THE BLOCK CLOCK

Bitcoin's cycle is not four years long. It never was.

The halving is an exact 210,000-block period, and measuring it in days is what makes it look loose. Here is what changes when you use the protocol's own unit.

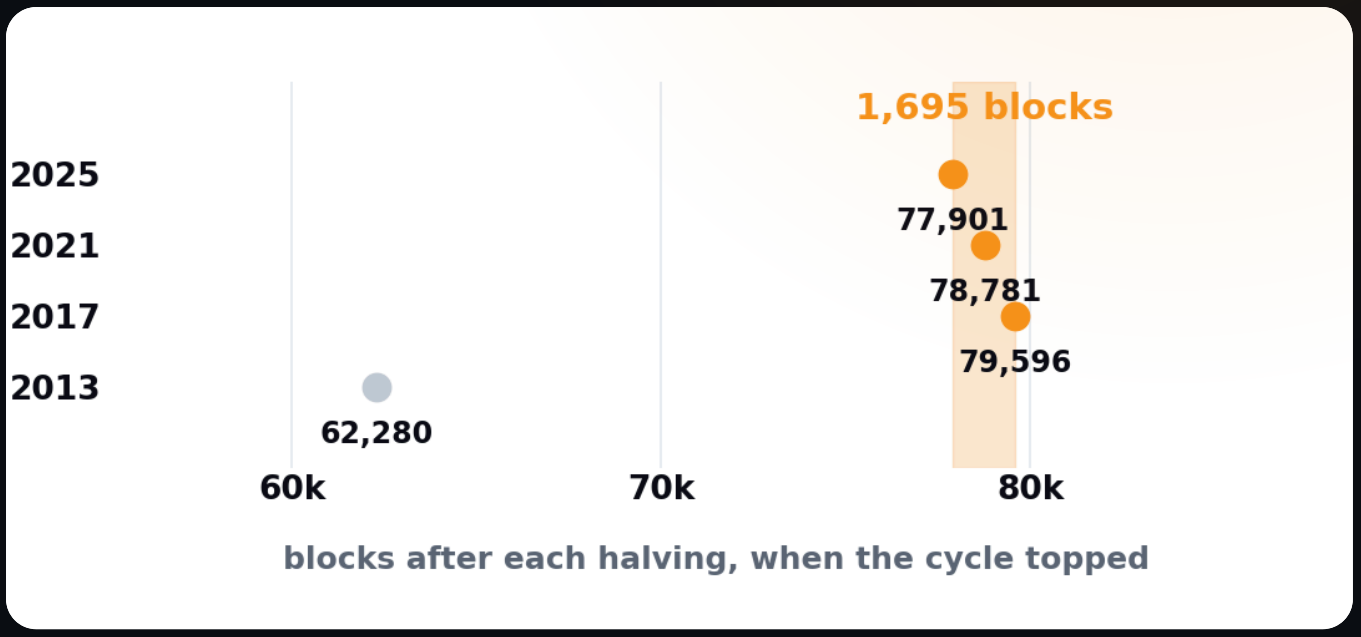

210,000
blocks between halvings, exactly

1,695
block spread across three cycle tops

0/10,000
random histories reproduce the top cluster

CHAIN TIP
962,455

INTO EPOCH 4
122,455

TO BOTTOM BAND
6,455

From the peer-reviewed research paper Bitcoin Runs on a Clock (Molnar, 2026). The full paper is free to read at arxiv.org/abs/2607.26188. Educational, not financial advice.

Bitcoin Jae
@bitcoin.daily

THE ORIGIN OF THE ERROR

WHERE THE FOUR CAME FROM

On 8 January 2009, five days after the genesis block, Satoshi Nakamoto described Bitcoin's monetary schedule to the Cryptography Mailing List. Coins would go to nodes that make blocks, **"with the amount cut in half every 4 years."**

That sentence is where the four-year cycle comes from. It is a plain-English summary written for an email list, and it is approximately right.

The code says something different. The rule is one constant, `SubsidyHalvingInterval = 210000`, fixed at block zero and adjustable by no one. Not four years. Not any number of years. A block count.

Four years is what 210,000 blocks works out to *if* the network produces a block every ten minutes on average. Do the arithmetic and it is not even four.

THE CONVERSION NOBODY CHECKS	
210,000 blocks at 10 minutes each	2,100,000 minutes
In days	1,458 days
In years	3.99 years, not 4

So the headline number was already rounded before the chain did anything unexpected. Then the chain runs fast or slow on top of that rounding, and the gap compounds. Every forecast built on the calendar version inherited an approximation from a mailing-list post and treated it as the rule.

THE WHOLE IDEA

The cycle is real. The unit is borrowed. Almost everything that looks messy about halving timing comes from that one substitution.

WHY THE CALENDAR BLURS

THE CHAIN DOES NOT RUN ON TIME

Bitcoin targets one block every ten minutes, which is 144 a day. It hits that target on average over the long run, because difficulty adjusts every 2,016 blocks to pull it back. But it drifts in between, and early on it drifted a lot.

Between 2011 and 2013 the network averaged roughly **170 blocks a day** against about **147** at maturity. Difficulty adjustments were chasing the arrival of the first serious mining hardware and could not keep up. **The chain was running about 15% fast.**

Why it drifts at all. Difficulty only re-tunes once every 2,016 blocks, roughly a fortnight. Between adjustments the network is running on a setting chosen for the hashrate that existed two weeks ago. If miners are switching on faster than that, every block in the gap arrives early, and the correction always lands after the damage.

In a mature network with stable hashrate that error is small and self-cancelling. In 2011, with the first purpose-built mining hardware arriving in waves, it ran in one direction for two years.

This is the mechanism behind everything else in this guide. If you measure a block-defined event in days, and the block rate drifts, your ruler stretches and shrinks underneath you. The cycle stays where it is. Your measurement of it moves.

The date-to-height map used in the paper is built from Coin Metrics daily block counts, cumulatively summed from genesis with no calendar gaps. It is validated two ways: all four protocol halving heights land within one day of their known dates, and eleven independent mempool.space lookups agree to a constant sampling offset that cancels in every range statistic.

THE RESULT

THE TOPS SNAP INTO FOCUS

Re-score the three mature cycle tops in blocks after their own halving and they land in a very narrow band.

CYCLE TOP	BLOCKS AFTER ITS HALVING
December 2017	**79,596**
November 2021	**78,781**
October 2025	**77,901**
Spread	1,695 blocks, about 12 days of mining

Heights from mempool.space point lookups. The paper's mechanical dating gives a 1,723-block spread, agreeing within end-of-day sampling.

The same three tops measured on the calendar spread across **21 days**. In blocks the spread is about 12 days' worth. Same three events, a third less scatter, purely from changing the unit.

How unlikely is that? The paper generates 10,000 synthetic Bitcoin histories by block-bootstrapping its actual daily returns, preserving volatility clustering and crash geometry, then runs the identical mechanical rule for finding tops. **Zero of the 10,000 reproduce the observed cluster.** Under a more forgiving construction that lets the null pick among candidate tops, 3 in 10,000 do.

AGAINST THE IMPOSTERS

BEATING THE LOOK-ALIKE CLOCKS

A four-year rhythm is not much of a finding on its own. Plenty of things happen every four years. So the halving is tested head to head against clocks that look just like it.

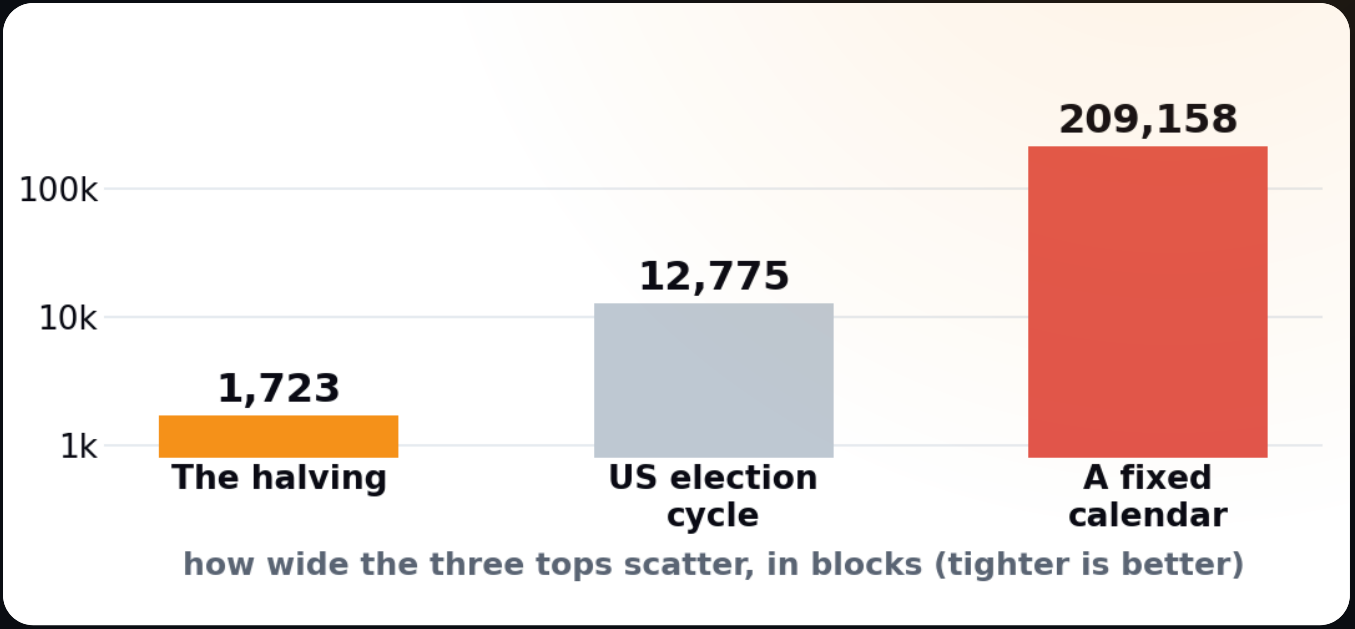

The halving holds the three tops inside **1,723 blocks**. The US election cycle, running on the same data, needs **12,775**. A fixed calendar period needs **209,158**, which is almost an entire epoch and means no alignment at all.

One result looks backwards until you see why. In days, a fixed four-year clock at a random starting phase never matches the halving's cluster, so the phase is carrying the signal. In blocks, a fixed 210,000-block clock at a random offset matches it 98.5% of the time.

That is not a failure, it is the point restated. In block space the halving is an *exact* period rather than a drifting one, so any clock with the right period lands on the tops regardless of where it starts. Draw the period at random instead and only 0.9% match. The period is the thing, and the protocol is the only thing that fixes it.

THE HARDER HALF

THE BOTTOMS, PARTLY RESCUED

The paper is careful about bottoms. Measured as a gap after the top, bear-market duration turns out to be largely intrinsic to how crashes behave, reproduced by about 40% of random histories. That statistic was demoted in the main paper and it stays demoted here.

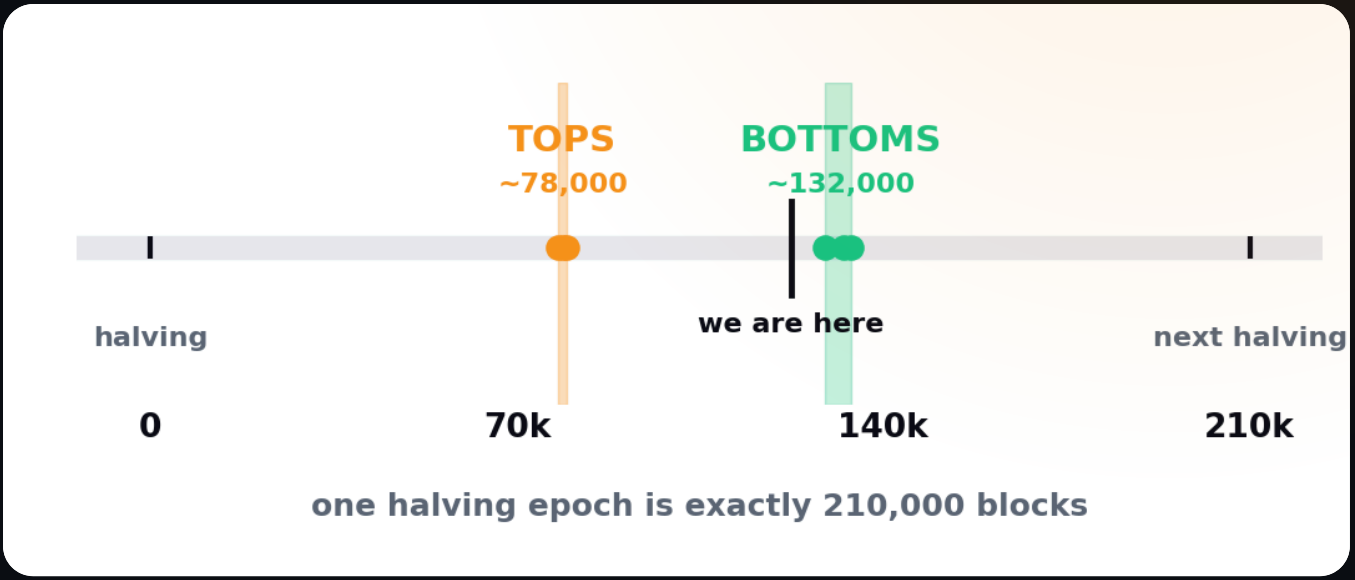

What the block coordinate adds is a phase statistic the calendar could not resolve. The three bottoms' blocks-after-halving spread is reproduced by **1.7%** of random paths in day phase and only **0.16%** in block phase. A tenfold sharpening, and it survives conditioning on cycle structure.

And an old argument dissolves. In days it was genuinely unclear whether bottoms anchored to the previous halving or the next one, and the two answers disagreed by months. In blocks the question cannot be asked, because blocks-since and blocks-until always sum to exactly 210,000. They are one statement written twice. The disagreement was the drifting calendar all along.

TOPS AGAINST BOTTOMS

HOW MUCH LOOSER IS THE BOTTOM

It is easy to say the bottoms are the weaker half and leave it there. Here is the actual size of the difference, measured the same way for both turns.

	BAND IN BLOCKS	ROUGHLY, IN MINING DAYS
Tops	1,695	12 days
Bottoms	5,024	35 days
Difference	3.0x wider	about three weeks more scatter

Band = the spread between the earliest and latest turn across three cycles, in blocks after each halving.

Three times wider. That is the honest size of the gap between how well this clock times tops and how well it times bottoms, and it is why the tops carry the paper's case while the bottoms are described as a weak prior worth watching.

It is still a band, and a 5,024-block band inside a 210,000-block epoch is a small target. But anyone treating a bottom call with the same confidence as a top call is ignoring a threefold difference in precision.

WHERE THAT LEAVES US

THIS EPOCH, IN NUMBERS

The current epoch opened at block 840,000 in April 2024. Add the historical bands and the two zones for this cycle are fixed, because the epoch length is fixed.

THIS EPOCH	BLOCK RANGE
Halving	840,000
Historical top band	917,901 to 919,596
Actual top, October 2025	917,901
Historical bottom band	968,910 to 973,934
Next halving	1,050,000

Bands are the observed range across the three prior cycles, not a forecast.

The October 2025 top landed at block `917,901`, inside the historical band and at its early edge. That is one more observation for the tops column, and it is the only part of this epoch that has already been settled.

A band is not a date. The chain reaches a given height when it reaches it, and the arrival time depends on hashrate, which nobody schedules. Treat these as zones to watch, not appointments.

THE ANOMALY

2013 STOPS BEING STRANGE

For years the 2013 top was the awkward one. It arrived early enough relative to its halving that it looked like it belonged to a different regime, and plenty of theories were built on treating it as an exception.

The 2013 top landed **62,280 blocks** after its halving, against a mature band of 77,901 to 79,596. Nobody is claiming it lands inside. The question is how far outside.

In days it sits at 0.693 of the mature average. In blocks it sits at **0.801**. Still early, but roughly a third of its apparent earliness was never the market at all. It was the 2011 to 2013 chain running 15% fast.

There is a robustness test hiding in this. Force the 2013 top into the mature set, which is the most hostile version of the analysis, and the null widens from zero to 8 per 10,000 in days but only to 3 per 10,000 in blocks. The native coordinate is more robust to the cycle that fits worst.

THE TAKEAWAY

Nothing about Bitcoin changed between those two numbers. Only the ruler did. That is the cleanest single demonstration of why the unit matters.

WHAT IT DOES NOT FIX

WHERE BLOCKS MAKE IT WORSE

A framework that improved everything it touched would be the more suspicious one. Here is the full list of what the block coordinate costs.

MEASURE	IN DAYS	IN BLOCKS
Cycle-shape overlay across cycles	0.72	0.69, worse
Volatility clock placebo	p = 0.19	p = 0.40, worse
Top-to-bottom lag sharpening	0.42%	1.11%, worse
Bear-market duration evidence	none	none, unchanged

Lower is better for the placebo and sharpening rows. The block clock loses on all four.

Read that top row carefully. The block clock times the cycle's discrete turns better than the calendar does. It does not organise the continuous path between those turns any better, and on volatility it is clearly worse.

And the sample is three cycles. Three tops and three bottoms. Every number in this guide rests on that, which is why the paper leans on nulls and placebo tests rather than on the tightness of the band alone. A tight band across three points is easy to produce by accident. A band that 10,000 synthetic histories cannot reproduce is harder.

Treat this as a strong prior about which unit to measure in, not as a machine that prints cycle dates. The honest claim is narrow. The halving keeps better time than the calendar does, and the calendar is what everyone has been using.

BITCOIN DAILY

COME FIND ME

This is the kind of breakdown I do every day. The hype brings the noise. I bring the data.

BD

bitcoin-daily.com
Every free guide, the live cycle charts, and the daily blog, all in one place
bitcoin-daily.com/guides

IG

Instagram
The live read every morning, daily charts, and the honest call on the market
@bitcoin.daily

YT

YouTube
Deep breakdowns like this one, every claim checked against the data
bitcoindailytv

X

X / Twitter
More data threads like this one
@BitcoinxDaily

Bitcoin Jae
@bitcoin.daily · Educational, not financial advice