Enigma in Economics
https://enigma.or.id/index.php/economy
e-ISSN: 3026-6696

Adaptive Quantile Calibration of Daily and Weekly Cycle-Low Forecasts in Bitcoin, S&P 500 Futures, and Gold

Muhammad Faiz[1*], Sonia Vernanda[2]

[1]*Department of Administrative Law, Seine Legal Institute, Manado, Indonesia*
[2]*Department of Mathematics Education, Emerald Education Center, Samarinda, Indonesia*

ARTICLE INFO

Article history:
Received: May 2, 2026
Accepted: June 1, 2026
Published: June 29, 2026

Keywords:
Bitcoin
Cycle theory
Market timing
Prediction interval
Quantile calibration

JEL classification:
G11; G12; G14; G17

***Corresponding author:**
Muhammad Faiz
muhammad.faiz@enigma.or.id

ORCID:
Muhammad Faiz: 0009-0004-8196-0632
Sonia Vernanda: 0009-0006-1036-2824

DOI:
https://doi.org/10.61996/economy.v4i1.131

All authors have reviewed and approved the final version of the manuscript.

ABSTRACT

Background. Market-cycle forecasts are vulnerable to hindsight because a low becomes identifiable only after subsequent price confirmation.

Objective. This study evaluated whether an adaptive, confirmation-aware interval could attain at least 80% chronological forecast precision for daily cycle lows (DCLs) and weekly cycle lows (WCLs) in Bitcoin, S&P 500 futures, and gold.

Methods. The Adaptive Quantile-Calibrated Cycle Window used only the latest 20 completed cycles. Its lower endpoint was the empirical 10th percentile of prior low-to-low durations, and its upper endpoint was the 90th percentile of prior-low-to-next-confirmation durations. Forecasts originating from 1 January 2021 through 14 July 2026 were evaluated sequentially, and the retrospective protocol was externally preregistered.

Results. Fixed clocks achieved 70.9% DCL precision and 55.6% WCL precision. The adaptive interval achieved 109/127 DCL hits (85.8%; 95% CI 78.7%–90.8%) and 27/27 WCL hits (100.0%; 95% CI 87.5%–100.0%). Mean window width increased from 14.7 to 32.8 days for DCL and from 4.0 to 11.7 weeks for WCL. A wider 5th–95th percentile band produced 93.7% DCL precision with a 95% lower confidence bound of 88.1%.

Conclusion. Adaptive interval calibration exceeded the 80% point target, but the gain depended on materially wider windows and a small WCL sample; prospective replication remains necessary.

1. Introduction

Market-cycle analysis translates irregular price paths into a hierarchy of daily, intermediate, and weekly lows. In practitioner Cycle Theory, counts are measured from low to low and are used as temporal priors rather than deterministic laws. The empirical difficulty is that completed charts permit several plausible retrospective labels. A credible evaluation must therefore define the price field, separate the date of the low from the later date on which it becomes knowable, preserve failed cycles, and freeze each forecast before its outcome is observed.

The three markets studied here have distinct temporal anchors. Equity research documents political and presidential return cycles and long-horizon valuation mean reversion; gold research identifies seasonal and macro-financial regularities; and Bitcoin alone has a mechanically scheduled halving approximately every four years[1-6]. These anchors are not commensurable: political policy, jewellery and hedging demand, and programmed issuance are different mechanisms. Their common implication is narrower—each market can contain non-random temporal structure without being governed by one universal clock.

Any timing claim must also confront efficient-market reasoning. The efficient market hypothesis predicts that stable historical rules should be competed away; the adaptive markets hypothesis instead treats efficiency as state-dependent, while limits to arbitrage explain why anomalies may persist when exploitation is risky and costly. Long-wave, prospect-theory, behavioral-finance, sentiment, and gold-survey literatures further show that slow adjustment and feedback can organize recurrent market phases[7-14]. This evidence supports a probabilistic and regime-conditioned interpretation of cycle counts, not a permanently accurate calendar.

The literature foundation is organized transparently in line with established narrative-review and evidence-reporting principles[15-16]. Foundational momentum evidence and critical-point models provide two complementary mechanisms: persistence can extend a move, whereas nonlinear feedback can terminate it abruptly[17-18]. For a low-window forecast, both mechanisms matter because the next trough can be delayed by persistence or accelerated by an unstable reversal.

Bitcoin's halving is the clearest scheduled cycle anchor, but recent studies show heterogeneous effects across abnormal returns, volatility, higher-order spillovers, cross-coin responses, speculative super-cycles, and market sentiment[19-24]. Four completed halvings remain too few to justify a clockwork price law. The schedule is known in advance, and its observed footprint is entangled with market maturation, liquidity, and macroeconomic conditions.

Time-varying efficiency provides the strongest direct motivation for adaptation. Very-long-run equity evidence, business-cycle predictability, and momentum studies show that forecasting strength changes with market state. Bitcoin research likewise progresses from early inefficiency findings to dynamic, alternating, liquidity-dependent, and long-memory efficiency regimes; wavelet evidence shows that dominant relationships change across both time and scale[25-38]. A rolling duration distribution is therefore more defensible than a fixed historical count, although rolling estimation cannot eliminate abrupt regime-break risk.

Technical timing evidence is similarly conditional. Classic moving-average and trading-range rules produced encouraging historical results, but data-snooping tests weakened many claims. Later equity and cryptocurrency studies recover predictive value in selected regimes, while demonstrating sensitivity to sample choice, risk premia, transaction costs, bubble states, and the difficulty of valuing the underlying asset[39-47]. These findings justify a transparent chronological evaluation and prohibit interpreting interval coverage as trading profitability.

Calendar effects provide a further caution. Halloween, January, weekend, seasonal-affective, day-of-week, and cryptocurrency calendar patterns have all been reported, but their persistence varies across subperiods and specifications. Momentum is more robust across equities, commodities, gold, and cryptocurrency, yet it also changes with market state[48-65]. Cycle translation can therefore be interpreted as a time-local momentum condition, whereas rigid day and week counts remain hypotheses to be tested rather than accepted periodicities.

Behavioral amplification supplies a plausible mechanism for irregular cycle lengths. Cryptocurrency herding, sentiment, liquidity, fear of missing out, and investor attention intensify or attenuate demand near local extremes. Bubble research using explosive-root and log-periodic approaches documents recurrent speculative regimes in equities and Bitcoin, but also shows substantial uncertainty around the timing of the critical reversal[66-81]. Consequently, a forecast window should accommodate recognition lag and uncertainty rather than promise an exact bottom date.

Macro-liquidity and safe-haven evidence reinforces the cross-asset design. Bitcoin responds to monetary-policy shocks and global investment demand; gold alternates among hedge, safe-haven, and ordinary-asset roles; and connectedness between cryptocurrency, equities, and gold changes during crises[82-95]. These state changes imply that an interval estimated from old cycles can become stale, making sequential updating and explicit sensitivity bands essential.

Finally, market structure is changing the information set used to identify cycles. On-chain indicators add observable cryptocurrency fundamentals, but stock-to-flow and simple network-scaling rules have weak predictive foundations. Volatility is regime-switching and structurally unstable, while spot Bitcoin exchange-traded products have altered flows and price discovery since 2024[96-112]. The implication is not that historical cycles are useless; it is that the timing rule must be causal, adaptive, and repeatedly re-evaluated as market structure evolves.

A fixed cycle clock is transparent and operationally convenient, but its binary success depends on the width assigned around the reference count. A narrow band may appear precise while systematically omitting regime-delayed lows. A very wide band may achieve high coverage while providing little timing information. The correct empirical question is therefore not whether a single count is always correct, but whether a sequential interval can remain calibrated while preserving useful sharpness. Calibration measures how often the defined target falls within the interval; sharpness measures how concentrated that interval remains. Both are necessary because a forecast covering most of the calendar would be statistically successful and operationally uninformative.

The previous fixed-window analysis did not sustain an 80% confirmation-aware rate. That negative result motivates a distributional reformulation. The latest 20 completed cycles are treated as a rolling empirical distribution. The lower endpoint estimates when the next low becomes plausible from prior low-to-low durations. The upper endpoint estimates when confirmation of that low should have occurred from prior low-to-next-confirmation durations. This asymmetric construction is intentional: an interval should not receive credit for a low that happened within the calendar band but could not yet have been recognized by its closing date.

The study addresses five questions. First, do the fixed Cycle Theory clocks remain below the 80% target in a 2021-2026 chronological evaluation? Second, does a common adaptive 10th-90th percentile interval exceed the target for both DCL and WCL forecasts? Third, are point estimates above 80% observed in

Bitcoin, S&P 500 futures, and gold without asset-specific retuning? Fourth, how much width is required to obtain that coverage? Fifth, does a wider 5th-95th percentile sensitivity band place the pooled DCL lower confidence bound above 80%?

The contribution is fourfold. First, the study formalizes confirmation-aware interval coverage for daily cycle lows (DCLs) and weekly cycle lows (WCLs). Second, it proposes an auditable Adaptive Quantile-Calibrated Cycle Window (AQCW) that can be reproduced without proprietary optimization. Third, it evaluates one common sequential algorithm over a 2021-2026 chronological pseudo-out-of-sample segment under an externally preregistered retrospective protocol. Fourth, it reports the cost of reaching the target through confidence intervals, asset-level results, sequential paths, and the complete precision-width frontier. The resulting claim is deliberately bounded: AQCW is a calibrated monitoring-window method, not proof of universal periodicity, exact bottom timing, or trading profitability.

2. Methods

2.1 Operational cycle hierarchy

A DCL is an important daily-chart Low followed by sufficient recovery evidence. A WCL is the corresponding higher-order Low on Monday-start weekly bars. An intermediate half-cycle state may organize translation and strength, but it is not treated as a substitute for the next DCL. The low date is the date of the minimum OHLC Low; the confirmation date is the first later date at which the stated recovery rule is satisfied. This separation removes same-bar hindsight.

The fixed reference counts are 60 calendar days for Bitcoin, 40 for S&P 500 futures, and 35 for gold. The corresponding DCL full widths are 14, 16, and 14 days. WCL counts are 30, 24, and 26 weeks, each with a fixed +/-2-week band. These counts are evaluated as theory priors, not as facts imposed on the labels.

The asset-specific clocks, interval widths, and reconstruction samples are detailed in Table 1. The table makes the design asymmetry visible: DCL reference counts differ by asset, whereas every fixed WCL comparator uses the same four-week full width. These settings are frozen and are not widened after a miss.

Table 1. Asset-specific fixed clocks, interval widths, and reconstruction samples.

Asset	DCL clock	DCL width	WCL clock	WCL width	Daily rows	Weekly origins
Bitcoin	60 days	14 days	30 weeks	4 weeks	4,318	17
S&P 500 futures	40 days	16 days	24 weeks	4 weeks	6,517	50
Gold	35 days	14 days	26 weeks	4 weeks	6,489	46

Note. Daily rows are the cleaned OHLC histories; weekly origins are completed primary WCL forecasts across the full reconstruction.

2.2 Testable hypotheses

H1 (fixed-clock limitation): fixed Cycle Theory windows do not sustain 80% confirmation-aware precision in the 2021-2026 chronological evaluation. H2 (adaptive calibration): the AQCW point estimate reaches at least 80% for both DCL and WCL in that evaluation. H3 (cross-asset replication): each asset records an adaptive point estimate of at least 80%. H4 (calibration-sharpness trade-off): achieving the target requires materially wider windows than the fixed clocks. H5 (robustness): a wider 5th-95th percentile band retains at least 80% precision with a pooled 95% confidence lower bound above 80% for DCL.

2.3 Data and event construction

The analysis uses daily OHLC histories for BTC-USD, continuous S&P 500 futures (ES=F), and continuous COMEX gold futures (GC=F), frozen on 14 July 2026. Samples begin on 17 September 2014 for Bitcoin, 18 September 2000 for S&P 500 futures, and 30 August 2000 for gold. Weekly bars are aggregated with Monday labels and contain the first Open, maximum High, minimum Low, and last Close of each week. Completed-candle overlaps were reconciled before event construction.

The primary daily detector identifies a candidate when OHLC Low is the minimum in a centered seven-bar neighborhood. Recognition is delayed until the right-hand neighborhood is observed. A candidate is confirmed within the next ten bars when the Close reclaims its 10-period moving average, closes above the short high preceding the low, or posts two consecutive higher closes. Weekly candidates use a three-week neighborhood and a four-week confirmation horizon. Every low date, confirmation date, lag, price, and selection status is stored separately.

Cycles advance sequentially. After a confirmed seed, a fixed-clock forecast is issued. If no eligible low is confirmed by the window end, the outcome is recorded as a miss and the first later confirmed low becomes a transparent reset anchor. This rule prevents the failed window from being backfilled after the fact. The same anchor path is used to estimate the adaptive duration distribution.

Before event construction, each asset file is sorted chronologically and checked against the retained source manifest for first date, final freeze date, and row count. Completed-candle overlaps are reconciled so that a timestamp contributes one OHLC record to the analytical series. Weekly aggregation uses only information contained in the underlying daily bars and never assigns a weekly low before that week's final completed observation. These controls are important because a duplicate bar, an incomplete terminal candle, or a shifted weekly label could change both the selected low and the measured confirmation lag.

The use of continuous futures series is intended to study timing structure rather than executable contract-level returns. No adjustment is made for transaction costs, roll slippage, financing, or venue-specific price differences because the outcome is an interval-coverage score rather than a trading profit. Bitcoin is represented by the retained BTC-USD history, while S&P 500 and gold are represented by continuous futures histories. The common event logic is applied to all three markets after cleaning; there is no asset-specific manual relabeling of lows in the 2021-2026 evaluation segment.

The analytical freeze date is 14 July 2026. Consequently, the final eligible forecast origin differs from the final price observation: an origin enters the evaluation only when the next low and its confirmation can be assessed under the frozen rules. This distinction prevents incomplete terminal events from being silently classified as misses or hits. The retained ledgers record the low date, known date, decision date, lag, selected status, and reset status, enabling the chronology to be reconstructed from the machine-readable files.

2.4 Confirmation-aware precision

A forecast is a hit only if the next selected cycle low lies on or after the forecast start and its confirmation occurs on or before the forecast end. Because every eligible origin issues exactly one interval and has exactly one next-cycle target, the reported precision is numerically equal to empirical interval coverage. It should not be confused with classification precision under selective abstention. The common target ledger and one-forecast-per-origin rule keep the fixed and adaptive procedures directly comparable.

Let S and E denote the forecast start and end, respectively; let the superscripts low and conf distinguish the realized low date from its later confirmation date. The binary score for origin t is defined as follows:

$$H_t = 1\{\tau_t^{\text{low}} \geq S_t \ \wedge \ \tau_t^{\text{conf}} \leq E_t\}.$$

The indicator equals one only when both inequalities hold. Precision is the arithmetic mean of these indicators over the eligible origins. Wilson 95% confidence intervals quantify finite-sample uncertainty because the pooled DCL and especially WCL samples are not large enough for normal approximations to be taken for granted. Mean and median window widths are reported in the same time unit as the forecast so that coverage cannot be interpreted independently of sharpness.

2.5 Adaptive Quantile-Calibrated Cycle Window

At each origin, AQCW uses only the most recent 20 completed cycles. For completed cycle i, low-to-low duration L and prior-low-to-next-confirmation duration C are defined separately. In daily analysis the unit is calendar days; in weekly analysis it is Monday-start weeks.

$$L_i = \tau_i^{\text{low}} - \tau_{i-1}^{\text{low}}; \quad C_i = \tau_i^{\text{conf}} - \tau_{i-1}^{\text{low}}.$$

Let a be the current confirmed anchor and let the script H denote the available rolling history. The primary AQCW interval adds the empirical 10th percentile of prior L values to the anchor for its start and the empirical 90th percentile of prior C values for its end:

$$W_t(0.10) = [a_t + Q_{0.10}(\mathcal{L}_t),\ a_t + Q_{0.90}(\mathcal{C}_t)].$$

The use of C at the upper endpoint incorporates recognition lag directly: a low that occurs within the calendar range but remains unconfirmed at the end is not counted as a success. At least eight completed cycles are required before forecasting begins. The 20-cycle cap allows old regimes to leave the estimation set while retaining enough observations to estimate empirical tails without an asset-specific model.

The rolling-history principle is motivated by evidence that return predictability, efficiency, long memory, and volatility change across market states[25-38,103-107]. AQCW remains an empirical quantile procedure: it does not claim distribution-free coverage or exchangeability. Its validity in this manuscript rests on chronological replay, transparent hit definitions, and fully reported interval width rather than on a formal asymptotic guarantee.

The primary specification uses a 20-cycle lookback and a 10th-90th percentile band. The training-validation segment begins in 2017 for DCL and 2018 for WCL and ends on 31 December 2020. Forecast origins from 1 January 2021 onward form a chronological pseudo-out-of-sample segment. Every forecast uses only previously completed durations. The study uses retrospective market histories under an externally preregistered protocol; the chronological segment is causal at each origin but is not equivalent to a live prospective trial. After a realized cycle is confirmed, its durations may enter the next rolling history; no future duration contributes to its own forecast.

The frozen implementation is summarized in Table 2. Reading the table vertically shows the complete information flow from confirmed anchor, to rolling history, to interval endpoints, to scoring, and finally to the sequential update. This compact specification is also the reproducibility contract used by the audit code.

Table 2. Frozen implementation of the Adaptive Quantile-Calibrated Cycle Window.

Component	Operational definition
Anchor	Confirmed OHLC Low; low date and confirmation date remain separate.
History	Latest 20 completed cycles; minimum eight cycles before first forecast.
Start	Current anchor + empirical 10th percentile of prior low-to-low durations.
End	Current anchor + empirical 90th percentile of prior-low-to-next-confirmation durations.
Hit	Next selected Low is not before start and its confirmation is not after end.
Update	Sequential after outcome confirmation; one rule is used across the evaluation.
Robustness	5th-95th percentile band and static training-distribution comparator.

Note. Durations are measured in calendar days for DCL and Monday-start weeks for WCL.

2.6 Baselines and robustness

The first baseline is the asset-specific fixed clock. The second is a static quantile interval estimated once from the pre-2021 segment and never updated. The static comparator distinguishes the value of rolling adaptation from the value of using a broad duration distribution. Sensitivity bands span central empirical ranges from 60% through 90%. Forecast quality is evaluated under the calibration-sharpness principle: precision is interpreted jointly with mean window width, and no band is preferred solely because it covers more outcomes.

The final reported model uses no price-derived machine-learning features and changes only timing uncertainty. It does not redefine lows, discard misses, alter the asset universe, or tune an entry/exit trading strategy. Nevertheless, alternative lookbacks and quantile bands were explored in this retrospective workflow; the sensitivity frontier is reported, and prospective confirmation is required.

3. Results

3.1 Adaptive windows exceed the 80% point target

H1 and H2 are supported. The fixed clocks record 90/127 DCL hits (70.9%; 95% CI 62.4%-78.1%) and 15/27 WCL hits (55.6%; 95% CI 37.3%-72.4%). The primary adaptive interval records 109/127 DCL hits (85.8%; 95% CI 78.7%-90.8%) and 27/27 WCL hits (100.0%; 95% CI 87.5%-100.0%). The gain is 15.0 percentage points for DCL and 44.4 points for WCL.

The pooled fixed-versus-adaptive contrast and its Wilson intervals are shown in Figure 1. The four plotted points reproduce the narrative values exactly: 70.9% and 85.8% for DCL, and 55.6% and 100.0% for WCL. The corresponding training-validation estimates, hit counts, confidence intervals, and mean widths are detailed in Table 3, allowing the visual comparison to be checked against exact tabulated values.

The hit-count decomposition clarifies the magnitude of the improvement. Relative to the fixed DCL clock, the primary adaptive DCL rule converts 19 of the 37 fixed-clock misses into covered outcomes, reducing the miss count from 37 to 18 over the same 127 forecast origins. For WCL, the adaptive rule covers all 27 outcomes compared with 15 under the fixed four-week clock, an increase of 12 hits. These are paired descriptive comparisons because both methods are evaluated on the same realized targets; they should not be read as independent samples. The Wilson intervals quantify uncertainty around each coverage proportion, while the paired hit counts show that the gain is not produced by changing the evaluation set. This common-target design is essential: allowing one method to omit difficult origins would inflate its apparent precision and break comparability.

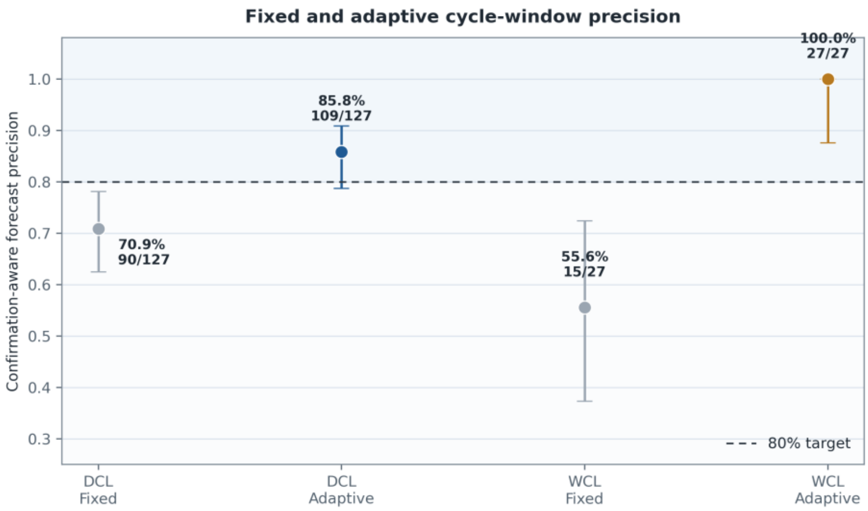

Figure 1. Confirmation-aware precision in the 2021-2026 chronological evaluation for fixed and adaptive DCL/WCL windows. Points show estimates; error bars show Wilson 95% confidence intervals; shading above the dashed line marks the 80% target region.

Table 3. Pooled confirmation-aware precision, hit counts, Wilson confidence intervals, and mean window widths for the fixed and adaptive specifications.

Method	Cycle	Evaluation	Hits	Precision	95% CI	Mean width
Fixed clock	DCL	2021-2026 eval.	90/127	70.9%	62.4%-78.1%	14.7 days
Adaptive 80%	DCL	Training validation	82/93	88.2%	80.1%-93.3%	34.0 days
Adaptive 80%	DCL	2021-2026 eval.	109/127	85.8%	78.7%-90.8%	32.8 days
Fixed clock	WCL	2021-2026 eval.	15/27	55.6%	37.3%-72.4%	4.0 weeks
Adaptive 80%	WCL	Training validation	14/14	100.0%	78.5%-100.0%	13.5 weeks
Adaptive 80%	WCL	2021-2026 eval.	27/27	100.0%	87.5%-100.0%	11.7 weeks

Note. Every forecast row uses prior completed cycles only; the retrospective protocol was externally preregistered. Precision requires both the low and its confirmation to meet the interval rule.

3.2 The result replicates across assets

H3 is supported by point estimates. Bitcoin, S&P 500 futures, and gold each exceed 80% for DCL, while all observed WCL windows in the chronological evaluation are hits. The confidence intervals are materially wider at asset level, especially for WCL, so cross-asset replication should be interpreted as a point-estimate result rather than a separate 95%-confidence proof for every market.

Exact asset-level estimates are detailed in Table 4, and the same estimates with Wilson intervals are visualized in Figure 2. Bitcoin has the highest adaptive DCL point estimate at 90.6%, followed by S&P 500 futures at 84.8% and gold at 83.7%. All three WCL point estimates are 100.0%, but the intervals in Figure 2 show that this equality in observed rates does not imply equal certainty.

Table 4. Asset-level confirmation-aware precision, Wilson confidence intervals, and mean window widths for the adaptive specification.

Asset	DCL hits	DCL precision	DCL 95% CI	DCL width	WCL hits	WCL precision	WCL 95% CI	WCL width
Bitcoin	29/32	90.6%	75.8%-96.8%	28.2 d	8/8	100.0%	67.6%-100.0%	11.2 w
S&P 500 futures	39/46	84.8%	71.8%-92.4%	38.3 d	10/10	100.0%	72.2%-100.0%	13.1 w
Gold	41/49	83.7%	71.0%-91.5%	30.5 d	9/9	100.0%	70.1%-100.0%	10.4 w

Note. d = calendar days; w = Monday-start weeks. WCL sample sizes are small and confidence intervals remain wide despite perfect point estimates.

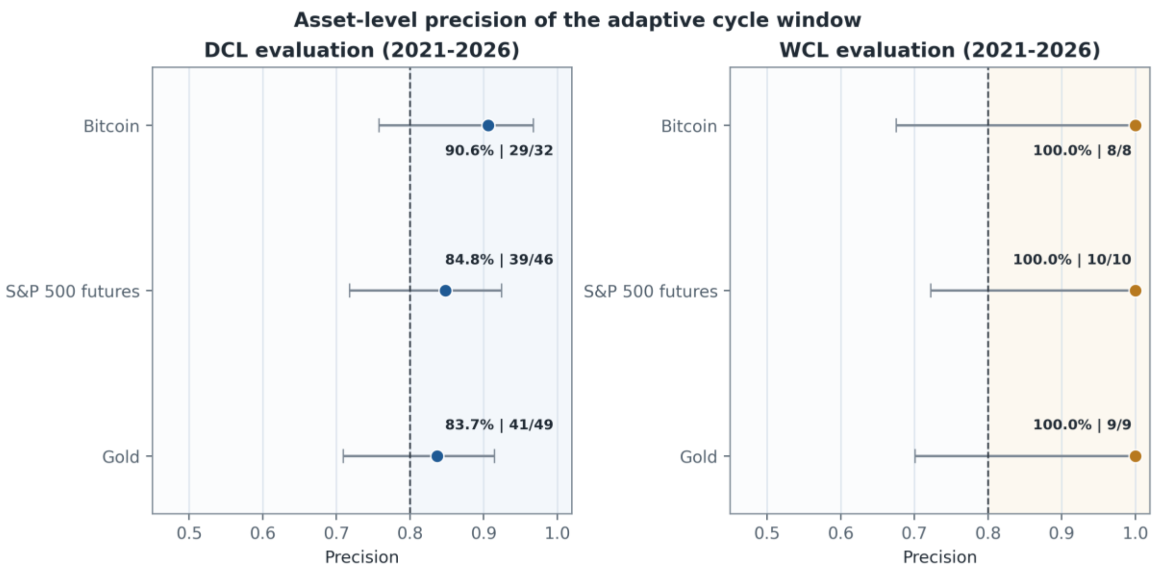

Figure 2. Asset-level precision of the adaptive cycle window in the 2021-2026 chronological evaluation. Points show estimates, horizontal bars show Wilson 95% confidence intervals, and labels show estimate and hit count.

3.3 Precision improves because uncertainty bands widen

H4 is supported. Mean DCL width rises from 14.7 to 32.8 calendar days, a 2.22x expansion. Mean WCL width rises from 4.0 to 11.7 weeks, a 2.92x expansion. Thus, the method reaches the target through better calibrated uncertainty, not through a more exact narrow turning-date estimate.

The static pre-2021 quantile interval obtains only 69.3% DCL precision, compared with 85.8% for rolling adaptation. For WCL, static precision is 88.9%, compared with 100.0% adaptively. Rolling updating therefore matters, particularly for daily cycles where duration regimes shift more visibly.

The full precision-width frontier is presented in Figure 3. Each successive point widens the central empirical band; the upward DCL path makes the cost of additional coverage visible, whereas WCL precision reaches 100% at several adjacent widths. Exact values for the fixed, static, primary adaptive, and high-confidence specifications are detailed in Table 5.

The frontier also prevents a misleading winner-takes-all interpretation. Moving rightward in Figure 3 purchases additional coverage with calendar time, so two specifications can serve different decision problems even when one has a higher hit rate. The primary 10th-90th percentile band is retained as the main specification because it clears the prespecified point target while remaining materially narrower than the 5th-95th percentile alternative. The latter is reported as a high-confidence option because its DCL lower confidence bound clears 80%, not because it is universally preferable. In applications where monitoring is costly, a narrower band may dominate despite more misses; where an omitted cycle is costly, the broader band may be rational. Table 5 makes this trade-off auditable by placing precision, uncertainty, and mean width in the same comparison rather than ranking models on coverage alone.

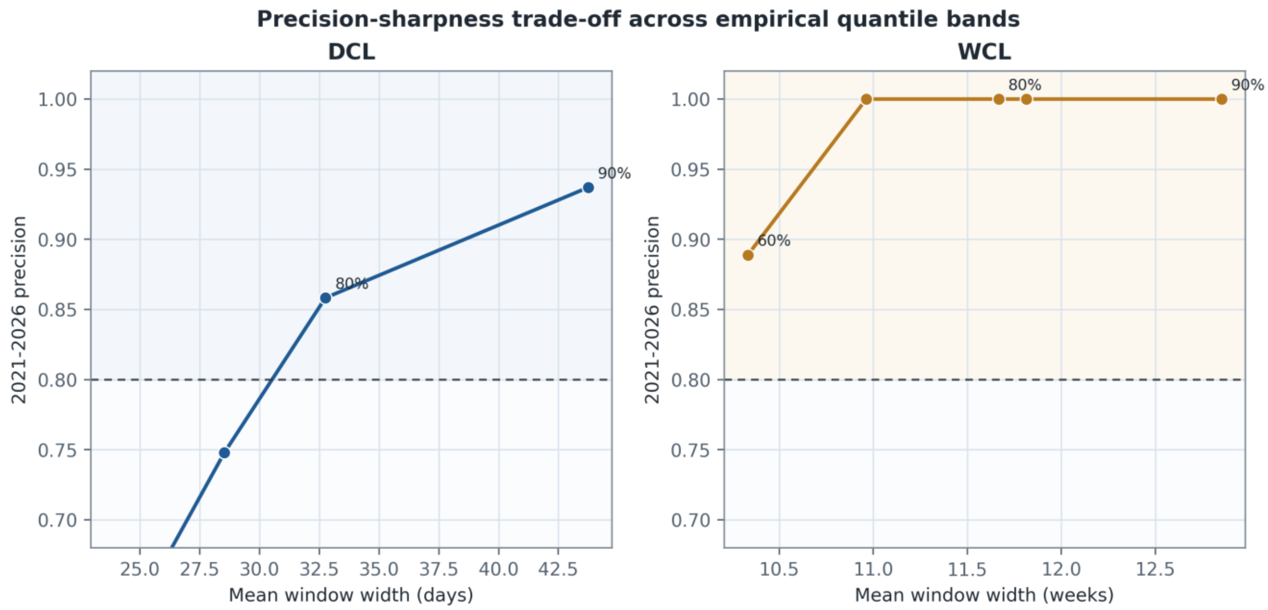

Figure 3. Precision-sharpness frontier in the 2021-2026 chronological evaluation across rolling empirical quantile bands. Percent labels denote the central empirical band; shaded regions mark precision at or above 80%.

Table 5. Precision–sharpness frontier: fixed, static, primary adaptive, and high-confidence specifications.

Method	Cycle	2021-2026 precision	95% CI	Mean width
Fixed clock	DCL	70.9%	62.4%-78.1%	14.7 days
Static training quantiles	DCL	69.3%	60.8%-76.6%	28.5 days
Rolling 10th-90th	DCL	85.8%	78.7%-90.8%	32.8 days
Rolling 5th-95th	DCL	93.7%	88.1%-96.8%	43.7 days
Fixed clock	WCL	55.6%	37.3%-72.4%	4.0 weeks
Static training quantiles	WCL	88.9%	71.9%-96.1%	11.0 weeks
Rolling 10th-90th	WCL	100.0%	87.5%-100.0%	11.7 weeks
Rolling 5th-95th	WCL	100.0%	87.5%-100.0%	12.9 weeks

Note. The 5th-95th percentile DCL band is the high-confidence specification; its lower Wilson bound exceeds 80%.

3.4 Sequential calibration remains above target by the end of the evaluation

The pooled DCL path temporarily falls below 80% during the early evaluation and then recovers as more completed cycles enter the rolling distribution. This behavior is consistent with regime adaptation rather than a one-time static fit. The WCL path records no miss in the 2021-2026 segment, but only 27 forecasts are available. The perfect WCL line should therefore be read together with its 87.5% lower pooled confidence bound and the much wider window.

The ordering of every realized forecast is retained in Figure 4. The DCL step path ends at 85.8% after 127 origins, while the WCL path remains at 100.0% through 27 origins. Showing the full path prevents the endpoint estimate from concealing the period in which cumulative DCL coverage was below the target.

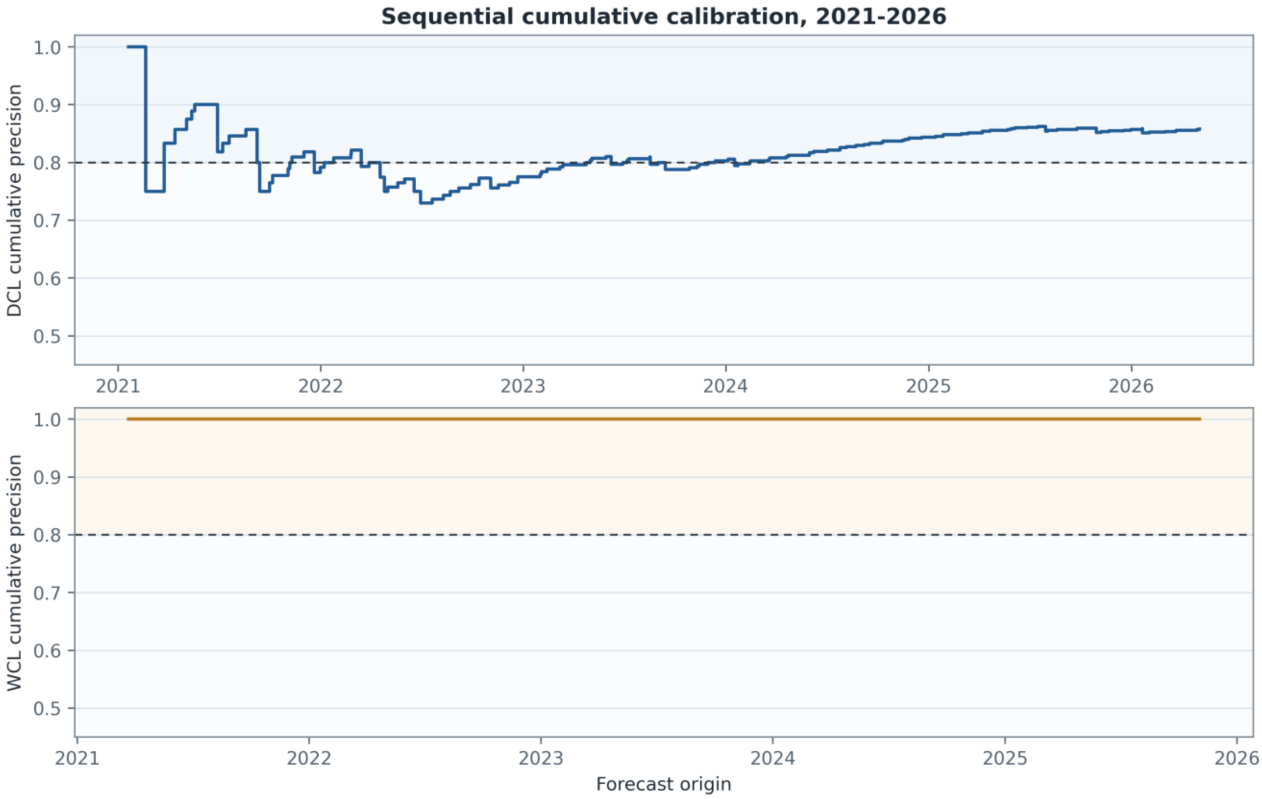

Figure 4. Cumulative confirmation-aware precision through the 2021-2026 chronological evaluation. The dashed line marks the 80% target; shaded regions identify target attainment.

3.5 High-confidence specification

H5 is supported at the pooled level. The wider 5th-95th percentile interval obtains 119/127 DCL hits (93.7%; 95% CI 88.1%-96.8%) and 27/27 WCL hits (100.0%; 95% CI 87.5%-100.0%). The DCL lower bound is 88.1%, above the 80% target. The cost is a mean width of 43.7 days for DCL and 12.9 weeks for WCL. This version is appropriate only when avoiding missed cycles is more important than date precision.

4. Discussion

4.1 What the 80% result establishes

The central result is positive but narrower than a claim of universal periodicity. Cycle Theory can produce an out-of-sample precision point estimate above 80% when the reference count is treated as a rolling duration distribution and the forecast endpoint explicitly allows for confirmation lag. The method is transparent, causal, and cross-asset: it uses the same lookback and quantile levels for all three markets and both timeframes.

The improvement also has a clear mechanism. Static DCL quantiles fail as market durations shift, while rolling quantiles adapt. Separating low-to-low duration from low-to-confirmation duration prevents recognition lag from turning a correctly timed low into a false success. The method therefore integrates the two objects that cycle practice often conflates: when the low occurs and when it becomes knowable.

The cross-asset result is consistent with evidence that efficiency and cycle anchors are not fixed properties. Bitcoin halving effects vary across returns, volatility, spillovers, currencies, and sentiment[19-24], while the broader efficiency literature documents alternating and state-dependent predictability[25-38]. A rolling distribution is therefore more defensible than a permanently fixed duration assumption, although the present study does not claim that timing adaptation alone explains market efficiency.

The primary DCL estimate of 85.8% is an endpoint of a sequential path, not a constant property of the full period. Figure 4 shows that cumulative coverage moved below and above the 80% reference before finishing above it. This temporal path is important: it demonstrates both the benefit and the limitation of delayed adaptation. The window updates only after a completed cycle becomes known, so abrupt duration shifts can still generate early misses before the new regime is represented in the rolling history.

4.2 Why calibration and sharpness must be reported together

A high hit rate is not sufficient. An interval covering most of the calendar would be perfectly calibrated and practically useless. The technical-timing literature also shows why apparent success must be read together with data-snooping risk, transaction costs, and regime dependence[39-47]. The primary DCL interval is more than twice the fixed width, and the primary WCL interval is almost three times wider. The 80% result should be interpreted as a calibrated risk window for monitoring and confirmation, not as a pinpoint bottom forecast.

This distinction explains why the high-confidence 90% band is statistically stronger but operationally less precise. Its pooled DCL confidence interval lies fully above 80%, yet the average interval spans more than six weeks. An analyst should choose the band according to the decision loss: the primary band balances timing and misses, while the high-confidence band prioritizes avoiding omission.

The difference between point-estimate success and confidence-bound success must also remain explicit. The primary DCL point estimate exceeds 80%, but its 78.7% lower Wilson bound does not. Only the wider 5th-95th percentile DCL specification places the lower bound above the target, at 88.1%. Conversely, the WCL estimate is 100.0%, but its lower bound is 87.5% because 27 forecasts cannot support a claim of near-certain future coverage. These distinctions are shown numerically in Tables 3 and 5 and visually through the interval bars in Figures 1 and 2.

A complete evaluation should consequently preserve three layers of evidence. The first is marginal coverage: the proportion of realized outcomes contained in the interval. The second is sharpness: the amount of time during which the forecast remains active. The third is sequential stability: whether calibration persists or is concentrated in one favorable subperiod. Tables 3 and 5 report the first two layers, while Figure 4 exposes the third. Reporting only the terminal percentage would hide the early DCL undercoverage; reporting only average width would hide whether the interval actually contains the target; and reporting only a confidence interval would not identify when misses occurred. Their joint presentation is therefore substantive rather than cosmetic. It allows a reader to distinguish a well-calibrated adaptive process from an excessively wide interval or a method whose aggregate success depends on a short episode.

4.3 Practical implementation

A practical implementation can maintain two clocks. The fixed theory count remains the interpretable center prior. Around it, the rolling duration ledger generates a calibrated monitoring window. During the window, candidate status should still require price evidence; time alone does not create a low. After confirmation, the realized low-to-low and low-to-confirmation durations enter the rolling sample, and the next interval is recomputed.

The method is suitable for research alerts, chart annotation, and risk-budget staging. It does not by itself specify position size, entry price, invalidation, or exit. A separate economic study should test whether the broader monitoring window produces value after transaction costs and whether signal concentration can be improved using pre-window volatility, trend, or liquidity features without sacrificing causal forecast integrity.

Operationally, the primary band is the more balanced research default because it limits the mean DCL window to 32.8 days and the WCL window to 11.7 weeks. The high-confidence band is better viewed as a loss-sensitive alternative for users who place a higher cost on a missed cycle than on prolonged monitoring. Neither choice should be converted into a trading instruction without a separate decision rule, transaction-cost model, and prospective execution record.

4.4 Limitations

Seven limitations matter. First, the target-anchor ledger is a rule-based reconstruction generated by the frozen Cycle Theory state machine rather than independently adjudicated ground truth. The comparison is internally fair because fixed and adaptive forecasts share that ledger, but an alternative event definition could produce different target dates. A useful replication should therefore publish both the detector and a blinded manual-adjudication protocol, then quantify agreement before comparing forecasting methods.

Second, the analysis is retrospective despite external preregistration, so live prospective performance remains untested. Chronological replay protects each origin from direct future-duration leakage, but it cannot reproduce operational data revisions, delayed feeds, contract-roll decisions, or human intervention. Third, WCL inference is based on only 27 forecasts; perfect observed precision is not proof of future perfection. Fourth, the primary DCL confidence interval still includes values marginally below 80%, although the wider robustness band clears the threshold.

Fifth, rolling quantiles react only after a new regime has generated completed cycles. The method can therefore remain temporarily miscalibrated after an abrupt structural break, as the early DCL path in Figure 4 illustrates. Sixth, continuous futures and convenience data feeds can differ from executable contracts and venue-certified histories, particularly around rolls and illiquid sessions. Seventh, widening the interval increases the chance that timing becomes less decision-useful even when statistical coverage improves. The present results should accordingly be evaluated as forecast-window calibration, not as evidence of implementable abnormal returns.

4.5 Research agenda

The next decisive study should extend the preregistration into a live event ledger containing every candidate, confirmation, failure, reset, low date, confirmation date, rule version, and manual intervention. Forecast intervals should be scored with interval scores as well as binary coverage, and every candidate parameter set should be frozen before prospective evaluation. A hierarchical adaptive model may partially pool the three assets while retaining asset-specific regimes. On-chain variables, volatility-state indicators, liquidity measures, and post-ETF market-structure variables provide prespecified candidates for testing whether calibration can be retained with narrower windows[96-112].

Prospective validation should also define success before any live forecasts are observed. A minimum protocol would timestamp the data snapshot, anchor, interval start and end, candidate-low date, confirmation date, and any reset at each origin. It would retain every forecast, including intervals that expire without confirmation, and would publish a locked rule for contract rolls, missing sessions, and revised source data. The primary analysis would then reproduce the present confirmation-aware hit indicator and Wilson interval, supplemented by interval score, mean width, miss direction, and time-to-confirmation. Prespecified quarterly calibration checks could flag drift without authorizing retrospective parameter changes. If an update is triggered, the new rule should begin a separately labeled evaluation era so that evidence from the original and modified procedures is not pooled. This design would turn the current chronological reconstruction into a genuinely prospective audit trail and would make successful replication more persuasive than another optimized historical backtest.

5. Conclusion

This study shows that Cycle Theory can exceed an 80% chronological pseudo-out-of-sample point target when fixed counts are converted into adaptive, confirmation-aware quantile intervals. The primary rolling specification reaches 85.8% for DCL and 100.0% for WCL in the 2021-2026 evaluation, compared with 70.9% and 55.6% under fixed clocks. The result replicates across Bitcoin, S&P 500 futures, and gold at the point-estimate level, but live prospective replication is still required.

The finding is inseparable from sharpness. Average widths expand to 32.8 days and 11.7 weeks. A wider robustness band raises DCL precision to 93.7% with a lower confidence bound above 80%, but it is less precise in calendar time. The defensible conclusion is therefore not that markets obey a narrow universal clock. It is that Cycle Theory becomes empirically calibratable when duration is modeled as a rolling distribution and confirmation delay is part of the forecast target.

Declarations

Funding. This research received no external funding.

Competing interests. The authors declare no competing interests.

Ethics approval. Not applicable; the study uses public market data and involves no human participants or animals.

Data availability. The analysis uses public market histories frozen on 14 July 2026. Source manifests, derived cycle ledgers, forecast tables, and validation notes are retained with the research files. Redistribution of raw data remains subject to provider terms.

Code availability. The complete causal reconstruction and adaptive-precision scripts are retained in the research package and can be supplied with a journal submission.

CRediT author statement. Muhammad Faiz: conceptualization, methodology, software, formal analysis, investigation, data curation, visualization, writing — original draft. Sonia Vernanda: methodology, validation, formal analysis, investigation, writing — review and editing. Both authors have reviewed and approved the final version of the manuscript.

Declaration of generative AI and AI-assisted technologies. The authors declare that no generative artificial intelligence or AI-assisted technology was used in the conception, analysis, or writing of this manuscript. The authors reviewed the methods, data, results, interpretations, and text, and accept responsibility for the publication.

References

1. Santa-Clara P, Valkanov R. The presidential puzzle: political cycles and the stock market. Journal of Finance. 2003;58(5):1841-1872. doi:10.1111/1540-6261.00590
2. Pástor Ľ, Veronesi P. Political cycles and stock returns. Journal of Political Economy. 2020;128(11):4011-4045. doi:10.1086/710532
3. Campbell JY, Shiller RJ. Valuation ratios and the long-run stock market outlook. Journal of Portfolio Management. 1998;24(2):11-26. doi:10.3905/jpm.24.2.11
4. Baur DG. The autumn effect of gold. Research in International Business and Finance. 2013;27(1):1-11. doi:10.1016/j.ribaf.2012.05.001
5. Erb CB, Harvey CR. The golden dilemma. Financial Analysts Journal. 2013;69(4):10-42. doi:10.2469/faj.v69.n4.1
6. Fabus J, Kremenova I, Stalmasekova N, et al. An empirical examination of Bitcoin's halving effects: assessing cryptocurrency sustainability within the landscape of financial technologies. Journal of Risk and Financial Management. 2024;17(6):229. doi:10.3390/jrfm17060229
7. Fama EF. Efficient capital markets: a review of theory and empirical work. Journal of Finance. 1970;25(2):383-417. doi:10.2307/2325486

8. Lo AW. The adaptive markets hypothesis. Journal of Portfolio Management. 2004;30(5):15-29. doi:10.3905/jpm.2004.442611
9. Shleifer A, Vishny RW. The limits of arbitrage. Journal of Finance. 1997;52(1):35-55. doi:10.1111/j.1540-6261.1997.tb03807.x
10. Kondratieff ND, Stolper WF. The long waves in economic life. Review of Economics and Statistics. 1935;17(6):105-115. doi:10.2307/1928486
11. Kahneman D, Tversky A. Prospect theory: an analysis of decision under risk. Econometrica. 1979;47(2):263-291. doi:10.2307/1914185
12. Shiller RJ. From efficient markets theory to behavioral finance. Journal of Economic Perspectives. 2003;17(1):83-104. doi:10.1257/089533003321164967
13. Barberis N, Shleifer A, Vishny R. A model of investor sentiment. Journal of Financial Economics. 1998;49(3):307-343. doi:10.1016/S0304-405X(98)00027-0
14. O'Connor FA, Lucey BM, Batten JA, et al. The financial economics of gold—a survey. International Review of Financial Analysis. 2015;41:186-205. doi:10.1016/j.irfa.2015.07.005
15. Baethge C, Goldbeck-Wood S, Mertens S. SANRA—a scale for the quality assessment of narrative review articles. Research Integrity and Peer Review. 2019;4(1):5. doi:10.1186/s41073-019-0064-8
16. Page MJ, McKenzie JE, Bossuyt PM, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ. 2021;372:n71. doi:10.1136/bmj.n71
17. Jegadeesh N, Titman S. Returns to buying winners and selling losers: implications for stock market efficiency. Journal of Finance. 1993;48(1):65-91. doi:10.1111/j.1540-6261.1993.tb04702.x
18. Johansen A, Ledoit O, Sornette D. Crashes as critical points. International Journal of Theoretical and Applied Finance. 2000;3(2):219-255. doi:10.1142/S0219024900000115
19. Veloso V, Gatsios RC, Magnani VM, et al. Is Bitcoin's market maturing? Cumulative abnormal returns and volatility in the 2024 halving and past cycles. Journal of Risk and Financial Management. 2025;18(5):242. doi:10.3390/jrfm18050242
20. Lashkaripour M. Some stylized facts about bitcoin halving. Finance Research Letters. 2024;69:106198. doi:10.1016/j.frl.2024.106198
21. Jiménez I, Mora-Valencia A, Perote J. Bitcoin halving and the integration of cryptocurrency and forex markets: an analysis of the higher-order moment spillovers. International Review of Economics & Finance. 2024;92:302-315. doi:10.1016/j.iref.2024.02.022
22. Liu W, Zhao X, Li W, et al. The effect of the cryptocurrency halving event. Pacific-Basin Finance Journal. 2025;94:102913. doi:10.1016/j.pacfin.2025.102913
23. M'bakob GB. Bubbles in Bitcoin and Ethereum: the role of halving in the formation of super cycles. Sustainable Futures. 2024;7:100178. doi:10.1016/j.sftr.2024.100178
24. Subawa NS, Mimaki CA, Mahendra IMO, et al. Bitcoin halving: how effective is it in driving cryptocurrency market dynamics? Journal of Risk and Financial Management. 2026;19(1):2. doi:10.3390/jrfm19010002
25. Urquhart A, Hudson R. Efficient or adaptive markets? Evidence from major stock markets using very long run historic data. International Review of Financial Analysis. 2013;28:130-142. doi:10.1016/j.irfa.2013.03.005
26. Kim JH, Shamsuddin A, Lim KP. Stock return predictability and the adaptive markets hypothesis: evidence from century-long U.S. data. Journal of Empirical Finance. 2011;18(5):868-879. doi:10.1016/j.jempfin.2011.08.002
27. Henkel SJ, Martin JS, Nardari F. Time-varying short-horizon predictability. Journal of Financial Economics. 2011;99(3):560-580. doi:10.1016/j.jfineco.2010.09.008
28. Chordia T, Shivakumar L. Momentum, business cycle, and time-varying expected returns. Journal of Finance. 2002;57(2):985-1019. doi:10.1111/1540-6261.00449
29. Urquhart A. The inefficiency of Bitcoin. Economics Letters. 2016;148:80-82. doi:10.1016/j.econlet.2016.09.019
30. Nadarajah S, Chu J. On the inefficiency of Bitcoin. Economics Letters. 2017;150:6-9. doi:10.1016/j.econlet.2016.10.033
31. Bariviera AF. The inefficiency of Bitcoin revisited: a dynamic approach. Economics Letters. 2017;161:1-4. doi:10.1016/j.econlet.2017.09.013
32. Tiwari AK, Jana RK, Das D, et al. Informational efficiency of Bitcoin—an extension. Economics Letters. 2018;163:106-109. doi:10.1016/j.econlet.2017.12.006
33. Khuntia S, Pattanayak JK. Adaptive market hypothesis and evolving predictability of bitcoin. Economics Letters. 2018;167:26-28. doi:10.1016/j.econlet.2018.03.005
34. Khursheed A, Naeem M, Ahmed S, et al. Adaptive market hypothesis: an empirical analysis of time-varying market efficiency of cryptocurrencies. Cogent Economics & Finance. 2020;8(1):1719574. doi:10.1080/23322039.2020.1719574
35. Chu J, Zhang Y, Chan S. The adaptive market hypothesis in the high frequency cryptocurrency market. International Review of Financial Analysis. 2019;64:221-231. doi:10.1016/j.irfa.2019.05.008
36. Jena SK, Tiwari AK, Doğan B, et al. Are the top six cryptocurrencies efficient? Evidence from time-varying long memory. International Journal of Finance & Economics. 2022;27(3):3730-3740. doi:10.1002/ijfe.2347
37. Mokni K, El Montasser G, Ajmi AN, et al. On the efficiency and its drivers in the cryptocurrency market: the case of Bitcoin and Ethereum. Financial Innovation. 2024;10(1):39. doi:10.1186/s40854-023-00566-3
38. Kristoufek L. What are the main drivers of the Bitcoin price? Evidence from wavelet coherence analysis. PLoS ONE. 2015;10(4):e0123923. doi:10.1371/journal.pone.0123923
39. Brock W, Lakonishok J, LeBaron B. Simple technical trading rules and the stochastic properties of stock returns. Journal of Finance. 1992;47(5):1731-1764. doi:10.1111/j.1540-6261.1992.tb04681.x
40. Sullivan R, Timmermann A, White H. Data-snooping, technical trading rule performance, and the bootstrap. Journal of Finance. 1999;54(5):1647-1691. doi:10.1111/0022-1082.00163
41. Neely CJ, Rapach DE, Tu J, et al. Forecasting the equity risk premium: the role of technical indicators. Management Science. 2014;60(7):1772-1791. doi:10.1287/mnsc.2013.1838
42. Hudson R, Urquhart A. Technical trading and cryptocurrencies. Annals of Operations Research. 2021;297(1-2):191-220. doi:10.1007/s10479-019-03357-1
43. Grobys K, Ahmed S, Sapkota N. Technical trading rules in the cryptocurrency market. Finance Research Letters. 2020;32:101396. doi:10.1016/j.frl.2019.101396
44. Detzel A, Liu H, Strauss J, et al. Learning and predictability via technical analysis: evidence from bitcoin and stocks with hard-to-value fundamentals. Financial Management. 2021;50(1):107-137. doi:10.1111/fima.12310
45. Masuku KL, Gopane TJ. Technical trading rules' profitability and dynamic risk premiums of cryptocurrency exchange rates. Journal of Capital Markets Studies. 2022;6(1):6-32. doi:10.1108/JCMS-10-2021-0030
46. Bazán-Palomino W, Svogun D. On the drivers of technical analysis profits in cryptocurrency markets: a distributed lag approach. International Review of Financial Analysis. 2023;86:102516. doi:10.1016/j.irfa.2023.102516
47. Svogun D, Bazán-Palomino W. Technical analysis in cryptocurrency markets: do transaction costs and bubbles matter? Journal of International Financial Markets, Institutions and Money. 2022;79:101601. doi:10.1016/j.intfin.2022.101601

48. Bouman S, Jacobsen B. The Halloween indicator, 'sell in May and go away': another puzzle. American Economic Review. 2002;92(5):1618-1635. doi:10.1257/000282802762024683
49. Zhang CY, Jacobsen B. The Halloween indicator, 'sell in May and go away': everywhere and all the time. Journal of International Money and Finance. 2021;110:102268. doi:10.1016/j.jimonfin.2020.102268
50. Keim DB. Size-related anomalies and stock return seasonality. Journal of Financial Economics. 1983;12(1):13-32. doi:10.1016/0304-405X(83)90025-9
51. French KR. Stock returns and the weekend effect. Journal of Financial Economics. 1980;8(1):55-69. doi:10.1016/0304-405X(80)90021-5
52. Kamstra MJ, Kramer LA, Levi MD. Winter blues: a SAD stock market cycle. American Economic Review. 2003;93(1):324-343. doi:10.1257/000282803321455322
53. Caporale GM, Plastun A. The day of the week effect in the cryptocurrency market. Finance Research Letters. 2019;31:258-269. doi:10.1016/j.frl.2018.11.012
54. Ma D, Tanizaki H. The day-of-the-week effect on Bitcoin return and volatility. Research in International Business and Finance. 2019;49:127-136. doi:10.1016/j.ribaf.2019.02.003
55. Aharon DY, Qadan M. Bitcoin and the day-of-the-week effect. Finance Research Letters. 2019;31:415-424. doi:10.1016/j.frl.2018.12.004
56. Kaiser L. Seasonality in cryptocurrencies. Finance Research Letters. 2019;31:232-238. doi:10.1016/j.frl.2018.11.007
57. Baur DG, Cahill D, Godfrey K, et al. Bitcoin time-of-day, day-of-week and month-of-year effects in returns and trading volume. Finance Research Letters. 2019;31:78-92. doi:10.1016/j.frl.2019.04.023
58. Kinateder H, Papavassiliou VG. Calendar effects in Bitcoin returns and volatility. Finance Research Letters. 2021;38:101420. doi:10.1016/j.frl.2019.101420
59. Sahu S, Fonseca Ramírez A, Kim JM. Exploring calendar anomalies and volatility dynamics in cryptocurrencies: a comparative analysis of day-of-the-week effects before and during the COVID-19 pandemic. Journal of Risk and Financial Management. 2024;17(8):351. doi:10.3390/jrfm17080351
60. Moskowitz TJ, Ooi YH, Pedersen LH. Time series momentum. Journal of Financial Economics. 2012;104(2):228-250. doi:10.1016/j.jfineco.2011.11.003
61. Asness CS, Moskowitz TJ, Pedersen LH. Value and momentum everywhere. Journal of Finance. 2013;68(3):929-985. doi:10.1111/jofi.12021
62. Cooper MJ, Gutierrez RC, Hameed A. Market states and momentum. Journal of Finance. 2004;59(3):1345-1365. doi:10.1111/j.1540-6261.2004.00665.x
63. Borgards O. Dynamic time series momentum of cryptocurrencies. North American Journal of Economics and Finance. 2021;57:101428. doi:10.1016/j.najef.2021.101428
64. Wen Z, Bouri E, Xu Y, et al. Intraday return predictability in the cryptocurrency markets: momentum, reversal, or both. North American Journal of Economics and Finance. 2022;62:101733. doi:10.1016/j.najef.2022.101733
65. Tan X, Tao Y. Trend-based forecast of cryptocurrency returns. Economic Modelling. 2023;124:106323. doi:10.1016/j.econmod.2023.106323
66. Jia B, Shen D, Zhang W. Extreme sentiment and herding: evidence from the cryptocurrency market. Research in International Business and Finance. 2022;63:101770. doi:10.1016/j.ribaf.2022.101770
67. Chen A-S, Nguyen HT. A new perspective on how investor sentiment affects herding behavior in the cryptocurrency market. Finance Research Letters. 2024;67:105737. doi:10.1016/j.frl.2024.105737
68. Bogdan S, Brmalj N, Mujačević E. Impact of liquidity and investors sentiment on herd behavior in cryptocurrency market. International Journal of Financial Studies. 2023;11(3):97. doi:10.3390/ijfs11030097
69. Wang J-N, Liu H-C, Lee Y-H, et al. FoMO in the Bitcoin market: revisiting and factors. Quarterly Review of Economics and Finance. 2023;89:244-253. doi:10.1016/j.qref.2023.04.007
70. Aslanidis N, Bariviera AF, López ÓG. The link between cryptocurrencies and Google Trends attention. Finance Research Letters. 2022;47:102654. doi:10.1016/j.frl.2021.102654
71. Lin Z-Y. Investor attention and cryptocurrency performance. Finance Research Letters. 2021;40:101702. doi:10.1016/j.frl.2020.101702
72. Phillips PCB, Shi S, Yu J. Testing for multiple bubbles: historical episodes of exuberance and collapse in the S&P 500. International Economic Review. 2015;56(4):1043-1078. doi:10.1111/iere.12132
73. Wheatley S, Sornette D, Huber T, et al. Are Bitcoin bubbles predictable? Combining a generalized Metcalfe's Law and the Log-Periodic Power Law Singularity model. Royal Society Open Science. 2019;6(6):180538. doi:10.1098/rsos.180538
74. Gerlach JC, Demos G, Sornette D. Dissection of Bitcoin's multiscale bubble history from January 2012 to February 2018. Royal Society Open Science. 2019;6(7):180643. doi:10.1098/rsos.180643
75. Shu M, Zhu W. Real-time prediction of Bitcoin bubble crashes. Physica A: Statistical Mechanics and its Applications. 2020;548:124477. doi:10.1016/j.physa.2020.124477
76. Yao CZ, Li HY. A study on the bursting point of Bitcoin based on the BSADF and LPPLS methods. North American Journal of Economics and Finance. 2021;55:101280. doi:10.1016/j.najef.2020.101280
77. Fry J. Booms, busts and heavy-tails: the story of Bitcoin and cryptocurrency markets? Economics Letters. 2018;171:225-229. doi:10.1016/j.econlet.2018.08.008
78. Geuder J, Kinateder H, Wagner NF. Cryptocurrencies as financial bubbles: the case of Bitcoin. Finance Research Letters. 2019;31:179-184. doi:10.1016/j.frl.2018.11.011
79. Chaim P, Laurini MP. Is Bitcoin a bubble? Physica A: Statistical Mechanics and its Applications. 2019;517:222-232. doi:10.1016/j.physa.2018.11.031
80. Bouri E, Shahzad SJH, Roubaud D. Co-explosivity in the cryptocurrency market. Finance Research Letters. 2019;29:178-183. doi:10.1016/j.frl.2018.07.005
81. Podhorsky A. Bursting the bitcoin bubble: do market prices reflect fundamental bitcoin value? International Review of Financial Analysis. 2024;93:103158. doi:10.1016/j.irfa.2024.103158
82. Karau S. Monetary policy and Bitcoin. Journal of International Money and Finance. 2023;137:102880. doi:10.1016/j.jimonfin.2023.102880
83. Ma C, Tian Y, Hsiao S, et al. Monetary policy shocks and Bitcoin prices. Research in International Business and Finance. 2022;62:101711. doi:10.1016/j.ribaf.2022.101711
84. Tosun TT, Uğurlu E. The impact of the Fed's monetary policy on cryptocurrencies: novel policy implications for central banks. Journal of Risk and Financial Management. 2025;18(7):393. doi:10.3390/jrfm18070393
85. Pietrzak M. What can monetary policy tell us about Bitcoin? Annals of Finance. 2023;19(4):545-559. doi:10.1007/s10436-023-00432-3
86. Hodula M. Does U.S. monetary policy sway global crypto investment demand? Finance Research Letters. 2025;80:107408. doi:10.1016/j.frl.2025.107408
87. Lin M, Liu Y, Sheng VNK. Analysis of the impact of macroeconomic factors on cryptocurrency returns—based on quantile regression study. International Review of Economics & Finance. 2025;97:103757. doi:10.1016/j.iref.2024.103757

88. Ahmed WMA. On the robust drivers of cryptocurrency liquidity: the case of Bitcoin. Financial Innovation. 2024;10(1):69. doi:10.1186/s40854-023-00598-9
89. Baur DG, Lucey BM. Is gold a hedge or a safe haven? An analysis of stocks, bonds and gold. Financial Review. 2010;45(2):217-229. doi:10.1111/j.1540-6288.2010.00244.x
90. Baur DG, McDermott TK. Is gold a safe haven? International evidence. Journal of Banking & Finance. 2010;34(8):1886-1898. doi:10.1016/j.jbankfin.2009.12.008
91. Beckmann J, Berger T, Czudaj R. Does gold act as a hedge or a safe haven for stocks? A smooth transition approach. Economic Modelling. 2015;48:16-24. doi:10.1016/j.econmod.2014.10.044
92. Bouri E, Molnár P, Azzi G, et al. On the hedge and safe haven properties of Bitcoin: is it really more than a diversifier? Finance Research Letters. 2017;20:192-198. doi:10.1016/j.frl.2016.09.025
93. Conlon T, McGee R. Safe haven or risky hazard? Bitcoin during the Covid-19 bear market. Finance Research Letters. 2020;35:101607. doi:10.1016/j.frl.2020.101607
94. Ji Q, Bouri E, Lau CKM, et al. Dynamic connectedness and integration in cryptocurrency markets. International Review of Financial Analysis. 2019;63:257-272. doi:10.1016/j.irfa.2018.12.002
95. Corbet S, Meegan A, Larkin C, et al. Exploring the dynamic relationships between cryptocurrencies and other financial assets. Economics Letters. 2018;165:28-34. doi:10.1016/j.econlet.2018.01.004
96. Grobys K, Näsman S, Sandretto D. Using on-chain data to predict Bitcoin cycles. Research in International Business and Finance. 2026;89:103486. doi:10.1016/j.ribaf.2026.103486
97. Kim G, Shin DH, Choi JG, et al. A deep learning-based cryptocurrency price prediction model that uses on-chain data. IEEE Access. 2022;10:56232-56248. doi:10.1109/ACCESS.2022.3177888
98. King JC, Dale R, Amigó JM. Blockchain metrics and indicators in cryptocurrency trading. Chaos, Solitons & Fractals. 2024;178:114305. doi:10.1016/j.chaos.2023.114305
99. Sakkas A, Urquhart A. Blockchain factors. Journal of International Financial Markets, Institutions and Money. 2024;94:102012. doi:10.1016/j.intfin.2024.102012
100. Morillon TG, Chacon RG. Dissecting the stock to flow model for Bitcoin. Studies in Economics and Finance. 2022;39(3):506-523. doi:10.1108/SEF-10-2021-0409
101. Shelton A. Bitcoin return prediction: is it possible via stock-to-flow, Metcalfe's Law, technical analysis, or market sentiment? Journal of Risk and Financial Management. 2024;17(10):443. doi:10.3390/jrfm17100443
102. Alabi K. Digital blockchain networks appear to be following Metcalfe's Law. Electronic Commerce Research and Applications. 2017;24:23-29. doi:10.1016/j.elerap.2017.06.003
103. Chkili W. Modeling Bitcoin price volatility: long memory vs Markov switching. Eurasian Economic Review. 2021;11(3):433-448. doi:10.1007/s40822-021-00180-7
104. Ardia D, Bluteau K, Rüede M. Regime changes in Bitcoin GARCH volatility dynamics. Finance Research Letters. 2019;29:266-271. doi:10.1016/j.frl.2018.08.009
105. Shakourloo A, Azimli A. Regime-switching in bitcoin volatility under global uncertainty: Markov-switching GARCH and hidden Markov Copula approaches. Research in International Business and Finance. 2026;83:103295. doi:10.1016/j.ribaf.2026.103295
106. Shen D, Urquhart A, Wang P. Forecasting the volatility of Bitcoin: the importance of jumps and structural breaks. European Financial Management. 2020;26(5):1294-1323. doi:10.1111/eufm.12254
107. Mensi W, Al-Yahyaee KH, Kang SH. Structural breaks and double long memory of cryptocurrency prices: a comparative analysis from Bitcoin and Ethereum. Finance Research Letters. 2019;29:222-230. doi:10.1016/j.frl.2018.07.011
108. Babalos V, Bouri E, Gupta R. Does the introduction of US spot Bitcoin ETFs affect spot returns and volatility of major cryptocurrencies? Quarterly Review of Economics and Finance. 2025;102:102006. doi:10.1016/j.qref.2025.102006
109. Kia K, Liu B, Li Q, et al. Price discovery in Bitcoin ETF market. Financial Review. 2026;61(2):435-449. doi:10.1111/fire.70026
110. Mohamad A. Do Bitcoin ETFs lead price discovery following their introduction in the Bitcoin market? Computational Economics. 2025;66(1):947-969. doi:10.1007/s10614-025-10998-x
111. Frino A, Gaudiosi R, Webb RI, et al. Price discovery in Bitcoin spot or futures? The jury is out. Journal of Futures Markets. 2025;45(4):269-288. doi:10.1002/fut.22560
112. Oefele N. One year of bitcoin spot ETPs: a brief market and fund flow analysis. Economics Letters. 2025;250:112304. doi:10.1016/j.econlet.2025.112304