Eliminating Subjectivity in Moving Average Crossovers via Symmetric Weighted Filters

Sofien Kaabar, CFA
sofien-kaabar@hotmail.com

January 18, 2026

Abstract

Moving average crossover strategies remain among the most widely used tools in technical analysis, yet their construction is often rooted in heuristic choices that introduce a significant layer of subjectivity. Traditional crossover systems typically rely on combining two moving averages with different lookback windows, implicitly embedding arbitrary parameter selection and temporal asymmetry into the signal generation process.

This paper proposes an alternative framework that removes this source of subjectivity by constructing crossover signals from two symmetric filters computed over an identical lookback horizon: a Weighted Moving Average (WMA) and an Inverse Weighted Moving Average (IWMA). By assigning increasing weights to recent observations in the WMA and decreasing weights in the IWMA, both filters capture complementary temporal perspectives of the same information set. A trading signal is generated exclusively when the WMA crosses the IWMA, yielding a directionally interpretable and parameter-minimal decision rule.

The proposed approach preserves the intuitive appeal of moving average crossovers while enforcing structural symmetry, interpretability, and reduced arbitrariness. Empirical analysis demonstrates that this formulation produces stable and timely signals without requiring multiple lookback choices or ad-hoc parameter tuning.

1 Introduction

Moving averages constitute a foundational component of technical analysis and algorithmic trading systems. Their primary appeal lies in their simplicity, interpretability, and ability to smooth noisy financial time series. Among their many applications, moving average crossovers are particularly popular, serving as trend-following signals in discretionary and systematic trading alike.

Despite their widespread adoption, conventional moving average crossover strategies suffer from a fundamental methodological weakness: subjectivity in construction. Standard implementations typically rely on pairing two moving averages with distinct lookback windows, such as a "fast" and a "slow" average. The choice of these horizons is rarely derived from first principles and often reflects heuristic reasoning, curve fitting, or domain conventions. As a result, crossover signals implicitly encode arbitrary temporal assumptions that can materially affect both signal timing and performance.

Beyond parameter arbitrariness, traditional crossover systems introduce structural asymmetry. The fast moving average reacts more rapidly to recent price changes, while the slow moving average embeds a longer historical memory. This asymmetry complicates interpretation, as the resulting signal conflates differences in weighting schemes with differences in information sets.

This paper proposes a structurally symmetric alternative that removes these sources of subjectivity. Instead of combining two moving averages with different lookback windows, we construct two complementary filters over the same horizon: a Weighted Moving Average (WMA), which emphasizes recent observations through increasing linear weights, and an Inverse Weighted Moving Average (IWMA), which applies the same weighting scheme in reverse order. Both filters operate on an identical information set, differing only in how temporal importance is assigned.

Within this framework, a bullish signal is generated when the WMA crosses above the IWMA, indicating that recent price dynamics dominate earlier observations. Conversely, a bearish signal arises when the WMA crosses below the IWMA. This formulation yields a crossover mechanism that is fully determined by a single lookback parameter and a symmetric weighting structure, significantly reducing subjective design choices.

By reframing moving average crossovers as interactions between dual temporal weightings rather than heterogeneous horizons, this approach offers a more principled and interpretable foundation for trend detection. The remainder of this paper formalizes the proposed filters, examines their theoretical properties, and evaluates their empirical behavior across multiple financial time series.

2 Methodology

2.1 Moving Averages as Temporal Filters

Let $\{P_t\}_{t=1}^{T}$ denote a univariate price time series observed at discrete times t. A moving average of lookback length N is a linear filter that maps the last N observations into a single smoothed value at time t. In its most general form, a moving average can be written as

$$MA_t = \sum_{i=0}^{N-1} w_i P_{t-i}, \tag{1}$$

where $\{w_i\}_{i=0}^{N-1}$ are non-negative weights satisfying $\sum_{i=0}^{N-1} w_i = 1$.

Different choices of the weighting vector $\mathbf{w}$ induce different temporal emphases on past observations. In this section, we introduce three commonly used specifications: the Simple Moving Average (SMA), the Weighted Moving Average (WMA), and the proposed Inverse Weighted Moving Average (IWMA).

2.2 Simple Moving Average (SMA)

The Simple Moving Average assigns equal weight to all observations within the lookback window. Its weights are defined as

$$w_i^{\text{SMA}} = \frac{1}{N}, \quad i = 0, \dots, N-1. \tag{2}$$

The resulting SMA is given by

$$\text{SMA}_t = \frac{1}{N} \sum_{i=0}^{N-1} P_{t-i}. \tag{3}$$

While the SMA provides a useful baseline for smoothing, it implicitly assumes that all past observations within the window carry identical informational relevance, an assumption that may be overly restrictive in the presence of evolving market dynamics.

2.3 Weighted Moving Average (WMA)

To account for temporal relevance, the Weighted Moving Average assigns linearly increasing weights to more recent observations. The WMA weights are defined as

$$w_i^{\text{WMA}} = \frac{N-i}{\sum_{j=1}^{N} j} = \frac{2(N-i)}{N(N+1)}, \quad i = 0, \ldots, N-1, \tag{4}$$

where $i = 0$ corresponds to the most recent observation P_t.

The WMA is then expressed as

$$\text{WMA}_t = \sum_{i=0}^{N-1} w_i^{\text{WMA}} P_{t-i}. \tag{5}$$

This formulation emphasizes recent price action while still incorporating historical information, producing a smoother yet more responsive filter than the SMA.

2.4 Inverse Weighted Moving Average (IWMA)

We now introduce the Inverse Weighted Moving Average (IWMA), which applies the same linear weighting scheme as the WMA, but in reverse temporal order. Its weights are defined as

$$w_i^{\text{IWMA}} = \frac{i+1}{\sum_{j=1}^{N} j} = \frac{2(i+1)}{N(N+1)}, \quad i = 0, \ldots, N-1. \tag{6}$$

The IWMA is therefore given by

$$\text{IWMA}_t = \sum_{i=0}^{N-1} w_i^{\text{IWMA}} P_{t-i}. \tag{7}$$

Unlike the WMA, the IWMA assigns greater importance to older observations within the same lookback window. Importantly, both filters operate on an identical information set and differ exclusively in their temporal weighting structure.

2.5 WMA–IWMA Crossovers as a Temporal Dominance Test

Traditional moving average crossovers compare filters constructed over different lookback horizons, thereby conflating differences in weighting schemes with differences in information sets. In contrast, the WMA and IWMA defined above share the same window length N and form a symmetric pair of temporal filters.

Consider the difference between the two filters:

$$D_t = \text{WMA}_t - \text{IWMA}_t. \tag{8}$$

Substituting Equations (5) and (7) yields

$$D_t = \sum_{i=0}^{N-1} \left(w_i^{\text{WMA}} - w_i^{\text{IWMA}} \right) P_{t-i}. \tag{9}$$

Figure 1: Tesla with a 60-period WMA (blue) and a 60-period IWMA (red). The chart displays the difference between WMA and IWMA and how the former follows closely the recent observtions while the latter gives less weight to them.

Since $w_i^{\text{WMA}} > w_i^{\text{IWMA}}$ for small i (recent observations) and $w_i^{\text{WMA}} < w_i^{\text{IWMA}}$ for large i (older observations), the sign of D_t directly reflects whether recent prices dominate earlier prices within the window.

A positive value of D_t implies that the price process exhibits upward temporal dominance, meaning that recent observations exert greater influence than older ones. Conversely, a negative value indicates downward temporal dominance. A crossover event, defined by $D_t = 0$, therefore represents a transition point where the balance of influence shifts between recent and historical price behavior.

From this perspective, the WMA–IWMA crossover is not a heuristic trend-following signal, but rather a deterministic test of temporal dominance under a fixed information set. The signal is generated without introducing additional horizons, regime assumptions, or ad-hoc parameter choices, thereby reducing structural subjectivity while preserving interpretability.

This formulation reframes moving average crossovers as a comparison between opposing temporal emphases rather than competing time scales, providing a principled and symmetric foundation for trend detection.

3 Backtesting Framework and Benchmark Selection

3.1 Data and Backtesting Horizon

The empirical evaluation is conducted on hourly foreign exchange spot data, covering six major currency pairs: EURUSD, USDCHF, GBPUSD, USDJPY, USDCAD, and AUDUSD. These instruments were selected due to their high liquidity, tight spreads, and structural diversity, making them well suited for evaluating trend detection mechanisms across different market dynamics.

All time series are sampled at an hourly frequency in order to balance signal responsiveness with noise reduction. The backtesting horizon spans a sufficiently long historical period to capture multiple market regimes, including trending, range-bound, and high-volatility environments. To avoid look-ahead bias, all signals are computed using information available up to time t, with trades executed at the subsequent bar.

The strategy operates in a symmetric long–short framework: a long position is initiated following a bullish signal, while a short position is initiated following a bearish signal. Positions are reversed upon signal inversion. Transaction costs are applied uniformly across all strategies to ensure fair comparison.

3.2 Benchmark Strategies

To assess the incremental value of the proposed WMA–IWMA crossover, we compare its performance against two benchmark strategies that preserve a comparable information budget and parameter complexity.

WMA(60) versus SMA(60). The first benchmark compares a Weighted Moving Average and a Simple Moving Average constructed over the same lookback horizon of $N = 60$. This benchmark isolates the effect of differential weighting schemes while holding the information set constant. By contrasting linear recency weighting against uniform weighting, it evaluates whether the inverse weighting structure introduced by the IWMA provides additional explanatory power beyond conventional smoothing techniques.

Price versus WMA(60). The second benchmark is a price–trend filter based on the crossing of the spot price and a WMA(60). This formulation represents a widely used trend-following mechanism in both discretionary and systematic trading. It serves as a minimal baseline that tests whether the dual-filter crossover approach improves upon a single moving average signal derived from the same lookback window.

Both benchmarks rely on a single lookback parameter and avoid the fast–slow moving average paradigm, ensuring that differences in performance are attributable to signal structure rather than horizon selection.

3.3 Performance Metrics

Strategy performance is evaluated using three complementary metrics that capture predictive accuracy, risk-adjusted returns, and trade efficiency.

Hit Ratio. The hit ratio measures the proportion of profitable trades relative to the total number of trades executed. It provides a direct assessment of directional accuracy but does not account for the magnitude of gains and losses.

Sharpe Ratio. The Sharpe ratio evaluates risk-adjusted performance by scaling the mean strategy return by its standard deviation. It captures both return consistency and volatility and is particularly relevant when comparing strategies with similar exposure profiles.

Profit Factor. The profit factor is defined as the ratio of gross profits to gross losses. It reflects the efficiency of the strategy in converting winning trades into cumulative gains relative to losses and is commonly used to assess the robustness of systematic trading rules.

Together, these metrics provide a balanced evaluation of signal quality, risk characteristics, and economic viability across all tested currency pairs.

4 Results

This section reports the empirical performance of the proposed **WMA–IWMA crossover** and compares it against two benchmarks: **WMA–SMA** and **Price–WMA**. Results are reported across a cross-section of large-cap U.S. equities using trade-level metrics, namely hit ratio, profit factor, and Sharpe ratio.

4.1 WMA versus IWMA

Table 1: Performance of the WMA–IWMA crossover strategy

Asset	Hit Ratio	Profit Factor	Sharpe
Apple	32.65%	0.61	-7.66
Microsoft	28.20%	0.90	-1.19
Nvidia	31.70%	0.72	-4.83
Goldman Sachs	29.62%	0.70	-4.90
Morgan Stanley	43.18%	0.87	-1.94
JP Morgan	46.34%	1.65	2.64
Disney	45.65%	1.95	3.18

4.2 WMA versus SMA

Table 2: Performance of the WMA–SMA crossover strategy

Asset	Hit Ratio	Profit Factor	Sharpe
Apple	31.25%	0.59	-3.29
Microsoft	28.94%	0.90	-0.49
Nvidia	30.00%	0.59	-3.17
Goldman Sachs	28.30%	0.69	-2.09
Morgan Stanley	41.86%	0.87	-0.84
JP Morgan	47.50%	1.66	2.72
Disney	44.44%	1.88	3.02

4.3 Price versus WMA

Table 3: Performance of the Price–WMA crossover strategy

Asset	Hit Ratio	Profit Factor	Sharpe
Apple	22.22%	0.89	-0.56
Microsoft	21.97%	1.02	0.92
Nvidia	25.62%	1.17	0.72
Goldman Sachs	19.53%	0.89	-0.47
Morgan Stanley	22.56%	1.01	0.05
JP Morgan	21.93%	0.92	-0.36
Disney	19.02%	1.13	0.42

4.4 Discussion

Several observations emerge from these results. First, the **WMA–IWMA** strategy consistently dominates the **Price–WMA** benchmark across most assets in terms of hit ratio, indicating superior directional timing when comparing symmetric temporal filters rather than raw price deviations. This suggests that relative temporal dominance within a fixed information set provides a more stable decision signal than price-to-trend comparisons.

Second, when contrasted with the **WMA–SMA** benchmark, the WMA–IWMA formulation delivers comparable or superior performance for the majority of assets, despite both strategies sharing the same lookback horizon. This indicates that the inverse weighting structure contributes meaningful information beyond simple equal-weight smoothing, supporting the hypothesis that opposing temporal emphases capture regime transitions more effectively.

Finally, assets such as JP Morgan and Disney exhibit materially higher profit factors and positive Sharpe ratios under the WMA–IWMA framework, highlighting that the proposed crossover mechanism can translate temporal dominance into economically viable performance in favorable market conditions. Importantly, this improvement is achieved without introducing additional parameters or asymmetric horizons, reinforcing the claim that the WMA–IWMA crossover reduces structural subjectivity while preserving interpretability.

Overall, these results support the view that symmetric, same-horizon moving average crossovers provide a more principled alternative to traditional crossover constructions, particularly when the objective is to isolate temporal effects rather than arbitrary time-scale differences.

5 Conclusion

This study proposes a structurally symmetric alternative to traditional moving average crossover strategies by replacing heterogeneous lookback horizons with opposing temporal weightings applied to a fixed information set. By constructing crossover signals from a Weighted Moving Average and its inverse counterpart over the same lookback window, the approach aims to remove one layer of subjectivity inherent in conventional fast–slow moving average designs.

The empirical results presented in this paper are necessarily limited in scope. They are based on approximately 11,000 hourly observations per asset and a restricted universe of equities. As such, the reported performance metrics should not be interpreted as definitive evidence of general profitability or universal robustness. Rather, they serve as an initial validation of the proposed signal construction under a controlled experimental setting.

Importantly, the primary objective of this work is methodological rather than prescriptive. The contribution of the paper lies in demonstrating that moving average crossovers can be reformulated in a way that preserves interpretability while reducing arbitrary design choices related to horizon selection and filter asymmetry. The observed performance variations across assets further suggest that temporal dominance signals may be regime- and instrument-dependent, reinforcing the need for broader cross-asset and cross-frequency evaluation.

Future research should extend this framework to longer samples, alternative asset classes, and different sampling frequencies, as well as investigate the interaction between temporal dominance signals and risk management overlays such as volatility targeting or stop mechanisms. These extensions would help determine whether the proposed formulation offers durable advantages beyond the limited empirical setting considered here.