The PRISM Signal Framework

Adaptive Momentum, Path-Constrained Divergence, and Cycle-Phase Conditioning for BTC/USD

PRISM CHARTS · Technical Note 2026-08 · August 8, 2026

Abstract. *PRISM is a signal framework for BTC/USD built on three observations that classical momentum tools do not accommodate: that a single lookback is a fixed compromise across regimes, that a fixed oscillator threshold carries different information at different realised volatilities, and that Bitcoin is the only major asset with a scheduled, pre-announced supply shock. The framework runs two momentum horizons concurrently and blends them continuously against a volatility-regime statistic; it qualifies turns against a recency-of-extreme condition; and it replaces the conventional two-point divergence test with a path-constrained sequence carrying an explicit invalidation predicate. A second, independent layer projects reversion levels by inverting a strength measure, and a confluence layer marks the bars on which the two agree. Section 7 states the classical divergence test formally, identifies three structural defects in it, and gives the constraint schema PRISM uses in its place. Calibration constants are withheld throughout; this note is a specification of behaviour, not of implementation. A companion strategy, Prism Cycle, applies the same market premise mechanically; its disclosed test record appears in Section 13.*

1 Why the classical tools underperform on this series

The momentum instruments on a typical crypto chart were designed between 1978 and the early 1980s for equity and futures markets. They are good instruments, and PRISM is not a replacement for any of them; a substantial share of our users run one alongside it. Three of their design assumptions, however, do not survive contact with this series.

1. **A single lookback is a fixed compromise.** Every classical oscillator commits to one responsiveness parameter. A lookback fast enough to resolve a turn during range expansion is intolerably noisy in compressed tape; one calm enough for compressed tape resolves late once range opens. The user selects a constant and accepts being mis-specified in one of the two regimes at all times.
2. **Fixed thresholds are not scale-invariant.** An oscillator reading of a given magnitude does not carry the same information in a market realising three percent daily as in one realising half a percent. Applying an identical numeric trigger to both is a category error the classics do not address.
3. **No instrument carries a calendar.** Bitcoin's issuance schedule is public and known years in advance. No general-purpose indicator has any mechanism for locating the current bar within that schedule, because none was designed for an asset that has one.

   PRISM addresses each directly.

2 Market premise: compression and expansion

BTC/USD spends the majority of its calendar time inside horizontal ranges and the majority of its price change inside a small number of short expansions. The distribution of returns is therefore severely non-uniform in time. A system correctly positioned for a handful of transitions per cycle captures a disproportionate share of the available move; one that transacts continuously through compression pays fees and slippage to remain approximately flat.

Two consequences follow, and they define the two products described here. First, *the transition dominates the trend*: being early into an expansion is worth a great deal, and marginal improvement in reading the middle of one is worth very little. This motivates a framework oriented around identifying turns rather than scoring trend strength. Second, *compression is itself information*: a range that has persisted is evidence about the conditional distribution of the next move, independent of any momentum reading. This motivates the mechanical system described in Section 12.

3 Architecture

The framework is constructed as two independent readings rather than one. The wave engine measures displacement of price against a reference and rescales it; the reference, the rescaling and the horizons are [WITHHELD]. The band engine proceeds from an unrelated premise: it inverts a strength measure to solve for the price levels at which that measure would attain stated extremes, and treats those levels as prospective reversion points.

The two layers can and do disagree, which is the point. Each emits its own event class, and a confluence layer identifies the bars on which they agree and a divergence structure has also completed. Two contextual inputs modulate the stack: a continuous volatility-regime statistic, and the position of the current bar within the issuance cycle.

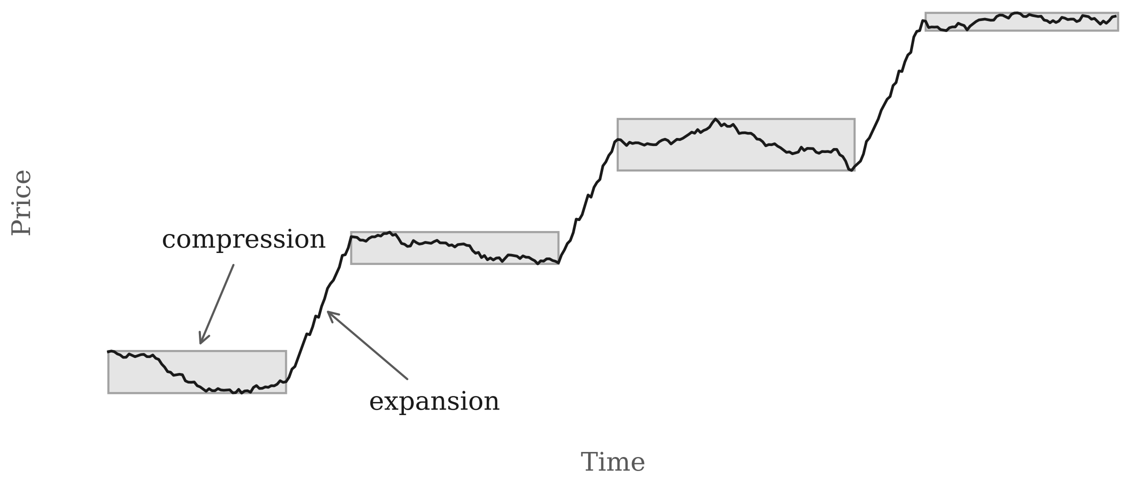

Figure 1. Schematic of the compression/expansion alternation that motivates the framework. Illustrative; no production data is shown.

4 Layer 1: regime-weighted dual horizon

The distinguishing feature of the current release is that the engine does not commit to a single lookback. It evaluates displacement at two horizons concurrently, one deliberately fast and one deliberately slow, and forms the output as a continuously weighted blend. Writing $w_t \in [0,1]$ for the weight on the fast leg and g_t for the regime statistic, the blend has the schematic form

$$o_t \;=\; w_t\,\phi_t^{\text{fast}} \;+\; \big(1-w_t\big)\,\phi_t^{\text{slow}}, \qquad\qquad w_t = W(g_t), \tag{1}$$

where ϕ^{fast} and ϕ^{slow} denote two independently parameterised displacement measures. The horizon lengths, the construction and rescaling of each measure, the definition of the regime statistic and the form of the weighting map W are [WITHHELD].

Three properties are worth stating. The weighting is continuous, so there is no bar at which the oscillator discontinuously changes character; a hard regime classifier would introduce a discontinuity precisely at the moments the system is relied upon. The blend is convex, so o_t is bounded by the two legs and cannot overshoot either. And w_t depends only on information available at t, so the construction is causal.

single-horizon engine over the full available history, the large majority of signals fall within two bars of their previous location, at negligible cost in turn quality as defined in Appendix A. The modification is surgical: it alters behaviour where the fixed-lookback formulation misbehaves and leaves it alone elsewhere.

5 Signal qualification

A raw momentum turn is not a signal. Let $\mathscr{C}$ denote the set of bars at which a turn in the wave is registered. The qualification stage admits $t \in \mathscr{C}$ only if the oscillator visited an extreme zone within a bounded recent window:

$$t \in \mathscr{S} \iff t \in \mathscr{C} \,\wedge\, \exists s \in [t-m, t] : \; o_s < \theta^- \qquad \text{(buy side; sell side mirrors)}, \tag{2}$$

with the zone level θ^- and the memory length m [WITHHELD]. The effect is to eliminate the dominant category of false positive in oscillator trading, the mid-range crossover: if the market never became stretched, there is no reversion to detect. Optional trend-alignment and minimum-movement filters are available and ship disabled; default behaviour does not depend on them.

Members of $\mathscr{S}$ are rendered as *qualified signals*: large dots in the oscillator pane, mirrored as triangles on price.

6 The divergence framework

PRISM recognises a divergence only when a four-step structure completes in order.

1. **Anchor.** A qualified signal prints at an extreme and becomes the reference point. Nothing is drawn.
2. **Transit.** The wave must travel fully through the midline and establish itself on the opposite side. A shallow crossing does not count; the requirement is enforced with hysteresis.

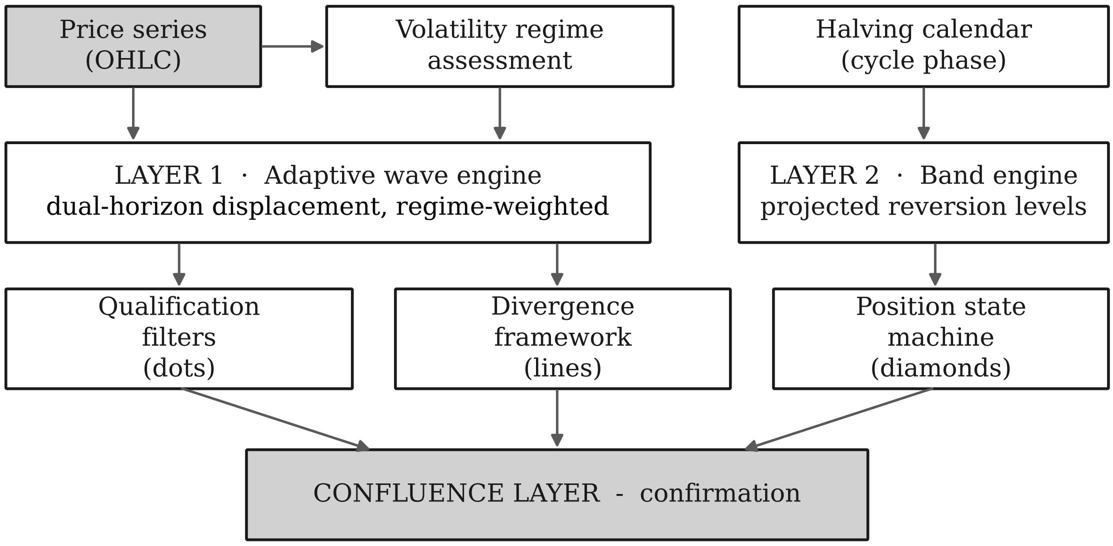

Figure 2. Layer structure. Two independent readings of the market are qualified separately and reconciled by a confluence layer.

3. **Return.** The wave returns toward the original extreme but falls short of it by a required margin, while price may have made an equal or worse extreme.
4. **Recognition.** Only now is the structure marked, in the pane and on price simultaneously.

Three behaviours distinguish the implementation. The connection *walks* to a better candidate rather than committing to the first eligible one. The structure can be *invalidated*: a genuinely deeper excursion than the anchor demonstrates the move was still running, and the anchor is retired, with already-drawn lines preserved as record. And same-structure connections are *stripped* when a new anchor confirms, because they were bounces within one leg rather than a two-leg divergence.

7 Formal statement: the classical test, and the constraint schema that replaces it

This section is the technical core of the note. It states the conventional divergence test in full, identifies three structural defects, and gives the schema PRISM substitutes. No calibration constant is disclosed.

7.1 The classical test

Let p_t denote price and o_t an oscillator. Define a *pivot low of order w* and a *pivot high of order w* by

$$\Pi_w^-(t) \iff p_t = \min_{|s-t|\le w} p_s, \qquad \Pi_w^+(t) \iff p_t = \max_{|s-t|\le w} p_s. \tag{3}$$

For two consecutive pivots of matching type at $t_1 < t_2$ with $t_2 - t_1 \le L$, write $\Delta p = p_{t_2} - p_{t_1}$ and $\Delta o = o_{t_2} - o_{t_1}$. The four conventional classes are then

$$\begin{aligned} &\text{regular bullish}: \quad \Pi_w^-,\ \Delta p < 0 \wedge \Delta o > 0, \qquad &&\text{regular bearish}: \quad \Pi_w^+,\ \Delta p > 0 \wedge \Delta o < 0,\\ &\text{hidden bullish}: \quad \Pi_w^-,\ \Delta p > 0 \wedge \Delta o < 0, \qquad &&\text{hidden bearish}: \quad \Pi_w^+,\ \Delta p < 0 \wedge \Delta o > 0, \end{aligned} \tag{4}$$

all four of which reduce to the single sign condition $\operatorname{sgn}(\Delta p) \neq \operatorname{sgn}(\Delta o)$ evaluated at pivots of the appropriate type. This is the entire test as it is implemented in the overwhelming majority of published divergence indicators.

7.2 Three structural defects

(i) Pivot confirmation is not causal. By construction $\Pi_w(t_2)$ cannot be evaluated before bar $t_2 + w$. Any tool that displays the divergence at t_2 is plotting a decision it could not have made at t_2. On a historical chart the marker appears at the extreme; in real time it appears w bars later, or appears and then withdraws. This is the mechanism behind most repainting complaints, and it is structural rather than a defect of any particular implementation.

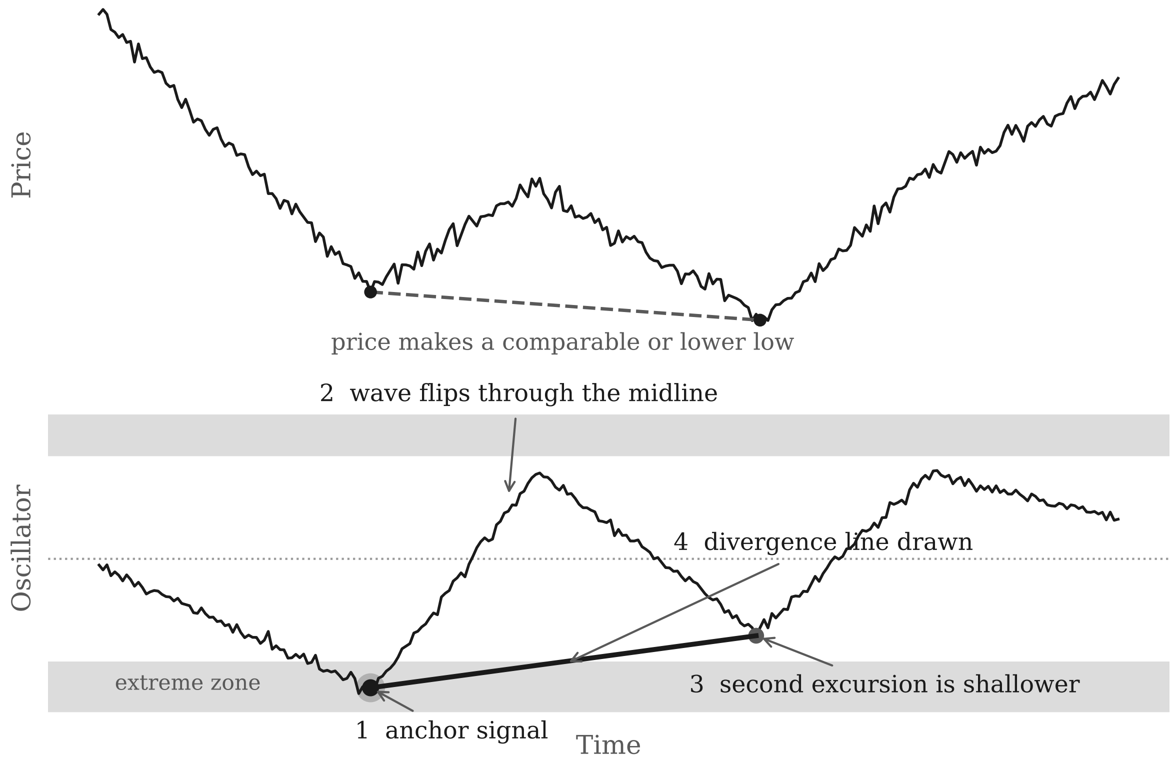

Figure 3. The path-constrained divergence sequence: anchor, full midline transit, shallower return, recognition. Schematic; synthetic data.

(ii) Pair eligibility is unconstrained. The test imposes no condition whatsoever on the path $\{o_\tau\}_{t_1<\tau<t_2}$. Consequently t_1 and t_2 may lie inside one continuous decline, where $\Delta o > 0$ arises mechanically from the oscillator's own bounded range rather than from any change in the underlying dynamics. The test therefore admits a large family of pairs that carry no information, and it has no vocabulary in which to reject them.

(iii) There is no magnitude requirement. The condition $\Delta o > 0$ is satisfied by arbitrarily small Δo. If the oscillator carries additive noise with non-degenerate variance, then as $|\Delta o| \to 0$ the probability that the observed sign reflects noise rather than signal approaches $\frac{1}{2}$. A test that fires on $\Delta o = \varepsilon$ is, in that limit, a fair coin.

7.3 The PRISM schema

PRISM replaces the two-point comparison with a path-constrained sequence. The recognition predicate takes the form

$$\mathscr{D}(t_1,t_2) \;=\; \underbrace{\mathbf{1}[\operatorname{sgn}\Delta p \neq \operatorname{sgn}\Delta o]}_{\text{classical condition}} \cdot \underbrace{\mathbf{1}[\,\exists\,\tau \in (t_1,t_2) : o_\tau > +\eta\,]}_{\text{midline transit, hysteresis } \eta} \cdot \; \underbrace{\mathbf{1}[\,|\Delta o| \geq \delta\,]}_{\text{separation margin } \delta} \;\cdot \underbrace{\mathbf{1}\left[\min_{t_1<\tau<t_2} o_\tau \;>\; o_{t_1}\right]}_{\text{invalidation predicate}} \tag{5}$$

evaluated on the buy side, with the sell side mirroring every term. Four differences follow directly.

- **The anchors are qualified signals, not price pivots.** $t_1, t_2 \in \mathscr{S}$ as defined in Section 5. Membership of $\mathscr{S}$ is decidable at the close of bar t, so defect (i) does not arise: every term of $\mathscr{D}$ is causal, and the predicate is evaluated only on confirmed bars.
- **The transit term addresses defect (ii) directly.** It is precisely a condition on the path between the anchors, and it is what the classical test lacks. Requiring the trajectory to clear $+\eta$ enforces that the two anchors belong to distinct legs rather than to one continuous move.
- **The margin term addresses defect (iii).** Requiring $|\Delta o| \geq \delta$ excludes the noise-dominated neighbourhood of $\Delta o = 0$.
- **The invalidation term has no classical counterpart.** If the path makes an excursion deeper than the anchor, the interpretation was premature; the anchor is retired and generates no further connections. Standard implementations continue to fire from reference points already falsified by subsequent price action.

The thresholds η and δ, the qualification parameters of Section 5, and the expiry bound on $t_2 - t_1$ are [WITHHELD]. The schema above is complete as a statement of what is required; it is deliberately incomplete as a statement of how it is computed.

7.4 Structural divergence

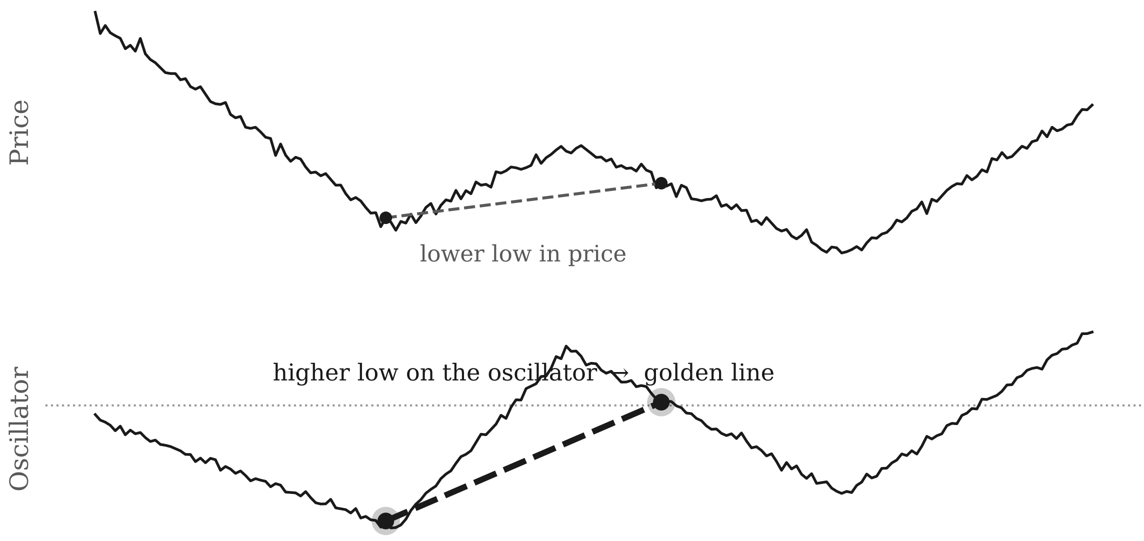

Figure 4. Structural divergence between consecutive qualified anchors. Both endpoints are members of $\mathscr{S}$, which is what makes the case rare.

A stronger case obtains when two consecutive anchors diverge against each other directly, that is when $t_1, t_2 \in \mathscr{S}$ are adjacent in $\mathscr{S}$ and $\mathscr{D}(t_1, t_2)$ holds. Because both endpoints are fully qualified rather than arbitrary local extrema, this is comparatively rare and correspondingly informative. It is rendered distinctly, alerted separately, and fires once per pair: an anchor that closes a structure cannot open the next one.

8 The confluence layer

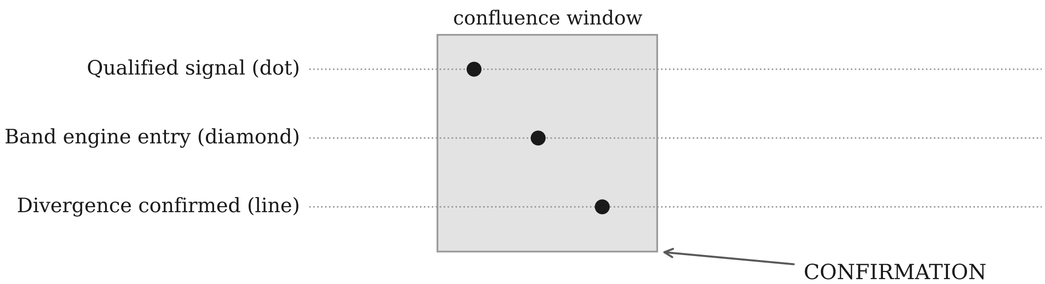

Figure 5. Confluence. Three independently generated event classes must fall inside a bounded window before a confirmation is marked.

A confirmation is marked when qualified signals, band entries and completed divergences each have a member inside a bounded window centred on the same bar. The width of that window is [WITHHELD]. We make no claim that confirmations are superior in expectation, and state plainly that they are rare enough that a user trading them exclusively will trade infrequently. The layer answers one question, whether a reading is supported by anything other than the oscillator, and answers it visibly rather than leaving it to discretionary judgement under pressure.

9 Layer 2: band engine and position framework

The second reading inverts a strength measure. Rather than reporting where strength currently stands, the band engine solves for the price levels at which strength would attain stated extremes and projects those levels forward. Two behaviours materially affect what the user sees.

Compression widening. When realised volatility contracts relative to its own baseline, the projected levels widen. This is an anti-chop measure: in precisely the conditions where a mean-reversion system produces its worst signals, the levels move out of reach. The trigger, the magnitude of the widening and its release are [WITHHELD].

Position state machine. Band events are not independent markers. The engine maintains explicit position state rather than evaluating each event in isolation, and manages an open position through to its exit. Repeated entry into a direction that has already failed is prevented. The state transitions, the bracket construction and the reset condition are [WITHHELD].

10 Cycle-phase conditioning

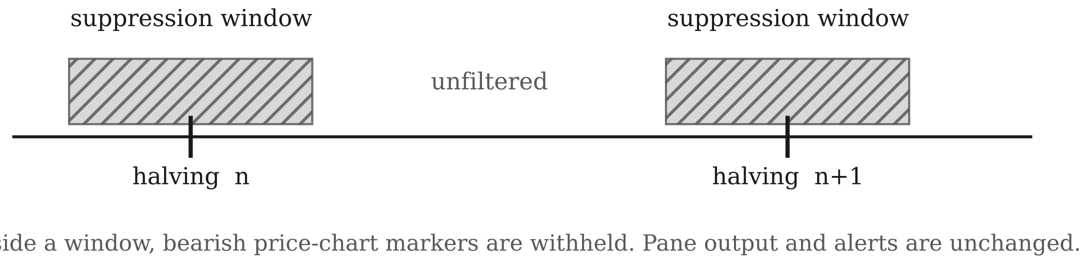

Figure 6. Cycle-phase windows. Inside a window, bearish price-chart markers are withheld; pane output and alerts are unaffected. Widths are withheld and not to scale.

The framework maintains the historical halving timestamps and the projected next one, and derives a phase position for every bar. This drives two independent mechanisms which should not be conflated.

- **Cycle filter.** Inside the expansion phase, short-side band entries are suppressed. This is a signal filter and it changes the behaviour of the position state machine.
- **Bitcoin Only.** A display switch, off by default, which withholds bearish markers from the price chart inside the window. Pane output, the underlying signal set and every alert are untouched. It engages only above a minimum resolution, on the reasoning that a multi-year thesis should not modify an intraday chart.

Both are switchable and both should be disabled on instruments without a scheduled supply mechanism.

11 Design commitments

11.1 Causality and non-repainting

Every predicate in this note is evaluated on confirmed bars, and every term is a function of information available at or before the bar on which it fires. A marker that has printed does not move, vanish or change class on a subsequent tick. This is an architectural commitment rather than a setting, and Section 7.2(i) explains why it is the single most consequential property to verify in any signal product.

11.2 Fixed calibration

The distributed build exposes display, direction and risk controls but not the values that define engine behaviour. The reasoning is stated openly rather than presented as a feature. A calibrated system produces one reading; when every parameter is exposed, two users see different signals from the same chart, support becomes intractable, and, most importantly, the user is implicitly invited to tune parameters until recent history looks favourable, which is the most reliable known method for destroying out-of-sample performance.

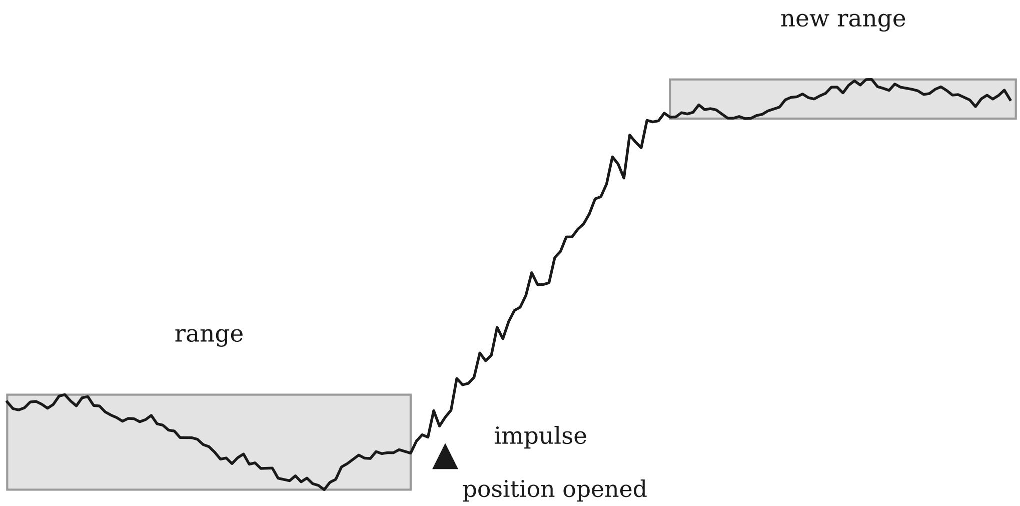

Figure 7. Prism Cycle: positioning inside a compression phase ahead of the expansion rather than entering after the breakout is established.

12 Prism Cycle

Prism Cycle is a separate product that applies the premise of Section 2 mechanically rather than through the oscillator framework. It maintains a directional state and takes positions against short-term moves within it, entering on weakness inside strength and on strength inside weakness. Its directional construction, its entry and exit conditions and its calibration are [WITHHELD]. It is directional-configurable: long only, short only, or both.

The long-only configuration is a systematic middle ground between two unsatisfactory alternatives, holding spot continuously through every drawdown and attempting to time the market discretionarily. It gives a mechanical rule for exposure during expansion phases and flatness during contractions, and composes naturally with a dollar-cost-averaging programme: the signal set can modulate accumulation rate rather than being used to trade in and out. It is intended for the daily resolution.

13 Disclosed test record

The following are results of Prism Cycle as measured in TradingView's Strategy Tester, reported with the conditions a reader requires in order to judge them.

Measured from	Trades	Profitable	Profit factor	Max drawdown	Net return
The 2013 top	150	39%	1.9	61.97%	+52,000%
The 2017 top	100	38%	1.9	37.84%	+8,100%
The 2022 top	70	37%	1.7	32.72%	+320%

Conditions. BTC/USD (Bitstamp), daily resolution. Signals evaluated on the close of a confirmed bar; orders filled at the open of the following bar, so the strategy never transacts at the price that generated its own signal. Commission 0.05% per side and two ticks of slippage applied to every fill. Position sizing as a percentage of current equity, so results compound. Reported figures are rounded and maximum drawdown is exact; the calibration is [WITHHELD].

13.1 Interpretation

These are hypothetical backtest results, not live trading results. *No representation is made that any account will or is likely to achieve profits or losses similar to those shown. Backtested performance has inherent limitations, does not reflect real-world liquidity constraints, exchange outages, funding costs or tax, and benefits from hindsight in the selection of the system's design.*

The hit rate is below 40% **and this is the expected profile.** A profit factor near 1.9 against a sub-40% hit rate describes a distribution in which winning trades are substantially larger than losing ones. This is the signature of trend and breakout capture and the correct shape for the premise of Section 2. A majority of individual positions lose; the edge is expressed across a sequence, not within any single trade.

Returns compound, so headline figures are not linear in window length. Percentage-of-equity sizing means reported net return reflects reinvestment throughout, and a longer measurement window mechanically produces a dramatically larger headline number. The three periods are therefore reported separately. The 2022 row is the most representative of a recent adopter's experience.

Drawdowns are severe. A maximum drawdown between roughly 33% and 62% is the central risk disclosure of this section, and any allocation should be sized on the assumption that a drawdown of this order recurs.

One instrument, one venue. Results on another exchange's feed for the same pair will differ, because the underlying series differs in gaps, wicks and session boundaries. Users should re-run on the venue they trade.

14 Provenance

Original to PRISM, and constituting the substance of the product, are: the regime-weighted dual-horizon adaptation (Section 4); the qualification stage and its recency-of-extreme condition (Section 5); the path-constrained divergence schema, including the transit, margin and invalidation terms (Section 7); the structural anchor-to-anchor case (Section 7.4); the confluence layer (Section 8); the position state machine, migrating brackets and direction locking (Section 9); compression widening (Section 9); the cycle-phase model and its two applications (Section 10); and the whole of Prism Cycle (Section 12).

A reader who concludes that the value of the framework lies in the layers above the oscillator rather than in the oscillator itself has understood the product correctly.

15 Risk disclosure

PRISM is analytical, decision-support software. It is **not financial, investment, tax or legal advice**, and nothing in this note is a recommendation, solicitation or offer to buy or sell any asset. Trading cryptocurrencies, equities, commodities and derivatives carries a substantial risk of loss and is not suitable for everyone; you may lose part or all of your capital. Signals are calculations, not instructions. All performance information in Section 13 is hypothetical and derived from backtesting; past performance never guarantees future results. You alone are responsible for your trading decisions.

A Evaluation metrics

The two statistics quoted in Section 4 are defined here so that the claims are checkable in principle.

A.1 Turn quality

For a signal set $\mathscr{S}$ and a symmetric window of half-width w, define for each buy signal the local percentile rank of price

$$F_t \;=\; \frac{1}{2w+1} \sum_{|s-t| \le w} \mathbf{1}\big[\, p_s \ge p_t \,\big], \qquad Q(\mathscr{S}) \;=\; \frac{2}{|\mathscr{S}|} \sum_{t \in \mathscr{S}} F_t, \tag{6}$$

with the inequality reversed for sell signals. Under signal placement independent of price, $\mathbb{E}[F_t] = \frac{1}{2}$ and hence $Q = 1$; under perfect turn selection $F_t = 1$ and $Q = 2$. The metric is deliberately defined without reference to any particular oscillator, so that competing engines can be scored on the same scale.

A.2 Signal agreement

For two signal sets and a tolerance of two bars,

$$A(\mathscr{S}_1, \mathscr{S}_2) \;=\; \frac{\big|\{\, t \in \mathscr{S}_1 \;:\; \exists u \in \mathscr{S}_2,\ |t-u| \le 2 \,\}\big|}{|\mathscr{S}_1|}. \tag{7}$$

The agreement figure quoted in Section 4 is this statistic evaluated between the current release and an unmodified single-horizon engine over the full available history.

A.3 Distinctness screen

Candidate engines are additionally screened against the standard oscillator family by Pearson correlation on the aligned series, with $\rho < 0.90$ against every reference member required before an engine is treated as materially distinct rather than as a reparameterisation. The screen is applied during development; it is reported here because a claim of novelty without such a test is not falsifiable.