Transactions on Quantitative Finance and Beyond

www.tqfb.reapress.com

Trans. Quant. Fin. Bey. Vol. 3, No. 2 (2026) 79–99.

Paper Type: Original Article

Performance-Driven Causal Signal Engineering for Financial Markets under Non-Stationarity

Lucas A. Souza*

Independent Researcher, Divinopolis, MG 35500-173, Brazil; lasouza@if.usp.br.

Citation:

Received: 03 May 2025 *Revised: 14 June 2025* *Accepted: 24 December 2025*	Souza, L. A. (2026). Performance-driven causal signal engineering for financial markets under non-stationarity. *Transactions on Quantitative Finance and Beyond, 3*(2), 79-99.

Abstract

We introduce a performance-driven framework for constructing strictly causal forward-oriented observables in strongly non-stationary time series. The method combines a robustly normalized composite of heterogeneous indicators with a causally computed derivative component, yielding a local phase-leading effect that is amplified near regime transitions while remaining fully causal. A hysteresis-based decision functional maps the observable into discrete system states, with execution delayed by one step to preserve strict temporal ordering. Adaptation is achieved through a walk-forward scheme, in which model parameters are selected using rolling train–validation windows and subsequently applied out-of-sample. In this setting, the validation segment acts as an internal performance screen rather than as a statistical validation set, and no claims of generalization are inferred from it alone. The framework is evaluated on high-frequency financial time series as an experimentally accessible realization of a non-stationary complex system. Under a controlled zero-cost setting, the resulting dynamics exhibit a pronounced risk-reshaping effect, characterized by smoother trajectories and reduced drawdowns relative to direct exposure, and should be interpreted as an upper bound on achievable performance. These results illustrate how causal signal engineering can generate anticipatory structure in non-stationary systems without relying on non-causal information, explicit horizon labeling, or high-capacity predictive models.

Keywords: Non-stationary systems, Causal observables, Financial market dynamics, Regime transitions, Walk-forward selection, Decision functionals.

1 | Introduction

Financial markets are paradigmatic examples of non-stationary complex systems, characterized by recurrent regime changes in volatility, trend persistence, and liquidity. From the perspective of statistical physics, these regimes can be viewed as metastable phases whose lifetimes and transitions carry economically relevant information. Early empirical work formalized this view by modeling bull and bear markets as persistent states with duration-dependent dynamics, showing that the probability of regime termination depends on its age rather than being memoryless [1].

Corresponding Author: lasouza@if.usp.br

https://doi.org/10.22105/tqfb.v3i2.65

While regime classification is valuable for ex-post analysis and risk assessment, it does not resolve the ex-anteproblem faced by trading systems: constructing decision rules that operate online, using only informationavailable at execution. This difficulty is amplified at intraday horizons, where the data generating process isstrongly non-stationary and market observables are endogenously coupled through feedback, reflexivity, and microstructure effects. In such conditions, price dynamics are often dominated by self excitation and internal interactions instead of by exogenous information flows, a behavior consistent with markets operating near a critical regime [2].

These features pose fundamental challenges to standard econometric approaches. In non-stationary settings, classical linear models and Granger-causality-based methods may produce unstable or misleading conclusions, as demonstrated in [3], [4]. More recent work has reframed non-stationarity as a potential source of identifiability for causal discovery [5]. However, such methods are primarily designed for structural inference and hypothesis

testing, and are not tailored to the construction of operational, real-time decision functionals. In parallel, machine-learning-based approaches have been increasingly applied to high-frequency financial data. Deep architectures with explicit causal constraints, such as dilated causal convolutions, have shown strong performance in volatility forecasting [6], while reinforcement learning and deep hedging frameworks aim to learn optimal trading or hedging policies directly from data [7], [8]. Despite their expressive power, these methods typically rely on high-capacity models and large training sets, making them sensitive to regime shifts and prone to overfitting in non-stationary environments [9]. Their interpretability and robustness under persistent structural change therefore remain open questions.

An alternative line of research emphasizes explicit causality and regime awareness in trading systems. Approaches based on optimal causal paths [10] or on regime-dependent portfolio allocation [11] highlight the importance of respecting temporal ordering and structural heterogeneity. These methods, however, generally rely on explicit regime identification or latent-state inference, which introduces additional modeling assumptions and estimation error.

In this work, we take an operator-design perspective. We construct a forward-oriented observable from past data only, and embed it in a walk-forward selection protocol so that execution is strictly out-of-sample. Building on a recently proposed causal signal framework [12], we consider a composite observable obtained from heterogeneous technical components, robustly normalized and causally smoothed. While the original formulation demonstrated that such observables can encode precursor information about regime transitions, its performance was found to be strongly regime- and scale-dependent.

Here, we extend this construction by embedding the causal observable within a performance-driven, walkforward selection framework. Model hyperparameters are selected exclusively from past data using rolling train–validation windows and then applied out-of-sample, preserving strict temporal ordering between information, signal construction, and execution. The validation block is interpreted as an internal performance control rather than as a statistical validation set, and is used to rank candidate configurations locally in time rather than to certify generalization. This procedure allows the decision functional to adapt to evolving market conditions without violating causality.

An empirical evaluation on high-frequency cryptocurrency markets illustrates how optimization-aware causal observables can systematically reshape risk and return characteristics, producing smoother equity dynamics and improved drawdown control. More broadly, the framework provides a bridge between causal signal engineering and performance-based adaptation in non-stationary complex systems.

The remainder of the paper is organized as follows. Section 2 introduces the causal signal construction and decision functional. Section 2.5 connects the forward-oriented observable to horizon-based trend objectives. Section 2.8 describes the walk-forward optimization protocol. Section 3 presents empirical results, and Section 4 concludes.

2 | Methodology

The proposed framework defines a fully causal decision functional designed to operate in strongly non-stationary environments with frequent regime changes. We construct a forward-oriented observable by transforming heterogeneous technical indicators into a single composite signal, whose parameters are selected directly through economic performance.

The methodology follows a modular pipeline (see *Fig. 1*). Raw prices are first mapped into standard indicators and processed through a robust causal normalization that removes scale effects and mitigates outliers. The normalized components are then aggregated and augmented by a derivative-based term that introduces controlled anticipatory structure without violating causality. The resulting signal is converted into trading actions via a hysteresis-based decision rule and evaluated under a walk-forward, strictly out-of-sample protocol. All steps are causal; out-of-sample blocks are concatenated to form the global equity.

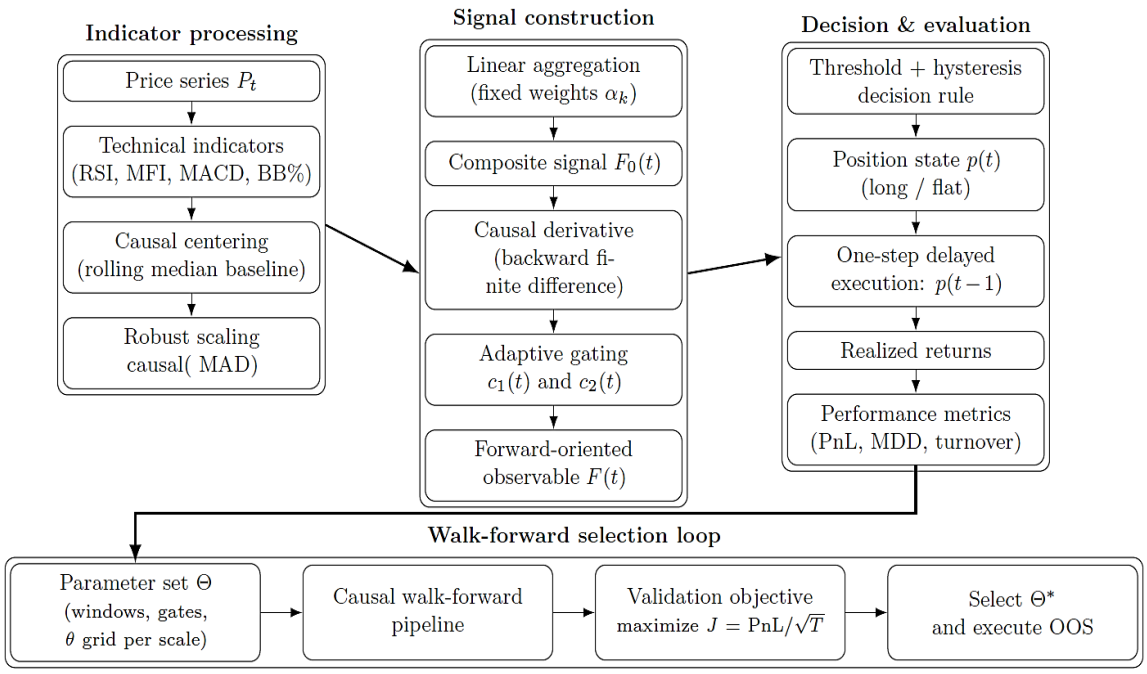

Fig. 1. Pipeline overview.

Prices are mapped to indicators, causally normalized, aggregated into $F_0(t)$, and transformed into a forward-oriented observable $F(t)$ via a gated derivative term. A hysteresis rule produces the binary exposure state, evaluated under a walk-forward protocol.

2.1 | Causal Robust Normalization

All technical indicators are first mapped to a common dimensionless scale through a two-step causal transformation based on rolling medians and Median Absolute Deviation (MAD). Let wnorm denote the normalization window.

It is worth noting that the indicators selected for this work differ in scale: RSI, MFI, and BB% operate within a fixed range of [0, 100], facilitating a uniform interpretation. By contrast, the MACD has no predefined bounds and scale proportionally to both the asset price and its volatility. See *Appendix A* for details. For each indicator $I_t^{(k)}$ and for $t \geq w_{norm}$, define the causal window

$$\mathcal{W}_t(\omega_{norm}) = \{\tau : t - w_{norm} \leq \tau < t\}.$$

containing the wnorm observations strictly prior to t. The causal rolling median (baseline) is

$$m_t^{(k)} = \text{median}\left\{I_\tau^{(k)} : \tau \in \mathcal{W}_t(\omega_{\text{norm}})\right\}.$$

and the centered series is

$$\tilde{I}_t^{(k)} = I_t^{(k)} - m_t^{(k)}.$$

A robust local scale is obtained via the causal MAD on the centered series:

$$s_t^{(k)} = \text{median}\left\{\left|\tilde{I}_\tau^{(k)}\right| : \tau \in \mathcal{W}_t(\omega_{\text{norm}})\right\} + \varepsilon$$

where $\varepsilon > 0$ is a small constant preventing division by zero. Both $m_t^{(k)}$ and $s_t^{(k)}$ depend only on past data ($\tau < t$), ensuring strict causality. The normalized indicator is then

$$Z_t^{(k)} = \frac{I_t^{(k)} - m_t^{(k)}}{s_t^{(k)}}. \quad (1)$$

This transformation yields a scale-free and outlier-robust representation of each indicator. All subsequent steps operate exclusively on the normalized series $\{Z_t^{(k)}\}$, ensuring that differences in scale or volatility do not dominate the composite signal.

2.2 | Composite Observable

Given the normalized indicators $\{Z_t^{(k)}\}$, we define a composite observable as a convex linear combination:

$$F_0(t) = \sum_{k=1}^{K} \alpha_k Z_t^{(k)}, \quad (2)$$

where $\alpha_k \geq 0$ and $\sum_{k=1}^{K} \alpha_k = 1$. This construction preserves the dimensionless nature of the normalized indicators and yields a composite signal interpretable as a consensus Z-score.

In this study, the aggregation weights are fixed and chosen uniformly, $\alpha_k = 1/K$. This design choice is deliberate. By holding the weights constant, we isolate the effect of the causal normalization, derivative-based enhancement, and regime gating mechanisms, reducing the risk of overfitting and improving interpretability. Allowing the weights to vary or be optimized is a natural extension of the framework and is left for future work. We employ four market indicators[1]: $Z_t^{(\text{RSI})}, Z_t^{(\text{MFI})}, Z_t^{(\text{BB\%})}, Z_t^{(\text{MACD})}$ whose formal definitions are provided in Appendix A. Each component captures a distinct aspect of market behavior: RSI measures momentum extremes, MFI incorporates volume-weighted pressure, BB% encodes volatility-adjusted price location, and MACD reflects trend acceleration. Their aggregation into $F_0(t)$ produces a multi-dimensional consensus signal. Coherent deviations across indicators reinforce the composite magnitude, while idiosyncratic noise from individual indicators is attenuated. This structure provides a stable foundation for the forward-oriented transformation introduced in the next subsection.

2.3 | Indicator Analysis

The causal robust normalization introduced in Section 2.1 maps each indicator into a dimensionless representation that is stable under non-stationarity and suitable for linear aggregation. *Fig. 2* illustrates this transformation for the four indicators used in this study (RSI, MFI, MACD, and BB%) on EURUSDT and BTCUSDT, using a 600-minute excerpt of one-minute data. The normalization window is set to wnorm = 5000 minutes, balancing statistical robustness against adaptation to evolving market conditions. The top row shows the centered indicators, obtained by subtracting a causal rolling median. This step removes slow-

[1] Throughout this work, "MACD" refers to the MACD histogram (i.e., the difference between the MACD line and its signal line), and "BB%" denotes the Bollinger %B indicator, also commonly referred to as %b in the technical analysis literature.

moving baselines and aligns all series around zero. In raw units, however, the indicators operate on markedly different numerical scales: RSI, MFI, and BB% are bounded by construction, whereas the

MACD is scale-dependent, as it is defined as a difference of moving averages of the underlying price. For EURUSDT, prices are of order unity, and the associated moving averages therefore remain close to one. Consequently, the MACD signal naturally lies at a much smaller scale, of order $10-4$, reflecting relative price variations rather than absolute levels. The middle row reports robust local scale estimates computed via the causal MAD. Since the MAD is not scale-invariant, the smaller values observed for the MACD reflect its lower numerical scale rather than reduced dynamical variability. In this example, the BB% exhibits the largest dispersion, followed by the MFI and the RSI.

After normalization (bottom row), all indicators fluctuate on comparable, dimensionless scales. Expressed as deviations from a local median in units of robust local variability, each component contributes symmetrically to the composite observable $F_0(t)$, preventing dominance by any single indicator and ensuring meaningful aggregation under non-stationarity.

2.4 | Forward-Looking Operator

To incorporate forward-oriented structure while preserving strict causality, we define a derivative-enhanced observable $F(t)$ from the composite signal $F_0(t)$. Let ndiff denote the lookback window used to compute causal finite differences, and let wMA be a short smoothing window applied to the resulting difference series. The discrete causal derivative is defined as

$$\partial_t^{(n_{difr})} F_0(t) = \frac{F_0(t) - F_0(t - n_{diff})}{n_{diff}},$$

where $n_{diff} \geq 1$ is the window length instead of the order of differentiation. Unlike the unit-lag differencing operator employed in ARIMA models (see Wilson [13] and Hamilton [14]), which is primarily designed to

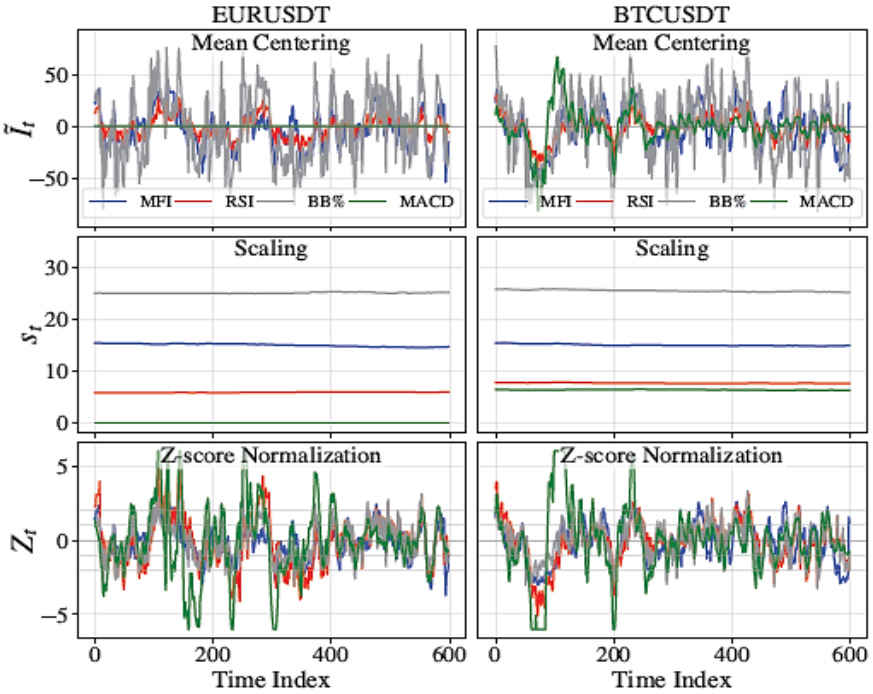

Fig. 2. Causal robust normalization applied to four technical indicators (MFI, RSI, BB%, and MACD) for EURUSDT (left) and BTCUSDT (right).

Top row: centered indicators, $\tilde{I}_t^{(k)} = I_t^{(k)} - m_t^{(k)}$, where $m_t^{(k)}$ denotes the causal rolling median computed over a 5000-minute window. Middle row: robust local scale estimates $s_t^{(k)}$, obtained as the rolling MAD of the

centered series. Bottom row: normalized indicators, $Z_t^{(k)} = \tilde{I}_t^{(k)}/s_t^{(k)}$ (dimensionless), mapped to a common dimensionless scale, enabling direct comparison andaggregation across indicators.

remove stochastic trends and enforce stationarity, the present operator uses a finite backward difference overan extended horizon to extract slope information that is informative about near-future signal evolution. The resulting quantity approximates the local rate of change of $F_0(t)$ based exclusively on past observations. To reduce high-frequency noise, the derivative series is further smoothed via a causal moving average MA of length ω_{MA},

$$\widetilde{\partial_t F_0}(t) = MA_{w_{MA}}\left[\partial_t^{(n_{dirr})} F_0\right],$$

The forward-oriented observable is then defined as

$$F(t) = c_1(t)F_0(t) + c_2(t)\,\widetilde{\partial_t F_0}(t), \tag{3}$$

with state-dependent mixing coefficients

$$c_1(t) = \tanh(|\lambda_1 F_0(t)|)\,; c_2(t) = A[1 - \tanh(|\lambda_2 F_0(t)|)],$$

where $\lambda_1, \lambda_2 > 0$ control sensitivity to the magnitude of the composite signal and $A > 0$ sets the overall scale of the derivative contribution. In strongly persistent regimes, where $|F_0(t)|$ is large, the coefficient $c_2(t)$ is suppressed and $F(t)$ behaves similarly to the level signal $F_0(t)$. Conversely, near transition regions ($F_0(t) \approx 0$), the derivative term is emphasized, yielding a causal approximation to a small forward shift of the composite signal. The derivativeinduced phase advance is illustrated in *Appendix B*. The operator does not forecast future values; it mixes level and slope information in real time. The link to horizon-based trend objectives is discussed in Section 2.5.

2.5 | Bridging to Horizon Based Trend Objectives

The forward-oriented operator $F(t)$ is designed to causally anticipate changes in the trajectory of the composite indicator $F_0(t)$. This objective is closely related to the notion of trend persistence and regime duration that underpins a large body of empirical work in financial economics. Seminal contributions by Pagan and Sossounov [15] and Lunde and Timmermann [1] formalize market trends as bull and bear phases and study their statistical properties ex post, including the duration dependence of regime termination.

Our approach shifts the focus from ex-post classification to ex-ante anticipation. We construct a causal signal that responds early to structural changes in market dynamics. Empirically, the non-causal advance $F_0(t+6)$ exhibits strong alignment with a forward-looking price-based trend metric, suggesting that the composite indicator contains information relevant to future regime transitions.

To formalize the trend target used in the literature, consider a price series $\{P_t\}$ and define the horizon-based future and past price averages as

$$\mu_H^{fut}(t) = \frac{1}{H}\sum_{i=1}^{H} P_{t+i}, \mu_H^{past}(t) = \frac{1}{H}\sum_{i=0}^{H-1} P_{t-i}.$$

The corresponding horizon-based percentage trend indicator is then given by

$$\Delta^{(H)}(t) = \frac{\mu_H^{fut}(t) - \mu_H^{past}(t)}{\mu_H^{past}(t)}, \tag{4}$$

which contrasts the average future price over the next H periods with the average past price over the preceding H periods. This construction, or close variants of it, is widely used as a benchmark for identifying local trend direction and strength.

Fig. 3 illustrates the empirical relationship between the contemporaneous composite signal $F_0(t)$, its non-causal advance $F_0(t+6)$, and the horizon trend indicator $\Delta^{(10)}(t)$. Despite being derived from distinct sources,

$F_0(t+6)$ and $\Delta^{(10)}(t)$ exhibit similar zero-crossing behavior, suggesting that advancing F_0 by a small number of steps captures information relevant for horizon-based trend objectives. Naturally, $F_0(t+6)$ is not available in real time and cannot be used operationally.

The two lower panels provide an illustrative mapping from these signals to directional regimes, shown side by side on a shared time axis. For simplicity, buy (green) and sell (red) regions are defined by sign changes of the respective signals, i.e., crossings of a zero reference level. This visualization is intended solely to highlight the similarity in regime timing induced by the two constructions and to serve as a conceptual bridge to the explicit decision rules introduced in the next section.

The derivative component in the forward-oriented operator $F_{(t)}$ provides a causal mechanism to approximate this anticipatory behavior. This approximation implicitly assumes that the composite signal $F_0(t)$ admits a locally smooth evolution, e.g., is locally Lipschitz-continuous over the relevant time scale, so that a first-order expansion provides a meaningful description of its short-horizon dynamics. Since $\partial_t F_0(t)$ captures the local rate of change of the composite signal, a first-order expansion yields

$$F_0(t+\delta) \approx F_0(t) + \delta\, \partial_t F_0(t). \tag{5}$$

showing that a suitably scaled derivative term can emulate a small forward shift. In our construction, the adaptive gating coefficient $c_2(t)$ effectively implements a regime-dependent δ, amplifying the derivative contribution near transitions and suppressing it during strong trending phases.

Although a fixed multi-step advance (e.g., six periods) is impossible under strict causality, the derivativeenhanced signal $F_{(t)}$ acts as a causal surrogate for limited forward anticipation. This provides a direct conceptual link between the proposed signal construction and horizon-based trend targets widely used in the literature. The economic relevance of this surrogate is ultimately assessed through the PnL-driven walk-forward optimization described in Section 2.8.

2.6 | Decision Rule and Performance Measurement

The forward-oriented observable $F_{(t)}$ serves as the primary input to a binary decision rule. Let $s_t = F(t)$ denote the signal value at time t. A position state $p_t \in \{0, 1\}$ is maintained, where $p_t = 1$ corresponds to a long position and $p_t = 0$ to a flat (cash) position.

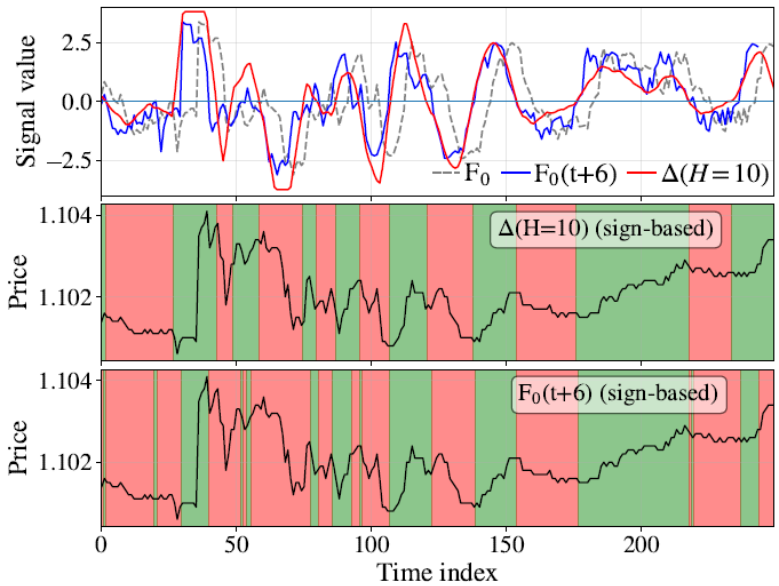

Fig. 3. Relationship between the composite indicator $F_0(t)$ (gray), its non-causal advance $F_0(t+6)$ (blue), and the horizon-based trend indicator $\Delta^{(10)}(t)$ (red).

The advanced signal exhibits alignment with the horizon trend, particularly around zero-crossings associated with local regime transitions. Lower panels illustrate the corresponding sign-based regime markings used for visualization.

To mitigate noise-induced whipsaws, state transitions are governed by a hysteresis mechanism with threshold $\theta > 0$:

$$p_t = \begin{cases} 1, & \text{if } p_{t-1} = 0 \text{ and } s_t > +\theta, \\ 0, & \text{if } p_{t-1} = 1 \text{ and } s_t < -\theta, \\ p_{t-1}, & \text{otherwise.} \end{cases} \quad (6)$$

This design ensures that a switch from flat to long occurs only when the signal surpasses a positive threshold, while a switch from long to flat requires the signal to drop below a negative threshold. The neutral zone between $-\theta$ and $+\theta$ prevents excessive trading when the signal oscillates near zero.

For a strictly causal implementation, the position state used to compute returns at time t is the previous state p_t-1. Let $r_t = (P_t - P_{t-1})/P_{t-1}$ be the simple return of the underlying asset. The strategy return at time t is then:

$$R_t = p_{t-1} \cdot r_t. \quad (7)$$

The cumulative equity curve is obtained by compounding:

$$V_t = V_0 \prod_{\tau=1}^{t} (1 + R_\tau). \quad (8)$$

With $V_0 = 1$ as the initial capital. Trading activity is quantified by the total number of position changes over a given horizon T:

$$\text{Turnover}(T) = \sum_{t=1}^{T} |p_t - p_{t-1}|. \quad (9)$$

2.7 | Signal Scaling and Gating Effects

Before introducing the parameter optimization framework, we examine how the scale of the constructed signal $F(t)$ depends on the gating parameters. This analysis is essential because trading decisions are based on thresholding $F(t)$: the frequency of trades and effective exposure depend on how often $|F(t)|$ exceeds a decision threshold θ. If the amplitude of $F(t)$ varies substantially with (λ_1, λ_2, A), jointly optimizing these parameters together with θ becomes ill-posed, as distinct configurations may generate similar trading behavior through simple rescaling. Characterizing the dependence of the signal magnitude on the gating parameters is therefore necessary to decouple signal construction from decision calibration.

To isolate gating effects, we fix equal indicator weights $\alpha_k = 1/K$ and compute $F(t)$ over the last 100,000 observations (approximately 70 days of one-minute data) for EURUSDT and BTCUSDT. Price data consist of one-minute OHLCV series obtained via the Binance public Spot API [16]. The sample spans January 2022 to December 2025. Missing candles are handled by forward-filling prices and setting volume to zero, ensuring a regular time grid without introducing look-ahead bias.

The derivative parameters are fixed at $n_{diff} = 2$ and $\omega_{ma} = 2$. For each experiment, one gating parameter is varied while the others are held constant, and the signal scale is summarized by median ($|F(t)|$).

Fig. 4 shows a smooth and monotonic dependence of the signal magnitude on the gating parameters. Increasing λ_1 raises median ($|F(t)|$) by strengthening the level component, while increasing λ_2 suppresses the

derivative contribution and reduces the overall amplitude. The parameter A scales the derivative channel almost linearly. These patterns are consistent across both assets, indicating that the robust normalization of $F_0(t)$ stabilizes the baseline scale, whereas (λ_1, λ_2, A) primarily control the effective amplitude of $F(t)$.

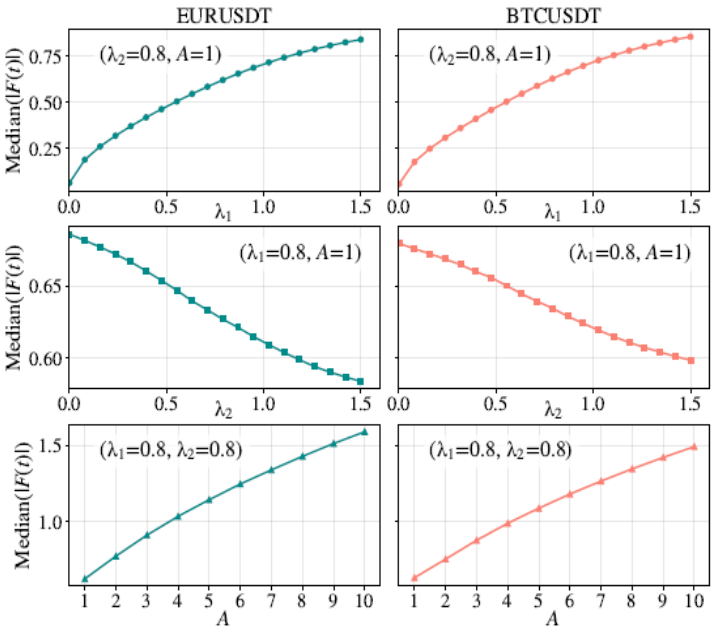

Fig. 4. Median magnitude of the constructed signal, median(|F(t)|), as a function of the gating parameters (λ_1, λ_2,A), computed over the last 100,000 observations for EURUSDT and BTCUSDT.

The derivative parameters are fixed at $n_{diff} = 2$ and $\omega_{ma} = 2$. The smooth and monotonic dependence indicates that the gating parameters predominantly determine the scale of $F(t)$.

Additional experiments (not shown) indicate that moderate variations around $n_{diff} = 2$ and $\omega_{ma} = 2$ do not materially affect the structure or scale of $F(t)$ relative to the dominant influence of the gating parameters. This separation motivates treating the decision threshold θ independently of signal-shaping parameters. The next section introduces the walk-forward optimization framework and describes how θ is selected in a manner consistent with the observed scaling behavior of $F(t)$.

2.8 | Parameter Optimization Framework

Given the forward-oriented signal $F(t)$ and the decision rule defined in Section 2.6, the trading strategy depends on a finite set of parameters that govern both signal construction and temporal adaptation. For clarity, we partition these parameters into two groups.

I. Signal construction parameters, which determine the shape and scale of the forward-oriented observable $F(t)$.

II. Walk-forward window parameters, which control how the model adapts to non-stationarity over time.

The signal construction parameters are collected in

$$\Theta_{sig} = \{n_{diff}, w_{ma}, \lambda_1, \lambda_2, A\}. \quad (10)$$

where n_{diff} and ω_{ma} define the causal derivative operator, and (λ_1, λ_2, A) govern the adaptive gating mechanism introduced in *Eq. (3)*. Threshold selection is treated separately (See Section 3.2) to avoid confounding signal scaling with execution filtering. The walk-forward protocol introduces two additional window parameters,

$$\Theta_{win} = \{\omega_{flt}, \rho\}, \quad (11)$$

where wfit denotes the length of the training window and ρ determines the validation horizon via $\omega_{val} = \lfloor \omega_{fit}/\rho \rfloor$. The out of sample execution window is set equal to the validation window (wexec = wval), ensuring a consistent evaluation horizon across epochs, as detailed in Section 2.9.

At each walk-forward epoch, the parameter set optimized is therefore

$$\Theta = \Theta_{sig} \cup \Theta_{win}.$$

The optimization is performed via a discrete grid search over predefined candidate values. The specific grids employed in the empirical analysis are reported in *Table 1*.

Objective function. For a given parameter configuration Θ, the pipeline generates a sequence of strategy returns $\{R_t(\Theta)\}_{t=1}^{T_{eval}}$. The terminal profit-and-loss over the evaluation window is

$$PnL(\Theta) = \prod_{t=1}^{T_{eval}} \big(1 + R_t(\Theta)\big) - 1. \quad (12)$$

Assuming unit initial capital. Here, Teval denotes the number of observations in the evaluation block on which the objective is computed. To enable comparisons across evaluation windows of different lengths, we maximize the time-normalized objective

$$J(\Theta) = \frac{PnL(\Theta)}{\sqrt{T_{eval}}}, \quad (13)$$

which penalizes configurations that achieve high PnL over very short horizons. This objective should be interpreted as a volatility penalized growth proxy, not a risk-adjusted estimator. When multiple parameter sets attain similar objective values, turnover measured as the total number of position changes over the evaluation window is used as a tie-breaking criterion to favor more stable trading behavior.

It is important to emphasize that this optimization should be interpreted as an internal, locally defined performance screen rather than as a mechanism for establishing statistically validated generalization. Because a finite grid of configurations is evaluated on each validation block, some degree of performance overfitting to local conditions is unavoidable. The role of the optimization is therefore not to identify a globally optimal or universally generalizing parameter set, but to maintain an internally consistent ranking of a small family of candidate configurations as market conditions evolve.

2.9 | Walk-Forward Protocol

Because the objective J(Θ) is evaluated through a fully causal backtest, gradient information is unavailable and parameter selection is performed via black-box optimization. To respect nonstationarity while preserving strict causality, we adopt a walk-forward protocol in which estimation, model selection, and execution are explicitly separated in time. In this protocol, the "validation" block serves purely as an operational performance filter that ranks candidate configurations based on recent behavior; it is not interpreted as an independent hold-out set providing formal guarantees about future generalization.

At each epoch, the data are partitioned into three contiguous past-to-present blocks (see *Fig. 5*):

I. A training block of length wfit, used exclusively for signal construction and state estimation.

II. A validation block of length wval, used solely for internal model selection and short-horizon monitoring, it is not used to report performance and does not constitute an out-of-sample test.

III. A subsequent test/forecast block of length wexec, on which the selected strategy is executed out-of-sample.

At an epoch boundary t, these blocks are defined as

$$\text{train: } [t - \omega_{ftt} - \omega_{val}, t - w_{val}). \quad (14)$$

$$\text{val: } [t - \omega_{val}, t). \tag{15}$$

$$\text{test: } [t, t + \omega_{exec}). \tag{16}$$

Crucially, parameter selection never uses information from the test block. The parameter vector Θ is chosen by maximizing the validation score computed exclusively on the validation block:

$$\Theta^*(t) = \arg\max_{\Theta \in \mathcal{G}} J_{val}(\Theta; \omega_{val}), \tag{17}$$

where $\mathcal{G}$ denotes a finite grid of candidate configurations (see *Table 1*). The selected configuration $\Theta*(t)$ is then frozen and applied without modification to the subsequent test block, producing the out-of-sample position sequence used to construct the global equity curve. A schematic overview of the rolling trainvalidation execution schedule, including the enforced horizon constraint $\omega_{exec} = \omega_{val}$, is shown in *Fig. 5*.

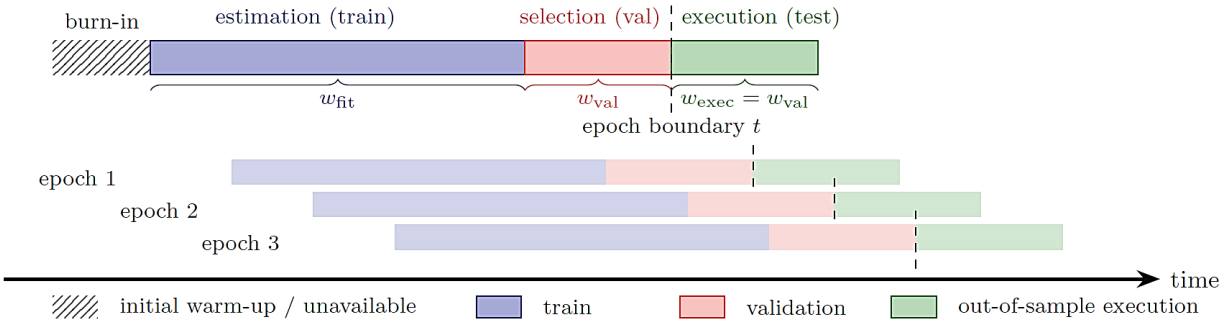

Fig. 5. Rolling window schedule used for causal walk-forward model selection.

At each epoch boundary t, parameters are estimated on a past training block, selected using the immediately preceding validation block, and then executed out-of-sample on the next block. We impose $\omega_{exec} = \omega_{val}$ so that the validation horizon matches the operational trading horizon, enabling directly comparable performance metrics.

In addition to optimizing Θ, the protocol allows the window lengths (ω_{fit}, ω_{val}, ω_{exec}) themselves to vary across epochs. Operationally, this is implemented as a discrete search over wfit and a ratio parameter ρ, which determines the validation length via

$$\omega_{val} = \text{round}\left(\frac{\omega_{ftt}}{\rho}\right). \tag{18}$$

We impose the constraint

$$\omega_{exec} = \omega_{val}, \tag{19}$$

which ensures that model selection and out-of-sample execution are evaluated over identical horizons. This constraint serves three purposes; 1) it enforces a clear separation between estimation and execution while maintaining a consistent evaluation scale, 2) it reduces the effective dimensionality of the window search, mitigating the risk of horizon overfitting, and 3) it yields a transparent operational interpretation in which the trading horizon is calibrated on the most recent validation period and immediately deployed on the next unseen block of the same duration. Under this walk-forward scheme, the loop advances in steps of wexec, producing non-overlapping out-of-sample test blocks whose concatenated positions define the global equity.

At this point, it is useful to contrast the present framework with earlier heuristic constructions in the literature. That study can be recovered as a particular and more restrictive instance of the current construction, with simplified preprocessing and empirically chosen, fixed hyperparameters. Related approaches based on derivativeaugmented indicators have been explored in prior work (see, [12] e.g.,), but without the robust normalization and systematic walk-forward calibration introduced here. Aggregation weights were adjusted to yield signals of similar visual amplitude, and key hyperparameters– including the derivative gain, smoothing

horizon, and mixing coefficients–were selected empirically based on qualitative inspection and preliminary performance.

While this heuristic approach proved effective and provided early evidence of the benefits of derivativeaugmented signals, it lacked an explicit normalization scheme, a principled calibration procedure, and a unified parameterization across indicators. The present work addresses these limitations by introducing robust MADbased normalization, a formally defined composite signal scale, and a systematic walk-forward calibration of all hyperparameters. Under this formulation, the framework of [12] emerges as a special case obtained by fixing the normalization to simple centering and selecting constant parameter values, thereby situating the earlier results within a broader and more robust causal signal engineering paradigm.

3|Experimental Results

At intraday horizons, the design and evaluation of trading strategies face substantial challenges due to the dominance of noise and quickly evolving structural changes, a point long emphasized in the empirical market microstructure literature, notably by Christensen et al. [17]. Cryptocurrency markets amplify these difficulties, exhibiting pronounced non-stationarity, frequent regime shifts, and distinctive microstructural dynamics, as documented by recent studies such as Easley et al. [18] and De Blasis and Webb [19].

Consistent with recent algorithmic trading literature that adopts rolling and walk-forward evaluation protocols on high-frequency cryptocurrency data (see, e.g., Stefaniuk and ´Slepaczuk [20] and Arian et al. [9]), we evaluate the proposed framework on two liquid spot cryptocurrency markets, EURUSDT and BTCUSDT, using oneminute sampled data. Specifically, we assess a fully causal decision functional optimized in a walk-forward fashion and evaluated at the one-minute frequency.

3.1|Experimental Setup and Walk-Forward Protocol

All stages of the pipeline indicator processing, normalization, signal construction, parameter selection, and trading execution are strictly causal, relying exclusively on information available up to time t. To isolate the intrinsic behavior of the signal construction and decision mechanism, all results reported in this section are obtained under zero transaction costs. This choice allows us to analyze the raw statistical and dynamical properties of the strategy independently of market-specific frictions, which are addressed in a separate robustness analysis.

Under the walk-forward framework introduced earlier, at each rebalancing time t, model parameters are estimated through a train–validation split relying exclusively on historical information (Section 2.8), after which they are held fixed and applied out-of-sample over a contiguous test window. The global equity curve is constructed by concatenating the corresponding out-of-sample position paths across successive windows. The walk-forward optimization framework described in Section 2.8 was implemented with the following parameter grids:

Table 1. Parameter grids used in the walk-forward selection.

Param.	Description	Grid values
Signal Construction Parameters		
n_{diff}	Derivative lookback window	{2}
w_{ma}	Derivative smoothing window	{2}
λ_1	Gating sensitivity (signal channel)	{0.01, 0.5, 1, 1.5}
λ_2	Gating sensitivity (derivative channel)	{0.01, 0.5, 1, 1.5}
A	Derivative amplitude scale	{0.75, 1, 2}
Walk-Forward Window Parameters		
w_{fit}	Training window length (minutes)	{720, 1440, 2880,7200, 12000}
ρ	Training-validation ratio	{2,3,5,6}

These parameter ranges were selected based on preliminary sensitivity analyses (Section 2.7) to cover a wide but computationally tractable search space. The fixed values for ndiff and ω_{ma} reflect the minimal configuration that provides stable derivative estimates while preserving high-frequency signal dynamics.

3.2 | Choice of the Hysteresis Threshold θ

The hysteresis threshold θ governs regime transitions in the final trading The choice is closely related to the empirical scale of the constructed signal. *Fig. 4* characterizes how the typical magnitude of the forward-oriented observable F(t) varies as a function of the gating parameters, summarized through median (|F(t)|). This analysis identifies a natural signal scale of order unity across assets and parameter configurations.

Operating directly at the level of the median signal magnitude would, however, result in excessive regime switching. Since |F(t)| exceeds its median value by construction approximately half of the time, setting θ near the median would induce frequent threshold crossings and a high turnover trading regime dominated by short-lived fluctuations. From a trading perspective, the hysteresis threshold must therefore be chosen above the typical median scale in order to enforce regime persistence and suppress noise-driven transitions. Accordingly, we consider a discrete set of threshold values

$$\theta \in \{0.6, 0.8, 1.0, 1.4, 1.6\},$$

which spans values below, near, and above the median signal magnitude reported in *Fig. 4*. This range allows us to study the sensitivity of trading behavior to the degree of hysteresis while remaining within a scale-consistent operating regime.

Importantly, the interpretation of *Fig. 4* does not rely on forward-looking information. While the figure is computed over a long sample for descriptive purposes, additional diagnostic checks (not shown) indicate that the empirical distribution of $|F_0(t)|$ and in particular its high percentiles remains remarkably stable over time for both assets considered. This temporal stability implies that the characteristic scale of the signal is a structural property of the normalization and aggregation procedure, rather than an artifact of a specific market phase or sample endpoint.

3.3 | Equity Curves, Trading Activity, and Risk Characteristics

Figure 6 reports the out-of-sample walk-forward performance of the proposed strategy for EURUSDT and BTCUSDT under zero transaction costs. For each asset, the figure compares the cumulative return of the proposed strategy across different values of the threshold parameter θ with a Buy-and-Hold benchmark, and displays the cumulative number of executed trades over time.

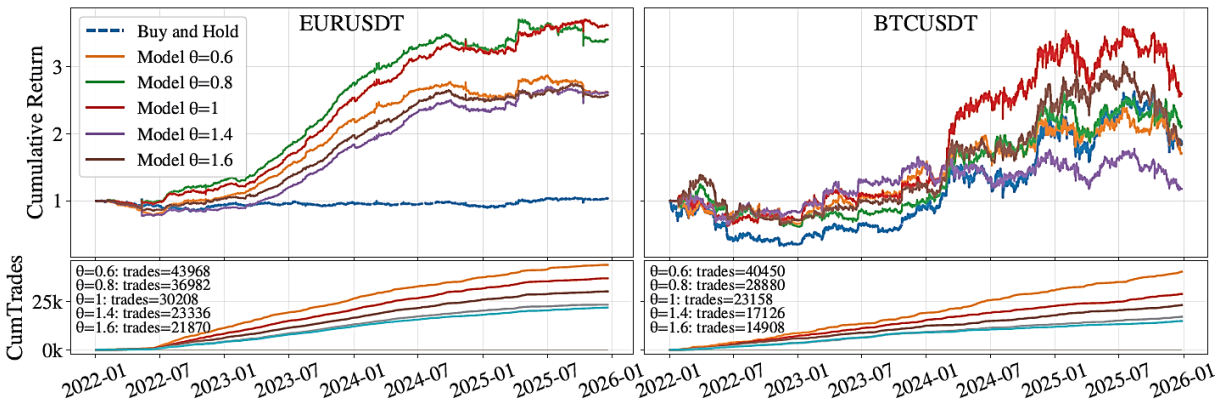

Fig. 6. Walk forward out of sample performance under zero transaction costs.

Strategy versus Buy-and-Hold for different values of θ, and the cumulative number of executed trades. The equity curves reveal distinct but economically meaningful effects across assets. For EURUSDT, the proposed strategy exhibits sustained growth and consistently outperforms Buy-and-Hold across a wide range of θ

values. For BTCUSDT, although Buy-and-Hold attains a higher terminal return for some thresholds, the proposed strategy produces markedly smoother equity curves, with visibly reduced drawdowns and lower path variability.

The lower panels of *Fig. 6* show that trading activity differs across markets, reflecting differences in volatility and regime dynamics rather than changes in signal construction, which is identical across assets. Importantly, variations in trading frequency across θ do not translate into increased instability of the equity process. Even in the more volatile BTCUSDT market, the strategy maintains controlled exposure and avoids the large drawdowns observed under Buy-and-Hold.

To quantify these qualitative observations, *Table 2* reports risk and return metrics for Buyand Hold and for the proposed strategy evaluated at the selected threshold value $\theta^* = 1.0$. This choice is not motivated by a maximization of out-of-sample performance. Rather, $\theta^* = 1.0$ provides a convenient and interpretable reference point at which trading activity, regime persistence, and signal magnitude are all clearly expressed, facilitating the analysis of parameter distributions, trade statistics, and regime behavior.

Importantly, the qualitative conclusions drawn in the subsequent sections do not depend on this specific value. As illustrated in *Fig. 6*, alternative threshold choices lead to predictable changes in trading intensity but preserve the core behavior of the strategy.

Table 2. Out-of-sample trading metrics for Buy & Hold and the proposed strategy under the walk-forward framework.

EURUSDT ($\theta^* = 1.0$)			
Period	**Metric**	**Buy & Hold**	**Proposed**
Jan 2022	Total Return	0.033304	2.620689
	Volatility	0.000272	0.000181
Dec 2025	Downside Volatility	0.000293	0.000266
	Max Drawdown	-0.200671	-0.114912
	Sharpe Ratio	0.000217	0.004859
	Sortino Ratio	0.000201	0.003312
	Calmar Ratio	0.165963	22.805979
	Ulcer Index	9.630334	2.309119
	Time Under Water	0.999107	0.966396
	Total Trades	-	30,208
	Trades per 1 k candles	-	20.3
BTCUSDT ($\theta^* = 1.0$)			
Jan 2022	Total Return	0.849965	1.596073
	Volatility	0.000845	0.000649
Dec 2025	Downside Volatility	0.000680	0.000688
	Max Drawdown	-0.677898	-0.440523
	Sharpe Ratio	0.000908	0.001306
	Sortino Ratio	0.001128	0.001232
	Calmar Ratio	1.253825	3.623134
	Ulcer Index	32.882349	16.633291
	Time Under Water	0.999351	0.983562
	Total Trades	-	23,158
	Trades per 1 k candles	-	15.5

Performance is computed by concatenating out-of-sample test windows across the full sample, with risk, drawdown, and activity metrics reported for $\theta^* = 1.0$. All metrics are computed at the native one-minute frequency. Sharpe and Sortino ratios are reported in non-annualized form. For comparison with annualized

benchmarks, these ratios can be scaled by $\sqrt{365 \times 24 \times 60}$. For standard definitions of the reported performance metrics, see, e.g., Bacon [21].

Taken together, the results indicate that the proposed framework primarily acts as a risk-reshaping mechanism. By combining causal signal construction, adaptive derivative leading, and hysteresis-based execution, the strategy systematically transforms raw market exposure into a smoother equity process sacrificing part of the upside in highly volatile environments while delivering strong absolute and risk-adjusted performance in more stable regimes.

3.4|Trading Regime Duration and Position Holding Structure

Beyond aggregate performance and risk metrics, the temporal structure of trading regimes provides insight into the economic nature and practical viability of the proposed strategy. In particular, the distribution of position holding times reveals whether performance is driven by sustained regime identification or by rapid oscillations more likely associated with microstructure noise.

Fig. 7 reports the empirical distribution of non-flat position durations (measured in one-minute candles) for EURUSDT and BTCUSDT under the selected walk-forward configuration. Each duration corresponds to a contiguous interval during which the strategy maintains a long position before reverting to flat.

Both assets display highly asymmetric duration distributions, characterized by a concentration of short holding times (on the order of tens of minutes) and a long right tail extending to several thousand minutes.

This structure is consistent with frequent local regime fluctuations coexisting with less frequent but persistent directional phases. Despite substantial differences in volatility and market microstructure, the qualitative shape of the duration distributions is remarkably similar across EURUSDT and BTCUSDT. This similarity suggests that the forwardoriented causal signal extracts regime information in a manner that is robust across asset classes, rather than being tuned to idiosyncratic price dynamics.

Table 3 summarizes key descriptive statistics of the position durations. The median holding time is approximately 22–23 minutes for both assets, while the upper quantiles reveal the presence of materially longer regimes. In particular, the 90th percentile lies around 77 minutes, and the maximum durations exceed 10,000 minutes in both markets, indicating that the strategy intermittently locks into persistent regimes.

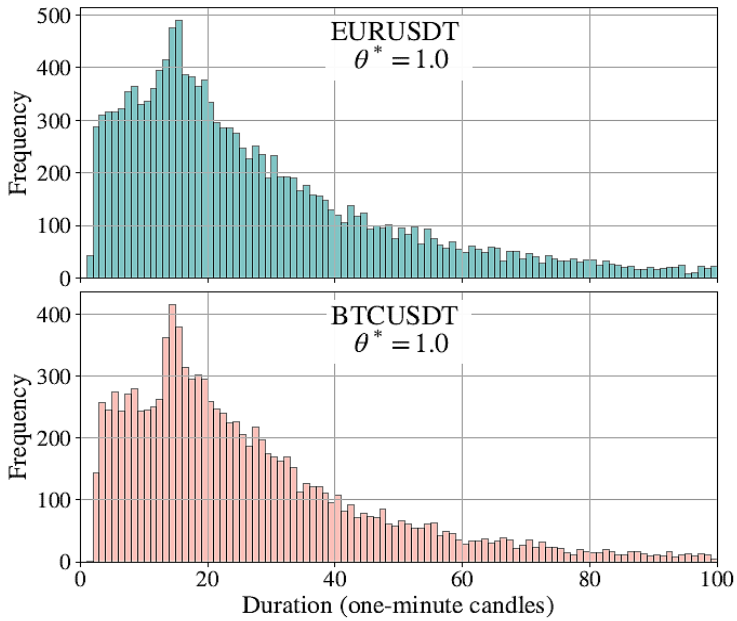

Fig. 7. Empirical distribution of position durations (non-flat segments) under the walkforward optimized strategy for EURUSDT and BTCUSDT.

Both assets exhibit strongly rightskewed distributions, with a large mass of short-lived positions and a non-negligible tail of persistent regimes.

Table 3. S ummary statistics of non-flat position durations (in one-minute candles) over the January 2022 to December 2025 period for θ∗ = 1.0.

Assest	Mean	Median	P25	P75	P90	Max
EURUSDT	52.1	23	13	42	77	11,086 (05-17 Mar/2025)
BTCUSDT	74.6	22	13	40	77	12,508 (17-30 Jan/2023)

From a practical perspective, the prevalence of short-duration positions highlights the importance of transaction costs and execution frictions. Since all results are reported under zero transaction costs, the observed performance should be interpreted as a measure of the signal's informational and regime-detection content, rather than as a directly tradable outcome.

At the same time, the presence of a non-negligible tail of long-duration positions indicates that the strategy is not merely exploiting high-frequency noise. Instead, it intermittently captures persistent regimes, precisely the environments in which transaction costs are less dominant and economic value is more likely to survive realistic execution conditions.

3.5. Summary on Selected Hyperparameters Across Walk-Forward Epochs

To provide a compact view of the tuning behavior induced by the walk-forward selection procedure, summarizes the empirical frequency with which each hyperparameter value is selected as optimal across epochs, separately for assets. Rather than reporting a single best configuration, these histograms reveal which regions of the search space are repeatedly preferred by the out-of-sample validation criterion, hence offering an interpretable diagnostic of hyperparameter stability under non-stationarity.

The top panels report the distribution of the rolling-window lengths used by the walk-forward protocol: the training window length wfit, wpred and the validation partition ratio ϱ. In our experiments, wval = wpred. The bottom panel reports the distribution of the selected values for the signal-construction parameters.

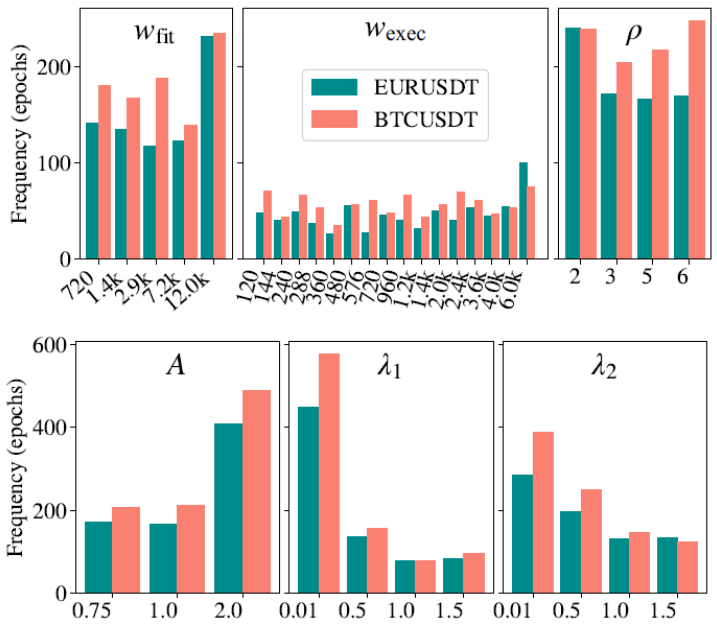

Fig. 8. Empirical selection frequencies of hyperparameters under the walk-forward protocol.

Each bar reports the number of epochs in which a given value is selected as optimal, shown separately for EURUSDT and BTCUSDT.

4 | Conclusion

We introduced and empirically tested a strictly causal operator that transforms heterogeneous indicators into a forward-oriented decision observable in non-stationary markets. The present framework is motivated by limitations observed in earlier fixed-parameter constructions, in which empirical performance was found to depend sensitively on market regime, particularly during extended low-volatility phases. By introducing adaptive window calibration and scale-stable normalization, the approach proposed here is designed to mitigate this sensitivity by dynamically realigning signal dynamics with recent market conditions.

The observed behavior across assets reflects differences in the underlying market structure. In EURUSDT, the strategy generates a stable cumulative return with substantially reduced volatility and drawdowns relative to Buy-and-Hold. In BTCUSDT, characterized by extreme volatility and deep benchmark drawdowns, the strategy reshapes exposure by compressing downside risk and smoothing equity dynamics, at the cost of reduced upside participation. These effects arise from systematic, causal adaptation to recent data rather than from any use of future information. At the same time, because parameters are selected from a finite grid based on past performance, some degree of local overfitting to the walk-forward objective remains possible.

At the same time, the effectiveness of the derivative-based leading mechanism relies on implicit structural properties of the underlying signal. The anticipatory effect is meaningful when the composite observable admits a locally smooth evolution, such that a low-order expansion of the form $F_0(t+\delta) \approx F_0(t) + \delta\, \partial_t F_0(t)$ provides a valid approximation. For signal geometries that deviate strongly from this structure e.g., highly irregular, discontinuous, or noise-dominated dynamics the derivative contribution may carry limited anticipatory information. Delineating the class of signal behaviors for which causal derivative operators remain effective is an important direction for future investigation.

A central limitation of the present study is the assumption of zero transaction costs. Under this idealized setting, trading activity is high, and reported performance should be interpreted as an upper bound. Nevertheless, the observed heavy-tailed distribution of holding times, with a non-negligible population of long regimes, indicates that the strategy intermittently captures persistent directional phases rather than reacting exclusively to short-scale noise.

Future extensions should incorporate transaction costs and market impact directly into the walk-forward objective, potentially using regime duration as an explicit regularization variable. Systematic comparisons with high-capacity horizon-based models, including convolutional and recurrent architectures, would further clarify the trade-offs between causal signal engineering and deep predictive approaches. An additional priority is to replace the relatively broad hyperparameter grids used here by a more structured selection of a small number of parameter families, emphasizing robustness and interpretability over marginal in-sample performance gains; this is expected to reduce the effective degrees of freedom of the framework and hence its susceptibility to overfitting.

Overall, the results demonstrate that carefully designed causal operators can generate effective forwardoriented structure under non-stationarity, offering an interpretable and adaptive alternative to non-causal or purely data-driven trading constructions.

Furthermore, all computations were implemented in Python using standard scientific libraries (NumPy, pandas) and the code used for data processing and backtesting will be made available upon request and released publicly after acceptance.

Funding

The author received no external funding for this research.

Acknowledgments

The author would like to thank the anonymous reviewers and members of the editorial board for their thoughtful comments and constructive insights, which helped to enhance the clarity, structure, and overall quality of this manuscript.

Data Availability

The data used in this study consist of high-frequency financial market datasets obtained from publicly accessible trading platforms. Processed data and implementation details supporting the findings of this research are available from the author upon reasonable request.

References

[1] Lunde, A., & Timmermann, A. (2004). Duration dependence in stock prices. *Journal of business & economic statistics*, *22*(3), 253–273. https://doi.org/10.1198/073500104000000136

[2] Hardiman, S. J., Bercot, N., & Bouchaud, J.-P. (2013). Critical reflexivity in financial markets: A Hawkes process analysis. *The european physical journal b*, *86*(10), 442. https://doi.org/10.1140/epjb/e2013-40107-3

[3] Papana, A.Yrtsou, C., Kugiumtzis, D., & Diks, C. (2016). Detecting causality in non-stationary time series using partial symbolic transfer entropy: Evidence in financial data. *Computational economics*, *47*(3), 341–365. https://doi.org/10.1007/s10614-015-9491-x

[4] Sionkowski, P., Bełdowski, P., Kruszewska, N., Weber, P., Marciniak, B., & Domino, K. (2022). Effect of ion and binding site on the conformation of chosen glycosaminoglycans at the albumin surface. *Entropy*, *24*(6), 811. https://doi.org/10.3390/e24060811

[5] Sadeghi, A., Gopal, A., & Fesanghary, M. (2025). Causal discovery from nonstationary time series. *International journal of data science and analytics*, *19*(1), 33–59. https://doi.org/10.1007/s41060-024-00679-7

[6] Moreno-Pino, F., & Zohren, S. (2024). DeepVol: Volatility forecasting from high-frequency data with dilated causal convolutions. *Quantitative finance*, *24*(8), 1105–1127. https://doi.org/10.1080/14697688.2024.2387222

[7] Jaisson, T. (2022). Deep differentiable reinforcement learning and optimal trading. *Quantitative finance*, *22*, 1429–1443. https://doi.org/10.1080/14697688.2022.2062431

[8] Buhler, H., Gonon, L., Teichmann, J., & Wood, B. (2018). Deep hedging. *Quantitative finance*, *19*, 1271–1291. https://doi.org/10.1080/14697688.2019.1571683

[9] Arian, H., Mobarekeh, D. N., & Seco, L. (2024). Backtest overfitting in the machine learning era: A comparison of out-of-sample testing methods in a synthetic controlled environment. *Knowledge-Based Systems*, *305*, 112477. https://api.semanticscholar.org/CorpusID:281092860

[10] Stübinger, J. (2018). Statistical arbitrage with optimal causal paths on high-frequency data of the S&P 500. *Quantitative finance*, *19*, 921–935. https://doi.org/10.1080/14697688.2018.1537503

[11] Nystrup, P., Madsen, H., & Lindström, E. (2016). Dynamic Portfolio optimization across hidden market regimes. *Quantitative finance*, *18*, 83–95. https://doi.org/10.1080/14697688.2017.1342857

[12] Souza, L. A. (2025). Forward-oriented causal observables for non-stationary financial markets. *arXiv preprint arXiv:2512.24621*. https://doi.org/10.48550/arXiv.2512.24621

[13] Tunnicliffe Wilson, G. (2016). Time series analysis: Forecasting and Control, *Journal of time series analysis*, *37*. https://doi.org/10.1111/jtsa.12194

[14] Hamilton, J. D. (1994). *Time series analysis*. Princeton University Press. https://doi.org/10.2307/j.ctv14jx6sm

[15] Pagan, A. R., & Sossounov, K. A. (2003). A simple framework for analysing bull and bear markets. *Journal of applied econometrics*, 18(1), 23-46. https://doi.org/10.1002/jae.664

[16] Binance. (2026). *Changelog for Binance's API*. https://developers.binance.com/docs/binance-spot-api-docs

[17] Christensen, K., Oomen, R. C. A., & Podolskij, M. (2014). Fact or friction: Jumps at ultra high frequency. *Journal of financial economics*, *114*(3), 576–599. https://doi.org/10.1016/j.jfineco.2014.07.007

[18] Easley, D., Yang, S., & Zhang, Z. (2024). *Microstructure and market dynamics in Crypto Markets*. https://doi.org/10.2139/ssrn.4814346

[19] De Blasis, R., & Webb, A. (2022). Arbitrage, contract design, and market structure in Bitcoin futures markets. *Journal of futures markets*, *42*, 492–524. https://doi.org/10.1002/fut.22305

[20] Stefaniuk, F., & Ślepaczuk, R. (2025). Informer in algorithmic investment strategies on high frequency bitcoin data. *arXiv preprint arXiv:2503.18096*. https://doi.org/10.48550/arXiv.2503.18096

[21] Bacon, C. R. (2021). *Practical risk-adjusted performance measurement*. Wiley. https://www.abebooks.com/Practical-Risk-Adjusted-Performance-Measurement-Bacon-Carl/31029014492/bd

[22] Appel, G. (1985). The moving average convergence-divergence trading method. *Trade republic*. https://openlibrary.org/books/OL9438619M

[23] Murphy, J. J. (1999). *Technical analysis of the financial markets: A comprehensive guide to trading methods and applications*. Penguin Publishing Group. https://books.google.nl/books?id=5zhXEqdr_IcC

[24] Wilder, J. W. (1978). *New concepts in technical trading systems*. Trend research. https://books.google.nl/books?id=WesJAQAAMAAJ

Appendix A

Technical indicators and composite observable

All indicators are computed using exclusively information available up to time t, ensuring strict causality, and follow standard definitions widely used in empirical finance and technical analysis (see Appel [22] and Murphy [23] for details). The indicators $\{Z_t^{(k)}\}$ are summarized in *Table 4*, together with their economic interpretation and defining equations. Each indicator captures a complementary aspect of market dynamics, as briefly summarized below.

Table A1. Technical indicators used to build the composite observable $F_0(t)$.

Indicator	Range	Definition
RSI	[0,100]	$100 \times \left(1 - \frac{1}{1 + G_t/L_t}\right)$
MFI	[0,100]	$100 \times \left(1 - \frac{1}{1 + PMF_t/NMF_t}\right)$
MACD diff.	Scale-dependent	$(\text{EMA}_{\text{fast}} - \text{EMA}_{\text{slow}})_t - \text{Signal}_t$
BB%	[0,100]	$100 \times \left(\frac{P_t - (\mu_t - k\sigma_t)}{(\mu_t + k\sigma_t) - (\mu_t - k\sigma_t)}\right)$

RSI, MFI and BB% are bounded in [0, 100], whereas MACD is scale-dependent. The Relative Strength Index (RSI), following Wilder [24], measures the balance between recent gains and losses and provides a bounded representation of short-term momentum. The Money Flow Index (MFI) extends this concept by incorporating traded volume, thereby capturing buying and selling pressure beyond price movements alone [23]. The MACD difference reflects changes in trend intensity by comparing short- and long-horizon exponential moving averages and their signal line [22]. Finally, the Bollinger Band Percent (BB%) expresses the relative position of the price within a volatility envelope, yielding a dimensionless measure of price extremeness conditional on recent volatility.

Appendix B

Derivative based signal leading: Illustrative examples

The forward-oriented operator introduced in Section 2.4 augments the composite signal with a causal derivative component in order to promote systematic phase advance while preserving strict causality. To illustrate the underlying mechanism in a controlled setting, we consider the demonstrative signal

$$f(t) = \sin(at + A_0 \sin(wt)) + mt - 1. \quad \textbf{(A1)}$$

with a=0.5, A_0=2, ω=1, and m=0.1. This construction combines oscillatory behavior with slow modulation and a linear trend, reproducing qualitative features typical of non-stationary signals. The causal derivative is approximated via backward finite differences,

$$\frac{df}{dt} \approx \frac{f(t) - f(t - n\Delta t)}{n\Delta t}, \quad \textbf{(A2)}$$

where Δt denotes the sampling interval and n=5 for this example.

Derivative leading and amplitude control. We first isolate the basic leading effect obtained by combining the signal with its scaled causal derivative,

$$f_{lead}(t) = f(t) + A \cdot \frac{df}{dt}. \quad \textbf{(A3)}$$

The top panel of *Fig. 9* shows that the derivative term induces a systematic phase advance relative to f(t). This effect is most pronounced near zero-crossings, where the local slope is largest. The amplitude parameter A controls the strength of the leading behavior: increasing A yields stronger anticipation but also amplifies high-frequency fluctuations.

Adaptive gating via λ parameters. In financial time series, derivative information is not uniformly informative across regimes. This motivates the adaptive gating mechanism,

$$f_{out}(t) = c_1(t)f(t) + c_2(t)f'(t). \quad \textbf{(A4)}$$

with,

$$c_1(t) = \tanh(|\lambda_1 f(t)|);\ c_2(t) = A[1 - \tanh(|\lambda_2 f(t)|)]. \quad \textbf{(A5)}$$

The bottom panel of *Fig. 9* illustrates the role of the parameter λ_2 for fixed A=10. As λ_2 increases, the derivative contribution is progressively suppressed when $|f(t)|$ is large (strong trends) and concentrated near zero-crossings (regime transitions). This mechanism focuses the leading effect precisely where anticipation is most valuable, while avoiding excessive amplification during persistent trends. Importantly, the figure also clarifies the role of λ_1. In the limit $\lambda_1 \to \infty$, the coefficient $c_1(t)$ converges to unity for all t, and the gated construction reduces to

$$f_{out}(t) = f(t) + c_2(t)f'(t). \quad \textbf{(A6)}$$

In this regime, the red curve in the right panel corresponding to $\lambda_2 = 0$ coincides exactly with the red curve in the left panel, which represents the ungated derivative-enhanced signal $f(t) + Adf/dt$. Thus, the left panel can be interpreted as a limiting case of the full gated model, providing a clean reference for understanding how adaptive gating modulates the leading effect.

These illustrative examples motivate the structure of the empirical trading signal used in this study: 1)causal derivatives can induce systematic phase advance without look-ahead bias, 2) the amplitude parameter A controls the strength of this anticipation, and 3) the gating parameters λ_1 and λ_2 regulate when derivative information is incorporated, emphasizing regime transitions while suppressing noise during strong trends.

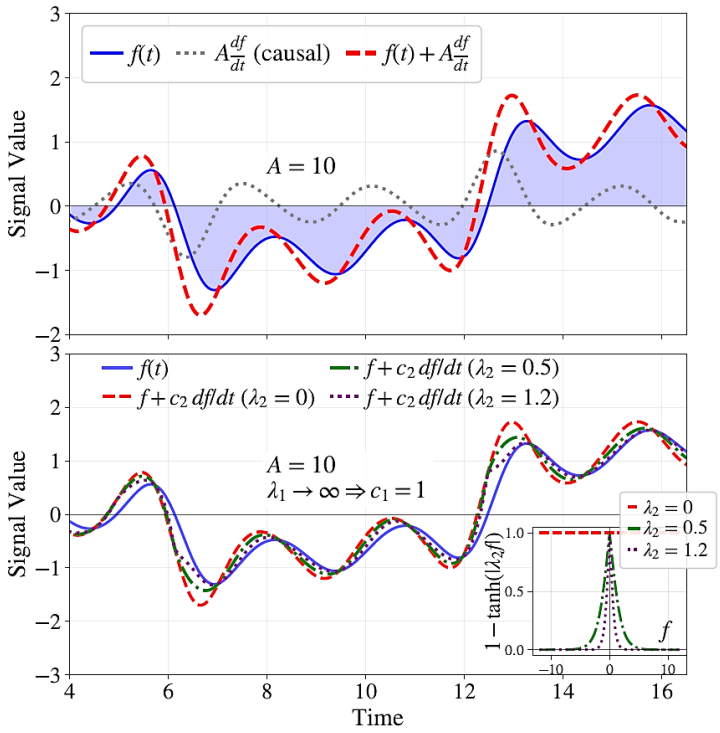

Fig. 9. Derivative-based signal leading under strict causality.

Top: the signal f(t) combined with its scaled causal derivative, f(t) + Adf/dt, exhibits a systematic phase advance, most visible near zero-crossings. Bottom: adaptive gating of the derivative term for different λ_2 with $A = 10$ and $\lambda_1 \to \infty$ ($c_1 = 1$). Increasing λ_2 suppresses the derivative contribution away from zero-crossings. For $\lambda_2 = 0$, the gated signal coincides with the derivative-enhanced signal shown in the top panel.