QUANTUM RESERVE PRESS · CRYPTOSOWHAT STRUCTURAL SERIES · WP 2026-09

The Cycle Is Real. The Number Is Not.

Bitcoin's 1,064-Day Pattern, and the Family of Approximate Regularities It Belongs To

Gregory S. Carmichael, DBA

First Edition · August 2026

A convincing scene assembled entirely from real elements that never occurred together. The resemblance to a cycle constant is the subject of this paper.

ABSTRACT. A widely circulated claim holds that Bitcoin tops are separated by roughly 1,064 days, with subordinate relationships at 512 and 256 days. This paper measures the claim rather than repeating it. The headline figure is misattributed: 1,064 days is the interval from a cycle *low* to the following top, not the interval between tops, and the three observed values are 1,068, 1,061, and 1,050 days. The interval between tops is 1,424 and 1,426 days across the two modern cycles. The 512-day figure is the interval from cycle low to the subsequent halving and fits to within four days twice. The 256-day figure has no referent in any construction reviewed here. The paper's central negative result is that these quantities are not independent: the top-to-top interval is identically equal to the halving interval plus the change in the halving-to-top lag, so no comparison among them can establish that one drives another. Three approximate regularities coexist—low-to-top near 1,060 days, halving-to-top near 536 days, and top-to-top near 1,425 days—and with three cycles the data cannot identify which, if any, is structural. The halving-to-top relationship, which the circulated framework omits, in fact produced the best out-of-sample call on the October 2025 peak. A preregistered test of the decline leg is specified in Section 8 with a scoring date of 31 March 2027, and a test that discriminates between the competing regularities resolves at the next cycle peak.

THE DINNER TABLE VERSION

Somebody at the table says Bitcoin tops every 1,064 days, and it sounds impressive because it is oddly specific. Odd specificity is what should make you check it. When you do, two things turn up. First, the number is measuring the climb from the last crash bottom to the next peak, not the gap between peaks. Second—and this is the part almost nobody gets to—there are three different "constants" floating around in this data, and they are not independent facts. They are three ways of describing the same handful of dates. Picking your favourite and calling it the real one is a preference, not a finding. With three cycles on the board, that is all anyone has.

1. The Claim as It Circulates

The circulated framework states that Bitcoin advances for approximately 1,064 days and declines for approximately 364 days. Its provenance is worth stating plainly: the model was posted anonymously in December 2023 and projected a cycle top on 6 October 2025.[1] Bitcoin recorded its all-time high of approximately $126,198 during that week. The forecast was made twenty-two months in advance and landed within days.

That record is striking and this paper does not dismiss it. But an accurate forecast is not a validated model, and the version now circulating contains a correctable error of attribution.

1.1 The error

The claim is routinely rendered as "1,064 days between tops." It is not. Every construction reviewed for this paper measures from a cycle *low* to the following top. The distinction changes the number by more than a year.

[1] CoinMarketCap Academy, "What's Really Going On With Bitcoin Price and the BTC 4-Year Cycle?" (2026). The archived original post has not been independently verified for this paper; the provenance rests on secondary reporting and should be treated accordingly.

Interval	Dates	Days
Top → Top	30 Nov 2013 → 17 Dec 2017	1,478
Top → Top	17 Dec 2017 → 10 Nov 2021	1,424
Top → Top	10 Nov 2021 → 6 Oct 2025	1,426
Low → Top	14 Jan 2015 → 17 Dec 2017	1,068
Low → Top	15 Dec 2018 → 10 Nov 2021	1,061
Low → Top	21 Nov 2022 → 6 Oct 2025	1,050

Table 1: The 1,064-day figure describes the lower block, not the upper one. Nothing in the dataset is 1,064 days between tops.

This is the paper's most durable and least interesting contribution. It is also the one most likely to be needed in a meeting.

2. The Measurement Ledger

2.1 Anchors and the assumption they carry

Date	Type	Level	Note
28 Nov 2012	Halving	—	Block 210,000
30 Nov 2013	Top	≈$1,163	Mt. Gox era; impaired tape
14 Jan 2015	Low	≈$152–166	Venue-dependent
9 Jul 2016	Halving	—	Block 420,000
17 Dec 2017	Top	≈$19,783	Some sources date 16 Dec
15 Dec 2018	Low	≈$3,122	
11 May 2020	Halving	—	Block 630,000
10 Nov 2021	Top	≈$68,990	
21 Nov 2022	Low	≈$15,476	
20 Apr 2024	Halving	—	Block 840,000 (UTC)
6 Oct 2025	Top	≈$126,198	Current all-time high

Table 2: Anchor set. Price anchors are intraday extremes on major USD spot venues.

CAVEAT

Material assumption, stated once and carried throughout. Every price anchor is an intraday extreme on a major USD spot venue. Substitute daily closes, or a different exchange, and each anchor moves by one to five days. Section 7 quantifies the effect.

Alternative datings circulating in secondary commentary are not admitted unless the venue and quote are documented. A 24 November 2022 low appears in several such sources; mainstream daily histories place the 2022 low on 21 November, and 24 November trades materially above it. It is not used here.

2.2 Every leg, measured

Leg	Cycle 1	Cycle 2	Cycle 3
Low → Top (“the 1,064”)	1,068	1,061	1,050
Top → Low (“the 364”)	410	363	376
Low → Halving (“the 512”)	542	513	516
Halving → Top (L)	526	548	534
Top → Top	1,478	1,424	1,426
Halving → Halving	1,319	1,402	1,440

Table 3: The complete interval ledger.

2.3 Why raw spread is the wrong statistic

The obvious way to compare these legs is by raw day-spread, and it is invalid. Raw spread ignores both interval length and sample size: an 18-day spread on a 1,060-day leg and a 2-day spread on a 1,425-day leg are not comparable quantities, and the second is computed from two observations while the first uses three. Any ranking built on raw spread will favour whichever leg happens to be longest and thinnest-sampled.

Computed instead as a coefficient of variation:

	Full sample (n=3)		Modern only (n=2)	
Leg	Mean	CV	Mean	CV
Low → Top	1,059.7	**0.70%**	1,055.5	0.52%
Halving → Top	536.0	1.70%	541.0	1.29%
Top → Top	1,442.7	1.73%	1,425.0	**0.07%**
Halving interval	1,387.0	3.64%	—	—

Table 4: Coefficient of variation, population form. On the full sample the low-to-top leg is the *tightest* of the three—the leg most often dismissed as a trend masquerading as a constant.

Two conclusions follow, and both cut against the intuition that a longer, steadier-looking interval is the better measurement.

First, on the full three-cycle sample the 1,064-day leg is the best-behaved quantity in the dataset, not the worst. Its 18-day spread looks poor only when compared against a spread measured on a longer interval.

Second, the top-to-top figure attains its apparent precision only after the 2013 observation is removed, which reduces it to two points. A coefficient of variation computed on two observations is a single difference divided by a mean. It is a description, not evidence.

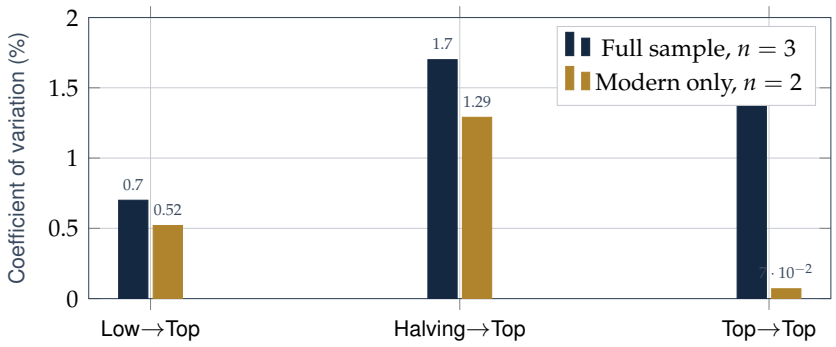

Figure 1: Which leg is "most stable" depends entirely on whether the 2013 cycle is admitted. That is a specification choice, not a property of Bitcoin.

2.4 The 2013 exclusion, and the standard it must meet

Bitcoin's 2013 peak printed on an exchange that was already impairing withdrawals and that halted trading on 24 February 2014—roughly twelve weeks later—before filing for bankruptcy in Japan that same week. Contemporary reporting documented the withdrawal problems and the resulting price premium on that venue. There is a real case for treating the observation as unreliable.

The case must be applied symmetrically. The 2015 low sits in the same era and is used in this paper without objection, because excluding it would destroy the 512-day relationship in Section 4. Any paper that excludes 2013 and retains 2015 is making a choice that helps its argument in both directions. The exclusion is therefore stated as a specification choice with a disclosed effect: **with 2013 included the top-to-top spread is 54 days, not 2**, and every claim built on the tighter figure is conditional on the exclusion.

3. Three Regularities, Not One Constant

3.1 The candidates

Three quantities in Table 3 are stable enough to be worth naming.

Regularity	Value	Observations	Status
Low → Top	≈1,060 d, declining ≈9 d/cycle	3	Circulated as "1,064"
Halving → Top	≈536 d, range 22 d	3	Largely omitted
Top → Top	≈1,425 d (modern)	2	Obscured by the above

Table 5: Three coexisting approximate regularities.

It is tempting to nominate one of these as the real constant and treat the others as derivative. On the evidence available, none of the three is identified as primary, and Section 6 explains why the data structurally cannot identify one.

3.2 The compensation, and what it does not prove

The top-to-top interval is the sum of the decline leg and the advance leg:

$$363 + 1{,}061 = 1{,}424 \qquad 376 + 1{,}050 = 1{,}426$$

The components moved in opposite directions and nearly cancelled. It is easy to read that as evidence of a protected total and to call the stability arithmetic. It is not. The identity guarantees only that the sum equals the parts; it says nothing about *why* the parts moved in offsetting directions by nearly equal amounts. That cancellation may be structural or it may be chance, and two observations cannot distinguish the cases.

THE DINNER TABLE VERSION

A commuter's drive takes the same total time most mornings even though the first half and second half vary a lot. You might conclude the total is somehow protected—traffic finds a way. Or the two halves might just have happened to offset twice. Watching two mornings does not tell you which. Watching two hundred would.

4. The 512-Day Rung

The 512-day figure measures from a cycle low to the *following halving*. Bitcoin Suisse identified it after the 2018 low, noting the bottom formed roughly 512 days ahead of the May 2020 halving.[2]

Low	Halving	Days	vs 512
14 Jan 2015	9 Jul 2016	542	+30
15 Dec 2018	11 May 2020	513	+1
21 Nov 2022	20 Apr 2024	516	+4

Table 6: The 512-day relationship: two close fits and one thirty-day miss.

One day and four days is a remarkable fit on two observations. The 2015 low missed by thirty. The usual defence—that early-cycle data is unreliable—is the same defence used to exclude the 2013 top in Section 2.4, and it cannot be deployed selectively. Either the pre-2016 tape is admissible or it is not. If it is admissible, the 512-day relationship holds twice out of three. If it is not, the relationship rests on $n = 2$ and so does the 1,425-day figure.

5. The 256-Day Problem

No construction reviewed for this paper assigns a leg to 256 days.[3] The figure travels with the framework without a referent, and that is the finding.

5.1 Reconstruction, correctly named

The reconstruction is often described as a "binary ladder," but the multipliers required to make the legs fit include 1.5 and 5.5, which are not powers of two. The hypothesis is properly described as a **256-day lattice**: cycle legs as integer or half-integer multiples of 256. A strictly binary variant (256×2^k) is tested separately below.

Tolerance must be fixed before classification. This paper adopts **1.0% mean absolute deviation relative to the hypothesised rung**, chosen because it corresponds to roughly five days on a 512-

[2] Bitcoin Suisse, "The Bitcoin Halving and the Bitcoin Market Cycle."

[3] This is a scoped negative: among the published and circulated constructions reviewed here, no explicit referent for 256 days was found. It is not a claim that none exists.

day leg—the paper's own stated anchor uncertainty. The threshold is disclosed as chosen after seeing the data; a genuine preregistration would fix it in advance.

Rung	Predicts	Days	Observed	MAD/rung	$< 1\%$?
2×256	Low $\rightarrow$ Halving	512	513, 516	0.49%	yes
5.5×256	Top $\rightarrow$ Top	1,408	1,424, 1,426	1.21%	no
4×256	Low $\rightarrow$ Top	1,024	1,068, 1,061, 1,050	3.48%	no
1.5×256	Top $\rightarrow$ Low	384	363, 376	3.78%	no
2×256	Halving $\rightarrow$ Top	512	526, 548, 534	4.69%	no

Table 7: The 256-day lattice tested rung by rung. MAD is mean absolute deviation from the hypothesised rung, expressed as a fraction of it.

CAVEAT

Threshold sensitivity, disclosed. At a 1.0% tolerance one rung passes. At 1.5% two pass. At 0.5% one passes marginally and the lattice has essentially no support. The conclusion "one rung fits" is therefore threshold-dependent across a plausible range, and the honest summary is weaker: *at most one of five rungs fits at any tolerance tight enough to be meaningful, and the same rung length of 512 days produces both the best and the worst fit in the table depending on which leg it is applied to.*

Under the strictly binary variant, only two legs have candidate rungs at all—512 for low-to-halving and 1,024 for low-to-top. The first fits; the second misses by 26 to 44 days. A two-rung ladder with one hit is not a structure.

5.2 The one forward-looking argument

The low-to-top leg is compressing at a steady rate:

$$1{,}068 \longrightarrow 1{,}061 \longrightarrow 1{,}050 \qquad \text{OLS slope} = -9.0 \text{ days per cycle}$$

At that rate it reaches 1,024 days in approximately 2.9 more cycles. The binary reading is not correct now but is approaching correctness monotonically.

CAVEAT

Fitting a slope to three points and extrapolating three cycles has no statistical power. Further, the usual mechanism offered—deeper liquidity, institutional participation, spot ETFs, faster price discovery—cannot explain a three-cycle monotone sequence, because U.S. spot Bitcoin ETFs launched in January 2024 and are present only in the final cycle. The compression is an observation. The explanation is speculation, and is labelled as such.

6. The Identity That Binds Them

A natural reading of Table 3 is that Bitcoin runs on two clocks—a protocol clock set by halvings and a price clock set by market behaviour—and that the two have begun to decouple. That reading is unavailable, and the reason it is unavailable is the most useful result in this paper.

6.1 The identity

Let H_i be a halving date, T_i the following top, and $L_i = T_i - H_i$ the halving-to-top lag. Then, as a matter of definition:

$$\underbrace{T_{i+1} - T_i}_{\text{top-to-top}} = \underbrace{(H_{i+1} - H_i)}_{\text{halving interval}} + \underbrace{(L_{i+1} - L_i)}_{\Delta L} \tag{1}$$

Substituting the measured values:

$$1{,}424 = 1{,}402 + (548 - 526) \qquad 1{,}426 = 1{,}440 + (534 - 548)$$

It is tempting to tabulate the differences between the halving interval and the top-to-top interval—+22 and −14 days—and read them as evidence that two independent clocks are separating. Those figures are exactly ΔL. They are an algebraic restatement of the phase relationship already contained in the data, not independent evidence of anything. No comparison among these three quantities can establish that one drives another, because the third is defined by the other two.

6.2 What the identity does establish

Equation (1) is not empty. It imposes a binding constraint: **the three regularities cannot all be stable at once unless they are mutually consistent**, and the halving interval is demonstrably not stable—1,319, 1,402, 1,440 days, a 121-day range and a CV of 3.64%.

It follows that top-to-top stability and halving-to-top stability are in direct competition. As the halving interval drifts upward, either the top-to-top interval must drift with it, or L must shorten to absorb the drift. Both cannot hold. Over the last cycle the halving interval rose 38 days while top-to-top rose 2, which means L absorbed 36 days of it. That is the opposite of decoupling. It is L acting as the accommodating variable.

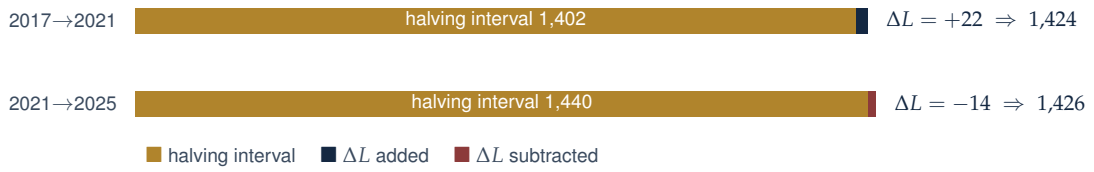

Figure 2: The top-to-top interval is the halving interval plus a small correction. It is not an independent measurement, and the near-constancy of the total across two cycles required L to swing 36 days in the opposite direction.

6.3 The leg the framework omits performed best

The halving-to-top lag has run 526, 548, and 534 days—a 22-day range around a 536-day mean. Used as a forecasting rule from the April 2024 halving, the two prior observations bracket a window of **28 September to 20 October 2025**. The actual top printed on 6 October 2025, comfortably inside it.

Rule	2025 out-of-sample call	Outcome
Low + 1,064 d	Point forecast: 20 Oct 2025	14 days late
Halving + 526–548 d	Window: 28 Sep – 20 Oct 2025	Contained the top

Table 8: The omitted relationship outperformed the circulated one on the only genuinely out-of-sample test available.

This is uncomfortable for a framework that treats the halving as a fading influence. The halving-to-top relationship is not proof of causation—22 days of range on three observations proves very little—but it is not weaker than the relationship the framework is built on, and it is the one nobody quotes.

6.4 A term worth getting right

A halving epoch is 210,000 blocks, which at the protocol's ten-minute block target is 1,458.3 days. This figure is frequently described as a ceiling. It is not one. Difficulty retargets every 2,016 blocks in *both* directions; if hashrate falls, an epoch can and will run longer than 1,458.3 days. The correct term is the **target duration at the ten-minute block interval**.

Epochs have historically run short of that target—1,319, then 1,402, then 1,440 days—because hashrate grew within each epoch and difficulty adjustment lagged the growth. As that growth decelerated, epochs lengthened toward the target. Note the direction: lengthening epochs mean block production has been *slowing* toward ten minutes, not accelerating.

SO WHAT

Do not read comparisons among these three intervals as evidence about which drives which. Equation (1) makes them algebraically dependent, so any such comparison is circular. The defensible conclusion is narrow: **the halving date is not a sufficient or mechanically deterministic timing signal**, and any model keying a price expectation to it should carry the 22-day dispersion of L explicitly rather than treating the relationship as exact. That is a very different statement from "stop using halving dates," which the evidence does not support.

7. Researcher Degrees of Freedom

7.1 What the anchors are, and are not

The six price anchors are sometimes described as free parameters. They are not fitted parameters in any statistical sense. The correct term is **researcher degrees of freedom**: the anchors are observations whose dating involves defensible but discretionary choices—exchange, intraday versus close, timezone—made by the analyst, potentially after seeing the target.

7.2 The count

Three complete cycles. Six price anchors, each movable by one to five days by specification choice. A researcher who selects anchors after seeing the target number will land on that number and cannot distinguish discovery from construction.

7.3 Sensitivity, bounded rather than sampled

Rather than test individual substitutions, the honest bound is the envelope. If every anchor can move ± 5 days, then any leg defined by two price anchors carries ± 10 days of specification uncertainty, and any leg with one protocol anchor carries ± 5.

Leg	Anchors	Envelope	Observed spread
Low → Top	price, price	±10 d	18 d ($n = 3$)
Top → Top	price, price	±10 d	2 d ($n = 2$)
Halving → Top	protocol, price	± 5 d	22 d ($n = 3$)
Low → Halving	price, protocol	± 5 d	29 d ($n = 3$)

Table 9: Specification envelope against observed dispersion.

The row that should trouble any reader is the second. The top-to-top spread of 2 days sits *inside* its own ±10-day specification envelope. The measurement is not precise enough to support the precision claimed for it. In this specific sense the tightest-looking number in the dataset is noise reported as signal—and that is a stronger objection to the 1,425-day figure than anything in this paper is able to raise against the 1,064-day one.

7.4 Scoring the forecast that made the framework famous

Low	Low + 1,064	Actual top	Error
14 Jan 2015	13 Dec 2017	17 Dec 2017	+4 d
15 Dec 2018	13 Nov 2021	10 Nov 2021	−3 d
21 Nov 2022	20 Oct 2025	6 Oct 2025	−14 d

Table 10: The 1,064-day rule scored. Two calls were fitted in hindsight; only the third was out-of-sample, and it was the worst.

Four days, three days, fourteen days is a better record than most published macro timing work. But two of the three were fitted after the fact. The single out-of-sample call had the largest error and was beaten by the halving-to-top window in Section 6.3.

8. Preregistration: The October 2026 Test

The framework is often described as resolving in October 2026. It does not. A low that forms in October cannot be known in October to have held, and any scoring rule requiring durability necessarily resolves later. The candidate window closes in October 2026; scoring occurs in 2027. This section fixes the protocol before the outcome rather than after it.

PREREGISTERED TEST — LODGED 9 AUGUST 2026

Series. BTC-USD daily closing price, Coinbase, UTC. Closes only; no intraday ticks.

Observation period. 6 October 2025 through 31 March 2027.

Definition of the low. L^* is the lowest daily close in the observation period. d^* is its date.

Hold condition. L^* qualifies as durable if no daily close in the 120 days following d^* is below L^*, and d^* falls at least 120 days before 31 March 2027.

Scoring date. 31 March 2027. No re-parameterisation before that date.

Mutually exclusive outcomes:

- **PASS** — d^* falls between 4 and 17 October 2026 inclusive, and the hold condition is met.
- **PARTIAL** — d^* falls between 1 September and 31 December 2026, hold condition met, but outside the 4–17 October window.
- **FAIL** — d^* falls outside 1 September to 31 December 2026, or the hold condition is not met.

Interpretation fixed in advance. PASS makes the decline-leg rule two-for-two out-of-sample and warrants instrumentation, not allocation. PARTIAL indicates the seasonal content is real and the day-level precision is not. FAIL retires the rule; it is not to be re-parameterised with a new constant and re-presented.

8.1 Position at lodging

Item (all as of 9 August 2026)	Value
All-time high	6 Oct 2025, ≈$126,198
Prior decline legs	363 d and 376 d
Candidate window	4–17 Oct 2026
Elapsed since high	307 days
Days to window open	56
Scoring date	31 Mar 2027
Spot (7 Aug 2026 reference)	≈$64,900, ≈48% below the high

Table 11: Open position of the framework at lodging. The spot reference is dated separately because it is a market observation rather than a computed interval.

8.2 The test that discriminates between the regularities

The October 2026 test scores the decline leg. It does not tell us which of the three regularities in Section 3 is primary. That question is decided at the next cycle peak, and the competing hypotheses make separable predictions.

Assume the next halving epoch runs approximately 1,450 days, placing the halving near 9 April 2028—longer than the last, consistent with the historical drift toward the 1,458-day target.

Hypothesis	Holds constant	Next top	Implies
H1	Top-to-top = 1,425 d	31 Aug 2029	L falls to ≈509 d
H2	L = 536 d	27 Sep 2029	Top-to-top rises to ≈1,452 d

Table 12: The two surviving hypotheses separate by roughly four weeks at the next peak. They cannot both be right.

Four weeks is a wide enough separation to score and narrow enough that neither can be dismissed now. It is the cleanest discriminating test the dataset offers, and it is three years out.

9. So What

SO WHAT

One. The headline is misattributed. 1,064 days is low-to-top, not top-to-top. This is the paper's most durable claim and the one most likely to be needed in a meeting. Correct it before repeating it.

Two. There is no single cycle constant. Three approximate regularities coexist—low-to-top near 1,060 days, halving-to-top near 536 days, top-to-top near 1,425 days—and Equation (1) makes them algebraically dependent. Any argument that one drives another is circular unless it brings evidence from outside these six dates. Nobody has.

Three. The relationship the framework omits is the one that performed best. Halving-to-top of 526–548 days bracketed the October 2025 peak out-of-sample; the 1,064-day rule missed it by fourteen days. If you carry one timing relationship, carry that one—with its 22-day dispersion attached, not as a point estimate.

Four. The halving is not disqualified as a timing input. The stronger version of this argument—stop using halving dates as timing signals—goes past the evidence. The defensible statement is that the halving date is not sufficient and not mechanically deterministic, and that models using it must carry the 22-day dispersion explicitly.

Five. 512 days (low to halving) fits twice out of three and depends on admitting the same early-cycle data the paper excludes elsewhere. 256 days has no referent in any construction reviewed. Do not repeat it without one.

Six. Precision claims in this literature routinely exceed the measurement, and the trap catches critics as readily as believers. A 2-day spread inside a ±10-day specification envelope is not a constant. Before quoting any figure here, ask whether its dispersion is larger than its specification uncertainty. For most of them it is not.

Seven. Two tests are lodged. The decline-leg test scores 31 March 2027 under the preregistered protocol in Section 8. The discriminating test between the surviving regularities resolves at the next peak, near September 2029. Write both down now and score against what was written. $n = 3$ does not support capital at risk under any of these readings.

Sources and Data Notes

Price anchors. Intraday extremes on major USD spot venues. The preregistered test in Section 8 uses daily closes instead, and specifies the venue and timezone, precisely because the intraday convention carries the ambiguity documented in Section 7.

Halving dates. Blocks 210,000 / 420,000 / 630,000 / 840,000, UTC timestamps.

Framework provenance. CoinMarketCap Academy, "What's Really Going On With Bitcoin Price and the BTC 4-Year Cycle?"; ShortSpike, "The 1064/364-Day Bitcoin Cycle"; CryptoPotato, "Bitcoin Cycle Timing: Could the Next BTC Bottom Arrive in October 2026?"; S. Tharayil, "The 1064/364 Bitcoin Theory: Genius Pattern or Beautiful Coincidence?" The archived original December 2023 post is not cited directly and has not been verified for this paper.

512-day relationship and halving-to-top intervals. Bitcoin Suisse, "The Bitcoin Halving and the Bitcoin Market Cycle"; J. Wedson / Alfractal, reported August 2025.

Spot reference. TheStreet Crypto and Yahoo Finance, 7 August 2026.

Reproducibility. All intervals are simple date differences reproducible from the anchor table in Section 2.1 with any date calculator. No proprietary data, smoothing, or transformation is used anywhere in the analysis.