

Is Bitcoin cyclical, and can its next bottom and top be forecast?

A backtested comparison of five model families

Badr Abardazzou 2026-09-13 babarda.com

Abstract. We ask two things of the daily Bitcoin price from 2010 to 2026: whether the level carries a multi-year cycle, and whether the date and level of the next cycle bottom and the next cycle top can be forecast with honest uncertainty. Cyclicity is supported when it is anchored to the halving schedule: a Rayleigh test rejects uniform phase for the four cycle tops ($R = 0.97$, $p = 0.012$) and for the three completed bottoms ($p = 0.037$), while a free-period structural model prefers a 2.75-year cycle and rejects four years at 5% (likelihood ratio 4.41), so the rhythm is halving-anchored rather than a free oscillation. Five forecasting families (calendar rules, a structural time-series model, a halving-phase template, the log-periodic power law and Markov regime switching) are compared in a rolling-origin backtest on every top and bottom since 2013. The halving-phase template has the smallest timing error (mean absolute error 43 days for bottoms and 55 days for tops in dollars) and, with the trended calendar rule, the smallest level error (0.30 and 0.65 in log price); the structural model has calibrated but wide intervals; the log-periodic model qualifies only on the 2017 run-up; regime switching captures volatility regimes of 8 to 34 weeks and has no multi-year content. The bottom of the current bust is treated as what it is, the minimum close of an open episode (2026-06-30 at 58,559 USD, day 801 after the halving, inside the tested 777-to-924-day window), and the forecast is the probability that it holds: 0.79 to 0.93 to the end of the halving window and 0.35 to 0.63 to the next halving under the two path-conditional models, against 3 of 3 under the cycle pattern. The pattern-based alternative, a new low at 31,162 to 41,856 USD by 2026-10-13 (the top cut by the 73% to 83% depth of the previous busts), carries a probability of 0.01 or less under the path models and is kept as the tested alternative. The next top falls in the second quarter of 2029 by three independent anchors, at 74,449 to 234,733 USD. Measured in gold, a robustness check, the current bust is an ordinary cycle-ending bust (433 days, 69% deep, ended 2026-02-23), which is why the dollar chart's unusually short and shallow bust is read as a yardstick effect. The conditions that refute each call are stated. Everything is reproducible from cached data with fixed seeds.

1 Introduction

Bitcoin's price history is short by the standards of asset pricing, sixteen years, and loud by the standards of anything else: four booms, four busts, each bust taking more than half the value off the top. One claim about that history circulates widely: that the price is cyclical, that a four-year rhythm tied to the halving of the block subsidy lets one date the next bottom and the next top. The claim is usually made from charts and rarely tested with the tools that exist for testing it.

This paper does two things. It tests the cyclical claim three independent ways: a periodogram against a red-noise null, a structural time-series model with a free cycle period, and a circular-statistics test of whether tops and bottoms cluster at a fixed phase of the halving

epoch. And it treats the forecasting question as a forecasting question: five published model families are fitted to the same data through one interface, backtested at rolling origins on every top and bottom we have, scored in days and in log price against naive baselines, and combined into an ensemble whose intervals come from its own backtest errors rather than from any model's internal standard errors. The result is one dated, falsifiable statement about the bottom of the current bust and the next top.

The analysis is in US dollars. The same bust measured in ounces of gold is used once, as a robustness check on whether the current bust is unusual: gold removes the dollar's own movement, and in gold the bust turns out to be ordinary.

The contribution is not a new model. It is the discipline of running the existing ones side by side, with the same events, the same horizons, the same scoring, and the same refusal to report a point without the backtest error that belongs to it.

Is Bitcoin cyclical, and can its next bottom and top be forecast? Five model families, one backtest, one dated call, in dollars and in gold.




Figure 1: How the study is built: one data layer, two questions, five model families, one backtest, one dated call, with gold as a robustness check.

2 Related work

Calendar anomalies in equities have a long literature, from the January effect (Rozeff and Kinney 1976) to the day-of-week effect (French 1980), and an equally long record of effects shrinking or vanishing once multiple testing is accounted for (Sullivan, Timmermann and White 2001). The permutation-with-correction design used here follows that lesson.

Periodicity in unevenly sampled or short series is the domain of the Lomb-Scargle periodogram (Lomb 1976; Scargle 1982), whose significance against autocorrelated noise must be assessed by simulation rather than by the white-noise formula. Stochastic cycles inside a structural time-series model are Harvey's (1989) formulation, estimated by the Kalman filter (Durbin and Koopman 2012); the profile likelihood over the cycle period is the natural way to show how well the period is identified. The Rayleigh test for circular uniformity (Mardia and Jupp 2000) is the standard test of whether events cluster at a phase.

Bubble timing by log-periodic power-law singularities goes back to Sornette, Johansen and Bouchaud (1996) and Johansen, Ledoit and Sornette (2000); the estimation used here is the reformulation of Filimonov and Sornette (2013) with the filters of Sornette et al. (2015), and the negative-bubble form for declines is from Yan, Woodard and Sornette (2012). Regime switching in returns is Hamilton (1989). Rolling-origin evaluation and the insistence on beating a naive benchmark follow Hyndman and Athanasopoulos (2021).

On Bitcoin specifically, the halving-cycle narrative is widely discussed in practitioner writing and rarely subjected to out-of-sample testing; the closest academic treatments look at bubbles in Bitcoin with LPPLS (Gerlach, Demos and Sornette 2019) and at calendar effects in cryptocurrency returns (Caporale and Plastun 2019), with mixed results that depend on the sample. This paper’s contribution relative to those is the joint design: the same events, both denominations, five families, one backtest.

3 Data and event definitions

3.1 Series

Daily Bitcoin prices from 2010-08-18 to 2026-09-13: blockchain.com’s averaged market price until 2014-09-16 and Yahoo Finance’s BTC-USD close from 2014-09-17, spliced with a 2.91% log gap on the splice day. For the robustness check in section 6.4, gold is COMEX front-month futures (Yahoo GC=F), forward-filled over weekends, and the gold-denominated series is ounces per bitcoin, $Q_t = P_t/G_t$. Halving dates at block heights 210,000, 420,000, 630,000 and 840,000: 2012-11-28, 2016-07-09, 2020-05-11, 2024-04-20. The fifth halving H5 is projected at 2028-02-21 (H4 plus the median inter-halving gap of 1,402 days; the minimum and maximum gaps give 2027-11-30 to 2028-03-30). Weekly series (Sunday close, log price, 839 weeks) serve the structural and regime models. Sources, fetch dates and hashes are in DATA_BASIS.md.

3.2 Events

A bust is one episode per all-time high: from the high to the day the running maximum is next exceeded, kept if the drawdown inside it goes below 50%. Within each halving epoch the cycle-ending bust is the one whose bottom lies furthest from the halving; its starting high is the cycle top T_k and its minimum the cycle bottom B_k . The rule has one parameter and recovers the dates practitioners name without tuning. The current bust has not recovered; its low so far is provisional and is never scored.

k	halving	top	bottom	H to T	H to B	T to B	ln(top/H)	depth
1	2012-11-28	2013-12-05	2015-01-14	372	777	405	4.51	84%
2	2016-07-09	2017-12-16	2018-12-15	525	889	364	3.40	83%
3	2020-05-11	2021-11-08	2022-11-21	546	924	378	2.06	77%
4	2024-04-20	2025-10-06	2026-06-30 (prov.)	534	801	267	0.65	53%

Table 1: Cycle events in dollars. Days between events; $\ln(\text{top}/H)$ is the log multiple of the top over the halving-day price; depth is the bust’s drawdown.

Two regularities in the table drive the modelling. Timing in halving phase is tight: tops at 372 to 546 days and bottoms at 777 to 924 days after the halving. Amplitude decays fast and roughly linearly in log units: the log top multiple falls 4.51, 3.40, 2.06, 0.65, and the log depth of the bust falls 1.85, 1.80, 1.45 over the three completed cycles. Timing and amplitude are therefore modelled separately, and amplitude in log units.

4 Methods

4.1 Part A: cyclicality

A1, periodogram. Lomb-Scargle on the log price after removing a long-run trend, over periods of 0.5 to 8 years, against 500 AR(1) surrogates fitted to the detrended series. Two detrendings are reported because they disagree: linear in time, and linear in log time (the decelerating growth curve). A block bootstrap (250-day blocks, 400 draws) gives the interval on the peak period.

A2, structural time series. Weekly log price minus a log-time trend, then Harvey's (1989) model with a random-walk level and a stochastic damped cycle, $\psi_{t+1} = \rho(\psi_t \cos \lambda + \psi_t^* \sin \lambda) + \kappa_t$, estimated by maximum likelihood through the Kalman filter (statsmodels UnobservedComponents). The cycle frequency λ is fixed on a grid of periods from 1 to 8 years in steps of 0.25 and the other parameters re-estimated at each point; the profile log-likelihood shows how well the period is identified and the likelihood ratio between the best period and four years tests the halving story as a free oscillation.

A3, phase concentration. Each top and bottom is placed at its phase of the halving epoch, $\theta = 2\pi(t - H_k)/L_k$ with L_k the epoch length (1,319, 1,402, 1,440 days; the median for the open epoch). The Rayleigh statistic $R = |\sum e^{i\theta}| / n$ tests uniform phase (Mardia and Jupp 2000, section 6.3, with the small-sample correction), cross-checked against 10,000 uniform draws. A1 to A3 are also run on the gold-denominated series as a robustness check.

4.2 Part B: five forecasting families, one interface

Every model implements `forecast(origin, kind, target)` and returns a date and a level, each with an interval, using only data up to and including the origin day. Horizons are shared: from a halving, a top is searched in the next 730 days and a bottom in days 365 to 1,460; from an origin after a top, a bottom in the next 730 days; from an origin after a bottom, a top in the next 1,100 days. Levels are in log price.

M0a, calendar rule (median). From a halving: top date H_k plus the median halving-to-top span over earlier epochs, bottom date likewise; level at the top $\ln P(H_k)$ plus the median log top multiple, at the bottom that minus the median log depth. Intervals are the minimum to maximum over earlier epochs. From an origin after a top: bottom date = top plus the median top-to-bottom span, level = observed top minus median log depth. From an origin after a bottom: the next top is anchored to the next halving. The baseline for timing.

M0b, calendar rule (trended). Same dates; levels from a linear trend across epochs in the log top multiple and the log depth, extrapolated one step, with a t -based 90% interval from the residuals (or the spread when only two epochs exist).

M0c, random walk with drift. Log price at the origin plus the drift of the last four years times the horizon to the calendar-rule date; interval $\pm 1.645\sigma\sqrt{h}$. The baseline for levels; it makes no date forecast of its own.

M1, structural time series. The Part A2 model, frequency free, re-fitted on the training window at each origin; 2,000 paths simulated 208 weeks ahead from the filtered end state, the trend added back, and the extremum inside the horizon read on each path. Point = median, interval = 5th to 95th percentile across paths.

M2, halving-phase template. For each earlier epoch j , $y_j(d) = \ln P(H_j + d) - \ln P(H_j)$ for $d = 0 \dots 1440$, each curve scaled by its own maximum; the template is their mean. Timing: the template's argmax and, after it, argmin; intervals from the per-epoch extrema. Amplitude: the log top multiple and the log depth extrapolated by their trend across epochs (M0b's rule). After a top has occurred, the observed top replaces the extrapolated one and the template's top-to-bottom span gives the bottom date.

M3, log-periodic power law. $\ln P(t) = A + B(t_c - t)^m + C_1(t_c - t)^m \cos(\omega \ln(t_c - t)) + C_2(t_c - t)^m \sin(\omega \ln(t_c - t))$ (Filimonov and Sornette 2013). The nonlinear triple (t_c, m, ω) is found by a grid of $15 \times 9 \times 8$ points followed by Nelder-Mead from the three best, the linear four by least squares at every evaluation. A fit qualifies if $0.1 < m < 0.9$, $6 < \omega < 13$, t_c lies within half the window length beyond its end, $B < 0$ for a bubble ($B > 0$ for the negative-bubble form

applied to declines, Yan, Woodard and Sornette 2012), and the damping $m |B| / (\omega C) \geq 0.8$ (Sornette et al. 2015). Nineteen windows ending at the origin, 180 to 720 days long, are fitted; the forecast is the median t_c across qualifying windows with their 5th to 95th percentile, and the share of qualifying windows is the LPPLS confidence. Windows are refitted every 30 days through each run-up and each decline to trace how t_c evolves.

M4, Markov regime switching. Two regimes on weekly log returns with switching mean and variance (Hamilton 1989; statsmodels `MarkovRegression`, 20 random starts). Expected regime duration $1/(1 - p_{ii})$; from an origin, the target date is the origin plus the expected remaining duration of the relevant regime, the interval from the geometric 5th and 95th percentiles, the level the origin price plus the regime mean times the duration. Specifications with three regimes and with monthly returns were also fitted.

4.3 Part C: rolling-origin backtest

Origins: each halving H_k ($k = 2, 3, 4$) forecasting T_k and B_k ; each top plus 60 days forecasting B_k ; each bottom plus 60 days forecasting T_{k+1} ; and two “late” families inside the move, $T_k - 60$ days forecasting T_k and $B_k - 90$ days forecasting B_k . A model at an origin sees only events completed before it. Scores per row: signed and absolute error in days and in log level, and whether the truth fell inside the model’s own interval. With three to nine scored rows per model and target no formal test of equal accuracy is possible (Diebold and Mariano 1995 needs far more), so every row is published and the summary is the mean absolute error, the bias and the coverage. The provisional current bottom is never scored.

4.4 Part D: ensemble

For each target, models whose backtest MAE is strictly better than the baseline’s (M0a for dates, M0c for levels) on the relevant origin families, with at least two scored rows, receive weight proportional to $1/\text{MAE}$; if none beats the baseline, the baseline alone is used. M0b and M0c inherit M0a’s dates and are excluded from date weighting. The ensemble forecast is the weighted mean of the models’ forecasts at the latest origin. Its interval comes from its own record: the ensemble’s signed error is recomputed with the same weights at every scored backtest origin at which every weighted model produced a forecast (origins in the first epoch, where the calendar models have nothing to draw on, are excluded, because an ensemble of the remaining models is not the ensemble being used), and the interval is the forecast minus the largest of those errors to the forecast minus the smallest, the full range, since there are four to six of them. No bias correction is applied with that few rows; the range carries the bias. The weights are in-sample; this is stated as a limitation. Falsification conditions are written from the interval.

5 Results


5.1 Cyclicity: halving-anchored, not a free oscillation

A1. Under the log-time trend the periodogram peaks at 3.64 years and clears the AR(1) 99% band; under a linear trend the largest peak is 6.95 years and does not clear 95%. The block-bootstrap interval on the peak period is 1.5 to 8.0 years.

A2. The structural model’s profile likelihood peaks at 2.75 years (log-likelihood 534.1 against 531.9 at four years, likelihood ratio 4.41, $p = 0.036$ on one degree of freedom); the free fit lands at 1.45 years with damping 0.965. In gold the peak is at 1.25 years (ratio 5.08 against

four years, free fit 1.32 years, damping 0.962). A model that lets the cycle float does not pick four years in either unit.

Profile likelihood of the cycle period in the structural model (red dotted: 95% likelihood cut-off)




Source: babarda.com, code/ in the article repo; data per DATA_BASIS.md

Figure 2: Profile log-likelihood of the cycle period in the structural model, in dollars (left) and, as a check, in gold (right). The four-year period (grey dashed) sits below the 95% cut-off (red dotted) in both.

A3. The four tops sit at phases 0.28, 0.37, 0.38, 0.38 of their epochs; the Rayleigh test gives $R = 0.966$, $p = 0.012$ (0.004 by simulation). The three completed bottoms sit at 0.59, 0.63, 0.64: $R = 0.989$, $p = 0.037$ (0.0095 by simulation). The same test in gold gives tops $R = 0.874$, $p = 0.035$ and bottoms $R = 0.970$, $p = 0.044$. Uniform phase is rejected for both event types, at conventional levels, from three and four points, and the result survives the change of yardstick.

Where in the halving epoch do tops and bottoms fall? Rayleigh test for phase concentration



Source: babarda.com, code/ in the article repo; data per DATA_BASIS.md


Figure 3: Phase of each top (red) and bottom (teal) within its halving epoch, in dollars (left) and in gold (right). The events cluster; the Rayleigh test rejects uniformity.

Taken together: the price does not oscillate freely at four years, but its extrema fall at a fixed phase of the halving epoch. The cycle is a calendar phenomenon anchored to the supply schedule, which is why the calendar-anchored models dominate the backtest.

5.2 Models

M2 template. The mean of the scaled earlier epochs peaks at day 546 and bottoms at day 924 after the halving in dollars; the individual epochs peak at 372, 525 and 546 days and bottom at 777, 889 and 924.


The halving-phase template: each epoch scaled to its own top, and their mean



Source: babarda.com, code/ in the article repo; data per DATA_BASIS.md

Figure 4: Each epoch scaled to its own top, and the template from epochs 1 to 3; dollars left, gold right (where the small fourth-epoch top multiple exaggerates that curve).

M3 LPPLS. Only the 2017 run-up qualifies. At the 2017-09-17 origin 21% of windows pass and the median t_c is 2017-11-01 (5th to 95th percentile 2017-10-17 to 2017-11-24) against the 2017-12-16 top; the signal creeps forward with the origin and stays three to six weeks early, the classic behaviour. No window qualifies on the 2013, 2021 or 2025 run-ups, and the negative-bubble form qualifies on none of the declines (median damping 0.1 to 0.5, below the 0.8 filter). LPPLS is therefore a 2017-only instrument in this sample.




Source: babarda.com, code/ in the article repo; data per DATA_BASIS.md

Figure 5: LPPLS critical time as the origin moves through the 2017 run-up in dollars: the median t_c (teal) advances with the origin and stays a few weeks ahead of the actual top (red dotted).

M4 Markov switching. The two regimes are a high-mean, high-variance state (mean +5.8% per week, standard deviation 24.9%, expected duration 7.7 weeks) and a low-mean, low-variance state (+0.7%, 7.5%, 33.0 weeks). No specification produces a persistent negative-mean regime.

The model describes volatility clustering on a horizon of weeks and carries no information about multi-year phase.

Two-regime Markov switching on weekly returns (dotted: cycle tops)



Source: babarda.com, code/ in the article repo; data per DATA_BASIS.md

Figure 6: Smoothed low-variance regime (shaded) on the weekly log price, in dollars (top) and gold (bottom). Regimes last weeks, not years.

5.3 Backtest

model	target	rows	MAE days	bias days	MAE log	date / level coverage
M0a calendar median	top	9	87	-87	1.93	0.33 / 0.00
M0a calendar median	bottom	6	50	-18	0.57	0.33 / 0.00
M0b calendar trended	top	9	87	-87	0.65	0.33 / 0.62
M0b calendar trended	bottom	6	50	-18	0.30	0.33 / 0.17
M0c random walk	top	9	87	-87	0.87	0.00 / 0.67
M0c random walk	bottom	6	50	-18	1.02	0.00 / 0.50
M1 structural	top	9	311	+240	0.90	0.67 / 0.78
M1 structural	bottom	8	180	-164	0.65	1.00 / 0.88
M2 template	top	9	55	-47	0.65	0.33 / 0.62
M2 template	bottom	6	43	-16	0.30	0.33 / 0.17
M3 LPPLS	top	1	59	-59	1.25	0.00 / n.a.
M3 LPPLS	bottom	0				
M4 Markov	top	9	486	-486	1.38	0.33 / 0.44
M4 Markov	bottom	8	285	-246	1.01	0.62 / 0.00

Table 2: Backtest, scored rows only. Rows = origins with a forecast (the calendar rules and the template need an earlier epoch, so they produce nothing at the first origins). MAE and bias in days are on the date; MAE log on the level; coverage is the share of truths inside the model's own interval.

Three findings. First, the halving-phase template is the best timing model and beats the calendar median it is built from (43 against 50 days on bottoms and 55 against 87 on tops). Second, on levels the trended amplitude rule (shared by M0b and M2) is the best (0.30 log on bottoms, 0.65 on tops), the median rule is badly biased upward because amplitude decays, and the random walk is the worst on bottoms. Third, the structural model is the only one whose own intervals are calibrated (date coverage 1.00 on bottoms, level coverage 0.88 to 1.00), at the price of timing errors of six months to a year; the calendar rules' min-to-max intervals from two or three epochs cover a third of the truths or fewer, which is why the ensemble takes its

intervals from backtest errors and not from the models. Every individual row is published in the supplementary results file (see the appendix).

6 The forecast, issued 2026-09-13

Price on the issue date: 76,851 USD. The current cycle top is 2025-10-06 at 124,753 USD.

6.1 What the data determines about the bottom

Three things are determined, in this order of strength.

The window. The tested regularity (section 5.2, A3) is that cycle bottoms fall 777 to 924 days after the halving. For the current epoch that is **2026-06-06 to 2026-10-31**. In every completed cycle the bottom fell inside its window and no lower close followed before the next all-time high (3 of 3).


The observed minimum. The minimum close since the top is not a forecast; it is the argmin of the price over the open episode: **2026-06-30 at 58,559 USD**, day 801 after the halving, inside the window. The five lowest closes of the episode all fall between 2026-06-25 and 2026-07-01 (58,559 to 60,004 USD).

The probability that the minimum holds. Whether 2026-06-30 is the cycle bottom is a question about the future path, so it is answered by the two models that condition on the path: the structural model simulated forward from its filtered state (4,000 paths) and a random walk with the drift and volatility of the last four years (first-passage probability, reflection principle). The probability that a close below 58,559 USD prints:

horizon	days	structural	random walk
end of the halving window, 2026-10-31	48	0.21	0.07
year end, 2026-12-31	109	0.39	0.18
next halving, 2028-02-21	526	0.65	0.37

Table 3: Probability of a close below the provisional bottom by each horizon. The structural model is the more pessimistic because its fitted cycle (1.4 years) can still swing down; the random walk carries the positive drift of the last four years (0.091% per day).

Probability that the 2026-06-30 low survives, by horizon (red: halving window closes, grey: next halving)



Source: babarda.com, code/ in the article repo; data per DATA_BASIS.md

Figure 7: The same as a survival curve: probability that no lower close has printed by each date. Red dotted line: the halving window closes; grey: the next halving.

So the statement the data supports is a probability, not a date: **the 2026-06-30 low is the bottom of the halving window with probability 0.79 to 0.93**, and the cycle bottom in the strict sense (no lower close before the next halving) with probability 0.35 to 0.63 under models that know nothing of the cycle, and 3 of 3 under the cycle pattern itself.

6.2 The pattern-based alternative (scenario A) and where its numbers come from

The calendar and template models do not condition on the path; they apply the previous busts' timing and depth to the observed top. Their arithmetic: the three completed cycle-ending busts had log depths of 1.85, 1.80 and 1.45 (84%, 83%, 77%). The median rule applies 1.80: $124,753 \times e^{-1.80} = 20,710$ USD. The trended rule fits a line through the three depths, falling 0.20 per epoch, and extrapolates to 1.30: $124,753 \times e^{-1.30} = 33,954$ USD, a 73% drawdown. The structural model's simulated median is 45,626 USD; the Markov model's 96,058 USD; the random walk (79,421 USD) is the baseline and takes no weight. With the backtest weights (M0a 0.24, M0b 0.30, M1 0.12, M2 0.30, M4 0.05) the weighted mean is 32,940 USD, and the ensemble's four backtest errors at the 2018 and 2022 bottoms (+0.06, +0.03, -0.14, -0.24 in log) give the range 31,162 to 41,856 USD. On timing the calendar rule gives 2026-10-19 raw; its four backtest errors were +41, +41, +6 and +6 days (late every time), so 2026-09-08 to 2026-10-13.

Scenario A is therefore: a close below 41,856 USD by 2026-10-13. The path-conditional models give it a probability of 0.01 (structural) and 0.00 (random walk): a 47% fall in a month from 76,851 USD. It is kept in the paper because it is what three cycles of dollar depths say, and because it is the alternative that the next month refutes or confirms.

model	B4 date	B4 level USD	T5 date	T5 level USD
M0a calendar median	2026-10-19	20,710	2029-08-03	1,604,251
M0b calendar trended	2026-10-19	33,954	2029-08-03	145,538
M0c random walk	2026-10-19	79,421	2029-08-03	201,529
M1 structural	2027-04-25	45,626	2028-12-10	364,495
M2 template	2026-10-19	33,954	2029-08-20	145,538
M3 LPPLS	no qualifying window		no run-up yet	
M4 Markov	2027-05-01	96,058	2026-11-06	120,359

Table 4: Each model's point forecast from the 2026-09-13 origin. The Markov model's "top" is the end of the current low-variance regime, not a cycle top, and it receives no weight.

6.3 The next top

Two anchors, both from tested regularities. From the halving: H5 (2028-02-21) plus 372 to 546 days gives **2029-02-27 to 2029-08-20**. From the bottom: the three previous bottom-to-top spans were 1,067, 1,059 and 1,050 days, the most stable quantity in the data; applied to 2026-06-30 they give **2029-05-15 to 2029-06-01**. The two anchors overlap in May 2029. The ensemble of models that beat the calendar baseline (M1 0.46, M2 0.54) gives 2029-04-25 raw and, after its backtest errors of -108 to +4 days, **2029-04-21 to 2029-08-11**. Three routes, one answer: the second quarter of 2029, conditional on the June low holding; if a lower low prints, the bottom-anchored date moves with it.

Level: the trended bottom-to-top rise (4.70, 3.04, 2.07 in log, extrapolated to 0.64) applied to the issue-date price gives $76,851 \times e^{0.64} = 145,538$ USD; the structural model gives 364,495 USD; the weighted mean over M0b (0.33), M1 (0.34) and M2 (0.33) is 199,049 USD with a range of 74,449 to 234,733 USD from five ensemble errors (-0.17 to +0.98). Nothing narrower is defensible; the log rise has fallen by 1.3 to 1.7 per epoch and the next value is an extrapolation from three points.

Where the bottom is and where the next top should be, issued 2026-09-13



Source: babarda.com, code/ in the article repo; data per DATA_BASIS.md

Figure 8: The forecast on the price history since 2022, real prices on a log scale: the halving window in which the bottom should fall, the lowest close so far with the chance it holds, the pattern-based alternative (faint red, 1% chance), and the next-top window with its level range.

6.4 The same bust measured in gold: a robustness check

Dividing the price by the gold price removes the dollar's own movement. Three facts survive that change of yardstick, each holding in every completed cycle:

cycle	top (gold)	bottom (gold)	H to T days	T to B days	depth
1	2013-12-05	2015-01-14	372	405	84%
2	2017-12-16	2019-02-07	525	418	83%
3	2021-10-20	2023-01-03	527	440	75%
4	2024-12-17	2026-02-23 (prov.)	241	433	69%

Table 5: Cycle events in ounces of gold per bitcoin. Rayleigh test on the halving phase in gold: tops $R = 0.874$, $p = 0.035$; bottoms $R = 0.970$, $p = 0.044$.

First, the halving-phase clustering holds in gold as in dollars. Second, the top-to-bottom duration is tighter in gold (405, 418, 440 days) than in dollars (405, 364, 378), and the current bust in gold ran 433 days, an ordinary value. Third, the depth decays monotonically (84, 83, 75, 69%). In gold, then, the current bust is a textbook cycle-ending bust that ended on 2026-02-23 at 12.4 ounces; in dollars it is the shortest and shallowest on record only because gold rose 64% over 2025 while the dollar chart was still climbing. That is the evidence that the dollar oddity is the yardstick, not Bitcoin, and it is why the June low is read as the bottom rather than as an interruption.

Gold is not carried further. The gold-denominated top of this cycle came 241 days after the halving instead of 525 to 546, so the timing models fail on gold tops (backtest MAE 159 to 172 days against 55 to 87 in dollars), and a gold-denominated forecast of the next top spans a factor of ten (34 to 350 ounces from the one model that beats the baseline). The full gold run is in the supplementary results.

6.5 The dated statement

Issued 2026-09-13. The cycle bottom is 2026-06-30 at 58,559 USD with probability 0.79 to 0.93 over the halving window and 0.35 to 0.63 to the next halving; the pattern-based alternative, a close between 31,162 and 41,856 USD by 2026-10-13, carries 0.01 or less. The next top falls in the second quarter of 2029, at 74,449 to 234,733 USD. Refutation: a close below 58,559 USD before 2028-02-21 refutes the bottom call; no all-time high inside the top window, or a top outside the level range, refutes the top call.

7 Discussion

What the cycle is. The evidence separates two things that popular writing runs together. A free four-year oscillation in the price is not supported: the structural model prefers shorter periods and rejects four years at 5%, and the periodogram's peak moves from 3.6 to 7 years depending on the trend removed. A halving-anchored calendar of extrema is supported: tops and bottoms fall at a fixed phase of the epoch, the Rayleigh test rejects uniform phase (and does so again in gold), and the models that exploit that calendar win the backtest. Bitcoin does not have a cycle in the sense of a pendulum; it has a schedule in the sense of a season, with the halving as the equinox.

Why the template wins. The template encodes the one thing the backtest shows to be stable, the phase of the extrema, and leaves amplitude to a separate trend. Models that try to learn the timing from the price path alone (the structural model, regime switching) have to estimate a period from three and a half cycles of noisy data and cannot. LPPLS needs a super-exponential run-up with log-periodic ripples; 2017 had one, 2021 and 2025 did not, and the declines never had the mirror image.

What gold settles. The current bust in dollars is the shortest and shallowest on record (267 days to its low, 53%). Measured in gold it is a textbook cycle-ending bust (433 days, 69%), because gold rose 64% over 2025 while Bitcoin stalled. That single comparison is what decides between the two readings of the June low: the pattern-based dollar models read it as unfinished business, the gold yardstick says the bust already ran its course, and the path-conditional models put the probabilities of section 6 on it. Gold is not used for the forecast itself: the gold-denominated top this cycle came 241 days after the halving instead of 525 to 546, so the timing models fail on gold tops (MAE 159 to 172 days) and a gold-denominated next-top range spans a factor of ten.

Amplitude. The log top multiple has fallen by about 1.1 to 1.4 per epoch (4.51, 3.40, 2.06, 0.65 in dollars). Extrapolated linearly it reaches zero within the next epoch, which is why the trended level rule places the next top at 145,538 USD, only slightly above the 2025 high, while the median rule places it at 1.6 million. The ensemble sits between and its interval is wide by construction. If amplitude decays geometrically rather than linearly, the trended rule is too pessimistic; the data cannot yet tell the two apart, and the next top will.

Limitations. Three and a half cycles. Every interval here is wide because of that and every claim of significance rests on three or four points, tested with methods built for small samples but still three or four points. The ensemble weights are chosen on the same backtest that reports the errors; a leave-one-epoch-out weighting would be cleaner and, with three epochs, unstable. The calendar models condition on the top having occurred without modelling the path since, which is why their dollar forecast of the current bust ignores the bounce from the June low and why the bottom call is stated as a path-conditional probability instead. The

probabilities themselves come from a structural model with a fitted 1.4-year cycle and a random walk, neither of which knows about the halving, so they bracket rather than pin the answer. Pre-2014 prices are an exchange average. Gold is front-month futures. The halving schedule is the only anchor considered; ETF flows, macro liquidity and reflexive belief in the four-year story could each produce the same phase clustering.

8 Conclusions

1. The price is cyclical in a specific sense: its tops and bottoms fall at a fixed phase of the halving epoch (Rayleigh $p = 0.012$ for tops and 0.037 for bottoms; 0.035 and 0.044 when the same test is run in gold), while a free-period model rejects a four-year oscillation. The cycle is a schedule, not a pendulum.
2. Of five model families backtested on every top and bottom since 2013, the halving-phase template times events best (MAE 43 days on bottoms, 55 on tops) and the trended amplitude rule levels them best (0.30 and 0.65 in log price). The structural model alone produces calibrated intervals. LPPLS qualifies only on the 2017 run-up. Regime switching has no multi-year content.
3. The bottom of the current bust is the minimum close of the open episode, 2026-06-30 at 58,559 USD, inside the tested halving window; the probability that it holds is 0.79 to 0.93 to the end of the window and 0.35 to 0.63 to the next halving under the path-conditional models, 3 of 3 under the cycle pattern. Measured in gold the bust is complete and ordinary (433 days, 69%, ended 2026-02-23), which is why the dollar chart's short, shallow bust is read as a yardstick effect. The pattern-based alternative, a new low at 31,162 to 41,856 USD by 2026-10-13, carries 0.01 or less.
4. The next top falls in the second quarter of 2029 by three independent anchors (halving plus 372 to 546 days, bottom plus 1,050 to 1,067 days, the model ensemble), at 74,449 to 234,733 USD. The level range is wide because the log rise has fallen 1.3 to 1.7 per epoch and three points cannot pin the next value.

Everything above is reproducible from the cached data with fixed seeds, and every number in this text is checked against the results files by `code/verify_text.py` before publication.

9 References

- Caporale, G. M. and Plastun, A. (2019). The day of the week effect in the cryptocurrency market. *Finance Research Letters*, 31, 258-269.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3), 253-263.
- Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*, 2nd ed. Oxford University Press.
- Filimonov, V. and Sornette, D. (2013). A stable and robust calibration scheme of the log-periodic power law model. *Physica A*, 392(17), 3698-3707.
- French, K. R. (1980). Stock returns and the weekend effect. *Journal of Financial Economics*, 8(1), 55-69.
- Gerlach, J.-C., Demos, G. and Sornette, D. (2019). Dissection of Bitcoin's multiscale bubble history from January 2012 to February 2018. *Royal Society Open Science*, 6(7), 180643.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357-384.

- Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65-70.
- Hyndman, R. J. and Athanasopoulos, G. (2021). *Forecasting: Principles and Practice*, 3rd ed. OTexts, section 5.10.
- Johansen, A., Ledoit, O. and Sornette, D. (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance*, 3(2), 219-255.
- Lomb, N. R. (1976). Least-squares frequency analysis of unequally spaced data. *Astrophysics and Space Science*, 39, 447-462.
- Mardia, K. V. and Jupp, P. E. (2000). *Directional Statistics*. Wiley, section 6.3.
- Rozeff, M. S. and Kinney, W. R. (1976). Capital market seasonality: the case of stock returns. *Journal of Financial Economics*, 3(4), 379-402.
- Scargle, J. D. (1982). Studies in astronomical time series analysis II. *Astrophysical Journal*, 263, 835-853.
- Sornette, D., Johansen, A. and Bouchaud, J.-P. (1996). Stock market crashes, precursors and replicas. *Journal de Physique I*, 6(1), 167-175.
- Sornette, D., Demos, G., Zhang, Q., Cauwels, P., Filimonov, V. and Zhang, Q. (2015). Real-time prediction and post-mortem analysis of the Shanghai 2015 stock market bubble and crash. *Journal of Investment Strategies*, 4(4), 77-95.
- Sullivan, R., Timmermann, A. and White, H. (2001). Dangers of data mining: the case of calendar effects in stock returns. *Journal of Econometrics*, 105(1), 249-286.
- Yan, W., Woodard, R. and Sornette, D. (2012). Diagnosis and prediction of rebounds in financial markets. *Physica A*, 391(4), 1361-1380.
- Data: Yahoo Finance (BTC-USD, GC=F) and blockchain.com charts API (market-price, total-bitcoins), fetched 2026-09-13; hashes in DATA_BASIS.md. Software: Python 3.14, statsmodels 0.15, scipy 1.18.

10 Appendix: reproducibility

All code is in `code/` of the article repository, run in this order with python: `utils.py` (loader, bust detector, events, weekly series, H5 projection), `cyclicity.py` (Rayleigh test), `m1_ucm.py`, `m2_template.py`, `m3_lppls.py`, `m4_markov.py`, `backtest.py`, `final_forecast.py`, `bottom_probability.py`, `make_report.py`, `verify_text.py`. Seeds are fixed at 20260913. Cached data with SHA-256 hashes and fetch dates are listed in `DATA_BASIS.md`. Every scored backtest row, every model's forecast at the latest origin and every ensemble weight are in the supplementary results file of the repository (reports folder) and the results folder. The one function the whole study rests on, the event detector:

```
def busts(dd, threshold=-0.5):
    """One episode per all-time-high peak: from the peak to the day the running peak
    is next
    exceeded (or to the end of the sample, flagged recovered=False). Kept if the
    drawdown
    inside the episode goes below threshold. Bottom = argmin of the drawdown in the
    episode."""
    at_high = dd == 0
    highs = np.flatnonzero(at_high.values)
    starts = [h for k, h in enumerate(highs) if k + 1 == len(highs) or highs[k + 1] != h + 1]
    rows = []
    for s_ in starts:
        later = highs[highs > s_]
        e = later[0] if len(later) else len(dd)
        seg = dd.iloc[s_:e]
        if seg.min() < threshold:
            bottom = seg.idxmin()
            rows.append(dict(peak=dd.index[s_], bottom=bottom, depth=float(-seg.min()),
                                recovered=bool(len(later))))
    return pd.DataFrame(rows)
```

The cycle-ending bust of each epoch is the one whose bottom lies furthest after the halving; its peak is T_k and its minimum B_k .