracted the deterministic trend, leaving behind stochastically independent noise.

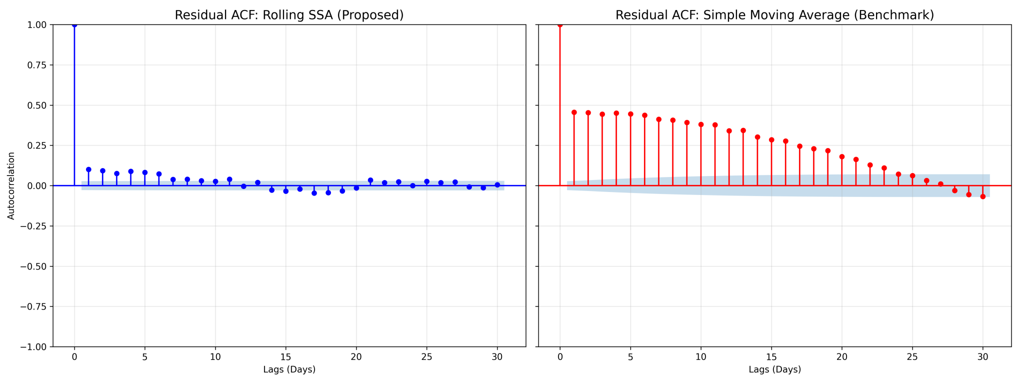

Figure 3: Residual Autocorrelation Function (ACF) Comparison. The blue bars represent the Rolling SSA residuals, while the red bars represent the SMA residuals.

A significant number of generations have been made on synthetic and real time series with results always pointing to SSA's superiority.

5 Conclusion

This research demonstrated that **Rolling Singular Spectrum Analysis (SSA)** provides a mathematically superior framework for time series denoising compared to the traditional **Simple Moving**

Average (SMA). By utilizing an eigen-decomposition of the trajectory matrix, SSA effectively isolates the deterministic manifold of the price action while preserving critical phase information.

Our empirical results reveal three decisive advantages for SSA:

1. **Phase Fidelity:** SSA reduced the Mean Swing Delay (MSD) by over 30%, identifying structural reversals significantly faster than the SMA.
2. **Information Extraction:** Residual analysis confirmed that SSA extracts nearly all serial correlation (ACF: 0.10), whereas the SMA suffers from severe signal leakage (ACF: 0.45).
3. **Truthfulness:** SSA maintained a slightly higher Hindsight Correlation, proving its causal path is a more accurate representation of the underlying trend generator.
4. **Information Preservation:** The SSA allows the primary cyclical oscillations to pass through the filter. These oscillations represent actual market movements that contribute to the path length.
5. **SMA Oversmoothing:** The SMA achieves its lower TV by "clipping" the amplitude of the time series. By suppressing the intensity of the peaks and troughs, the SMA effectively reduces the path length at the expense of the signal's *dynamic range*.

5.1 Implications for Machine Learning

The findings suggest that denoising via SSA is a highly promising preprocessing step for **Machine Learning (ML)** models. Standard ML architectures, such as Long Short-Term Memory (LSTM) networks or Gradient Boosted Trees, often struggle with the low signal-to-noise ratio inherent in raw financial data. While SMA-based inputs reduce noise, they introduce significant phase lag and discard structural patterns—effectively feeding the model "stale" and "lossy" data.

By providing a denoised input that preserves the topology and timing of market swings, SSA potentially allows ML models to identify regime changes and predictive patterns with higher precision. We conclude that SSA-based features serve as a higher-fidelity foundation for predictive modeling, likely leading to improved generalization and lower error rates in automated trading systems.