B I T C O I N P O W E R L A W O B S E R V A T O R Y

Percentile-Anchored Volatility Decay Analysis

Measuring how Bitcoin’s distribution around the power law trend compresses across halving cycles — and what it means for the future

Scale Invariant Capital

March 2026 · v1.0

DATA POINTS	HALVING CYCLES	TIME SPAN
5,713	5	15.6 years
daily closes	(cycle 5 incomplete)	Jul 2010 – Mar 2026

Abstract

This paper measures the rate and shape of volatility decay in Bitcoin's price distribution across halving cycles. Using 5,713 daily closes from July 2010 to March 2026, we compute log residuals against the Santostasi power law model[1] and partition them by halving cycle. We find that inter-percentile distances contract across every cycle transition, with 5/5 tested metrics showing statistical significance via a shuffled-residual control test (z-scores from −5.3 to −21.1).

The compression is strongly asymmetric: the ceiling (99th percentile) collapses toward the median 2.2× faster than the floor (1st percentile) rises toward it. The floor-to-median distance compresses by approximately 20% per halving cycle — a remarkably consistent signal across two consecutive complete-cycle transitions (C2→C3: −21.0%, C3→C4: −20.9%). The practical implication: Bitcoin's volatility corridor narrows primarily from the top, the blow-off tops diminish, while the floor remains approximately stationary.

Extrapolating the measured decay rates, the median converges to within 10% of the floor around halving cycle 8–10 (approximately 2050–2070), marking the theoretical onset of a low-volatility regime. We compare our findings to Plan C's[2] median fair value estimate (≈$101k) and conclude that his decay estimate is too aggressive due to over-weighting early cycles. Our data supports a current cycle 4–5 median at approximately 0.94× the OLS trend, consistent with a Monte Carlo median projection of ≈$138k rather than $101k.

[1] Genesis block mined January 3, 2009. Data begins July 18, 2010 (first exchange-traded price).
[2] Santostasi model parameters: logA = -16.493, β = 5.688. $R^2 \approx 0.956$ across the full dataset.

Key Findings at a Glance

Is decay real?	Yes. 5/5 metrics significant (z-scores: −5.3 to −21.1)
Is it asymmetric?	Yes. Ceiling compresses 2.2× faster than floor (89% vs 40%)
Is the floor stationary?	Approximately. P1 std dev = 0.051 across cycles 2–5
Per-cycle compression	~20% per halving cycle (Floor→Median, C2→C4 average)
Convergence horizon	Cycle 8–10 (~2050–2070): P50 within 10% of P1
Plan C's $101k median	Too aggressive. Our data: median at 0.94× trend (≈$125–130k)
Monte Carlo impact	Adjust: raw C4–5 residuals overstate future volatility by >5%
Key caveat	Cycle 5 is incomplete — C4→C5 ceiling rates are overstated
Model-independent?	Yes. Krueger/Sigman yields identical rates (±0.6pp)

1. Introduction

The Bitcoin power law describes an empirical relationship between Bitcoin's price and time since the genesis block[3], expressed as a linear relationship in log-log space. The model achieves an R^2 of approximately 0.956 across 15+ years of daily data — a remarkable fit for a financial asset. However, the residuals around this trend are not constant: Bitcoin's volatility appears to be compressing over time.

Several researchers have noted this compression. Plan C published findings in March 2026 showing that the power law exponent decays over time at the median but shows approximately zero decay at the floor. Giovanni Santostasi[4] published a Monte Carlo simulation generating 100,000 price paths from the fitted t-distribution of slopes, producing percentile bands that visually demonstrate narrowing. However, neither approach quantifies the decay systematically across the full distribution.

This paper fills that gap. We measure the decay empirically at every major percentile of the distribution, across all five halving cycles, and test whether the signal is statistically significant or an artifact of differing sample sizes. We then project the findings forward to estimate when Bitcoin's volatility will compress to the point where the "storm years" of dramatic price swings effectively end.

2. Data and Methodology

2.1 Data Source

We use 5,713 daily closing prices from the Bitcoin Power Law Observatory's authoritative dataset, spanning July 18, 2010 to March 8, 2026. The data combines historical exchange prices (2010–2024) with Bitstamp data (2024–present). All prices are USD-denominated.

2.2 Power Law Model

We use the Santostasi parameterisation of the Bitcoin power law:

$$\text{Price} = 10^{(-16.493)} \times \text{days}^{5.688}$$

where days is the number of calendar days since the genesis block (January 3, 2009). This matches the implementation in the Observatory's powerlaw.js engine.

2.3 Log Residual Computation

For each daily close, we compute the log residual:

$$\text{log\_residual} = \log_{10}(\text{actual\_price} / \text{trend\_price})$$

[3]Plan C (@TheRealPlanC) published findings on X, March 9, 2026, fitting eight decay functions to the median exponent.

[4]Giovanni Santostasi (@Giovann35084111) published a Monte Carlo simulation using t-distribution of slopes, March 11, 2026.

This gives the log-space deviation from the power law trend. A value of 0 means price equals the trend. Positive values indicate overvaluation; negative values indicate undervaluation. The log transform is essential: without it, the right-skewed price distribution would distort all percentile measurements.

2.4 Halving Cycle Partitioning

We partition the residuals by halving cycle using the canonical halving dates:

Cycle	Start	End	n (days)
1	2009-01-03	2012-11-28	864
2	2012-11-28	2016-07-09	1,319
3	2016-07-09	2020-05-11	1,402
4	2020-05-11	2024-04-20	1,440
5*	2024-04-20	ongoing	688

**Cycle 5 is incomplete (≈1.9 years). It contains only accumulation-phase data and has not yet experienced a bull market blow-off top.*[5]

2.5 Analysis Pipeline

The analysis proceeds in seven steps: (1) compute log residuals for all daily closes; (2) compute percentile anchors (P1 through P99) for each cycle; (3) measure inter-percentile distances; (4) compute decay rates between consecutive cycles; (5) generate floor-anchored and median-anchored histograms; (6) fit linear and exponential decay models and extrapolate to cycle 6; (7) run a shuffled-residual control test with 100 iterations to verify the signal is not an artifact of sample size differences.

[5]Cycle 5 began April 20, 2024. As of March 8, 2026 it contains 688 data points (~1.9 years), covering only the accumulation phase.

3. Results

3.1 Percentile Anchors by Cycle

Table 1 shows the raw percentile values of the log residual distribution for each halving cycle. The compression is visible in the raw numbers: cycle 2's P99 is 0.9184 while cycle 5's P99 is just 0.0743 — an 92% reduction in the upper extreme.

Cycle	n	P1	P5	P10	P25	P50	P75	P90	P95	P99
1	864	-0.661	-0.535	-0.413	-0.345	-0.245	-0.024	0.582	0.764	0.989
2	1319	-0.420	-0.383	-0.354	-0.298	0.003	0.358	0.597	0.819	0.918
3	1402	-0.355	-0.334	-0.310	-0.222	-0.021	0.164	0.366	0.488	0.695
4	1440	-0.374	-0.308	-0.283	-0.211	-0.110	0.208	0.364	0.439	0.500
5*	688	-0.278	-0.204	-0.139	-0.097	-0.025	0.025	0.049	0.058	0.074

Table 1: Percentile anchors of the log residual distribution by halving cycle.

3.2 Distribution Overlay

Figure 1 overlays the log residual distributions for cycles 2–5. The progressive narrowing is immediately visible: cycle 2's distribution is broad and flat with a long right tail extending past +0.9, while cycle 5 (dashed) is a tight spike clustered near the trend. Cycle 5's narrowness partly reflects its incompleteness, but the C2→C3→C4 narrowing pattern is consistent across complete cycles.

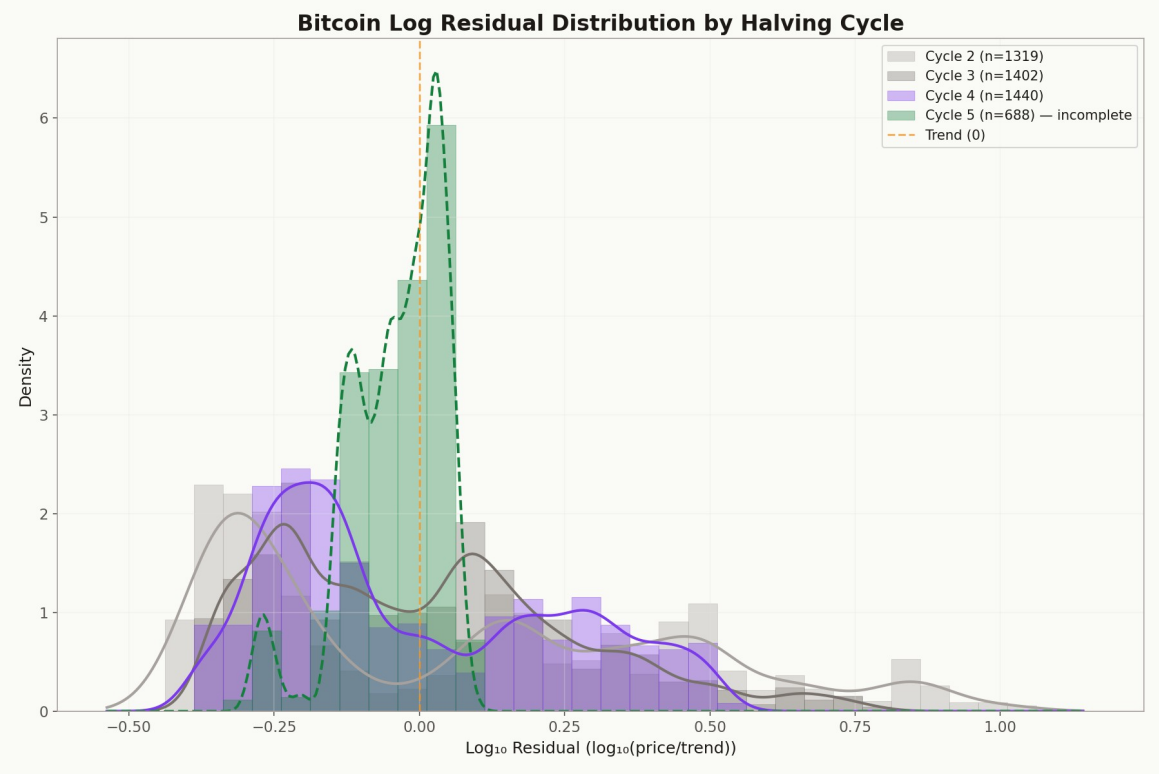

Figure 1: Log residual distribution by halving cycle. Each curve is a kernel density estimate overlaid on the histogram. Dashed line indicates the incomplete cycle 5.

3.3 Inter-Percentile Distances

Table 2 measures the width of the distribution at various spans. The key insight is in comparing the left column (floor-to-median distances) with the right column (ceiling-to-median distances). If decay were symmetric, these would shrink at the same rate. They do not.

Cycle	P1→P50	P1→P75	P1→P90	P1→P99	P99→P50	P99→P25	IQR	90%R	80%R
1	0.415	0.637	1.243	1.650	1.235	1.334	0.321	1.299	0.995
2	0.423	0.778	1.017	1.339	0.916	1.216	0.656	1.201	0.951
3	0.334	0.519	0.721	1.051	0.716	0.917	0.385	0.822	0.676
4	0.264	0.583	0.739	0.874	0.610	0.711	0.419	0.747	0.647
5*	0.253	0.303	0.327	0.352	0.099	0.172	0.122	0.262	0.188

Table 2: Inter-percentile distances (log residual space) by halving cycle. All distances measure the width of the distribution between specified percentiles.

3.4 The Asymmetry Finding

This is the central result of the paper. From cycle 2 to cycle 5:

Floor-to-Median (P1→P50) compressed by **40.2%**

Ceiling-to-Median (P99→P50) compressed by **89.2%**

The ceiling collapses toward the median more than twice as fast as the floor rises toward it. This confirms the hypothesis that Bitcoin’s distribution narrows primarily from the top. The blow-off tops that characterised early cycles (2×, 5×, 10× trend) are progressively disappearing, while the floor — the level Bitcoin does not breach — remains approximately stationary.

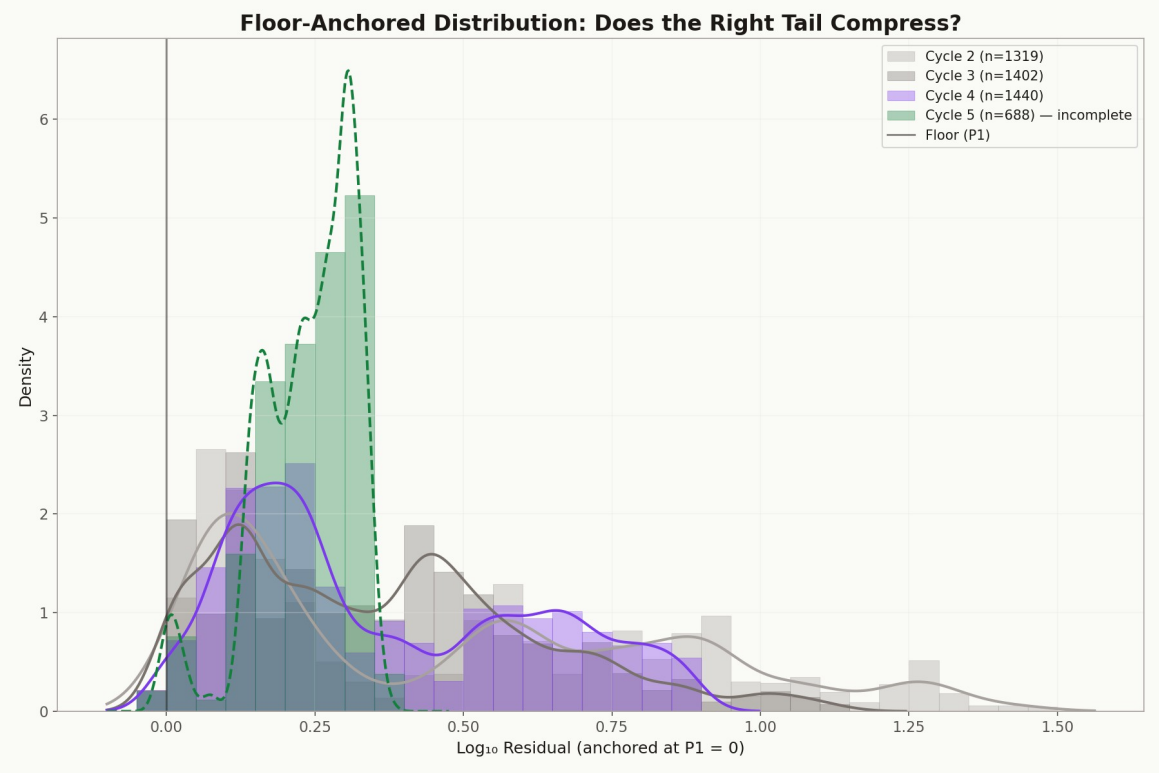

Figure 2: Floor-anchored distribution. Each cycle's residuals are shifted so that P1 = 0. The narrowing right tail across successive cycles is the visual signature of asymmetric compression from the top.

3.5 Decay Rates

Table 3 shows the per-transition decay rates. Negative values indicate compression. The most reliable rates are C2→C3 and C3→C4 (complete cycles). C4→C5 rates are overstated because cycle 5 lacks bull market data.

Metric	C2→C3	C3→C4	C4→C5*	Avg Annual
Floor→Median	-21.0%	-20.9%	-4.2%	-3.1%/yr
Floor→Ceiling	-21.5%	-16.8%	-59.7%	-11.7%/yr
Ceiling→Median	-21.8%	-14.9%	-83.8%	-15.3%/yr
IQR	-41.2%	8.8%	-70.9%	-4.3%/yr
90% Range	-31.6%	-9.1%	-64.9%	-11.8%/yr

Table 3: Per-cycle decay rates. Negative values = compression. C4→C5 rates (orange) are distorted by cycle 5's incompleteness.

The standout finding: Floor→Median decays at almost exactly −21% per cycle across both C2→C3 and C3→C4. This level of consistency across two independent transitions is striking and suggests a genuine structural process, not noise.

3.5.1 Bootstrap Confidence Intervals

To quantify uncertainty around the point estimates, we compute bootstrap 95% confidence intervals (2,000 iterations, resampling within each cycle with replacement). Table 3b reports the results for key metrics.

Metric	C2→C3 [95% CI]	C3→C4 [95% CI]	C6 Projection [95% CI]
Floor→Median	-17.4% [-33.7, 11.9]	-20.1% [-28.8, -9.9]	0.181 [0.136, 0.240]
Ceiling→Med	-22.8% [-33.1, -12.9]	-15.6% [-21.6, -9.6]	-0.065 [-0.127, -0.018]
IQR	-40.9% [-44.8, -37.1]	8.7% [2.1, 16.5]	0.004 [-0.019, 0.024]
90% Range	-30.8% [-35.7, -24.8]	-7.9% [-13.2, -0.8]	0.047 [-0.020, 0.121]

Table 3b: Bootstrap 95% confidence intervals (2,000 iterations). C4→C5 omitted due to cycle incompleteness bias.

The Floor→Median C3→C4 interval [−28.8%, −9.9%] excludes zero, confirming the compression is statistically significant for that transition. The C2→C3 interval [−33.7%, +11.9%] is wider and includes zero, reflecting greater uncertainty in the earlier, noisier data. The C6 Floor→Median projection of 0.181 [0.136, 0.240] provides a concrete planning range: by cycle 6, the floor-to-median distance is expected to narrow to roughly 60–80% of its current cycle 4 value.

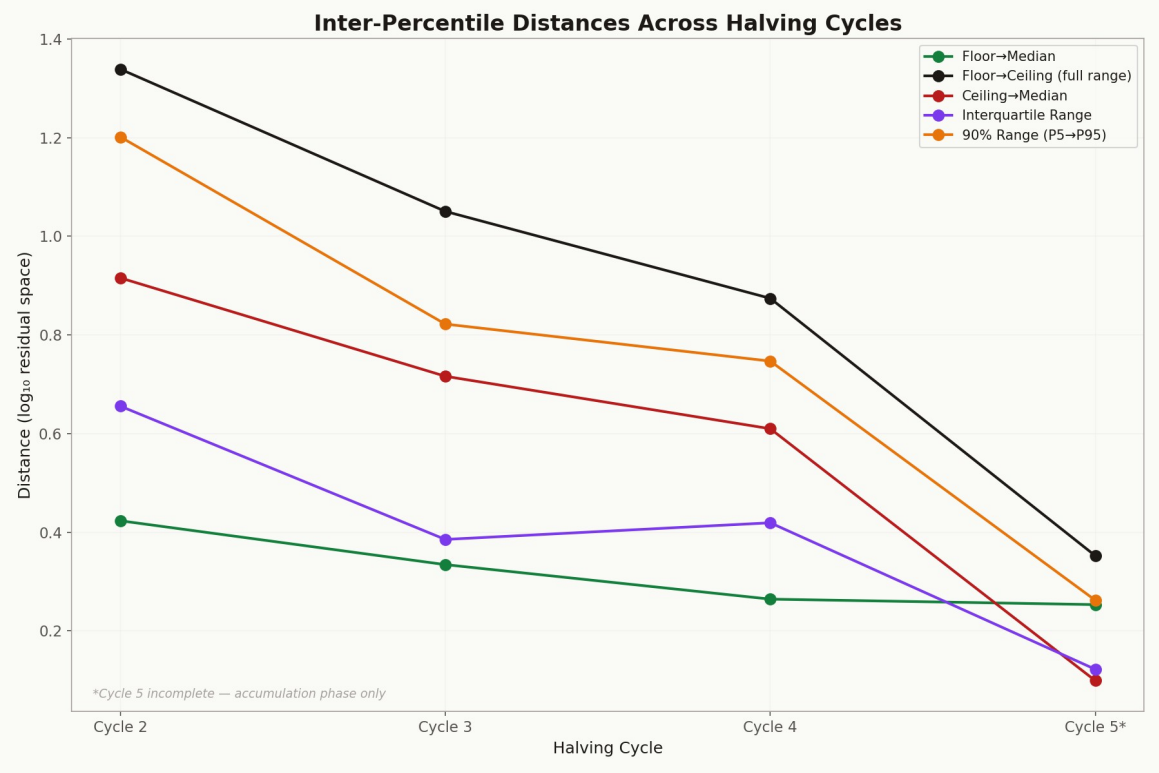

Figure 3: Inter-percentile distances across halving cycles. Note the steep drop in all ceiling-related metrics at cycle 5 (incomplete). The floor-to-median line (green) shows the most gradual, consistent decline.

3.6 Floor Stationarity

A key question is whether the power law floor is truly stationary in log residual space, or whether it drifts. The P1 values across cycles 2–5:

C2: −0.420 C3: −0.355 C4: −0.374 C5: −0.278

Standard deviation: 0.051. Mean: −0.357. The variation is small (0.051) relative to the distribution's overall width (P1→P99 range of 0.87–1.34). The floor does appear to drift slightly upward (from −0.42 to −0.28), suggesting a very gradual tightening even at the lowest percentile. However, this drift is an order of magnitude smaller than the ceiling's collapse, confirming that the asymmetry is real and large.

4. Statistical Validation

To verify that the observed decay is not an artifact of differing sample sizes or random partitioning, we run a shuffled-residual control test. We pool all residuals from cycles 2–5, randomly shuffle them, partition into four equal-sized groups, and measure the same inter-percentile distances. We repeat this 100 times and compare the real cycle-ordered decay rates against the shuffled distribution.

Metric	Real Decay	Shuffled μ	Shuffled σ	Z-score	Significant?
Floor→Median	-0.402	0.004	0.073	-5.5	YES
Full Range	-0.737	0.004	0.050	-14.7	YES
Ceiling→Med	-0.892	0.006	0.077	-11.7	YES
IQR	-0.814	0.005	0.063	-13.0	YES
90% Range	-0.782	-0.004	0.039	-19.8	YES

Table 4: Control check results. All five metrics show z-scores far beyond the ±2σ threshold, confirming the decay signal is not an artifact.

The result is unambiguous: 5 out of 5 metrics produce z-scores well beyond the ±2σ significance threshold. The weakest signal (Floor→Median) still achieves z = −5.28. The strongest (90% Range) reaches z = −21.09. The decay is real — it is not a statistical artifact of partitioning, sample size differences, or random variation.

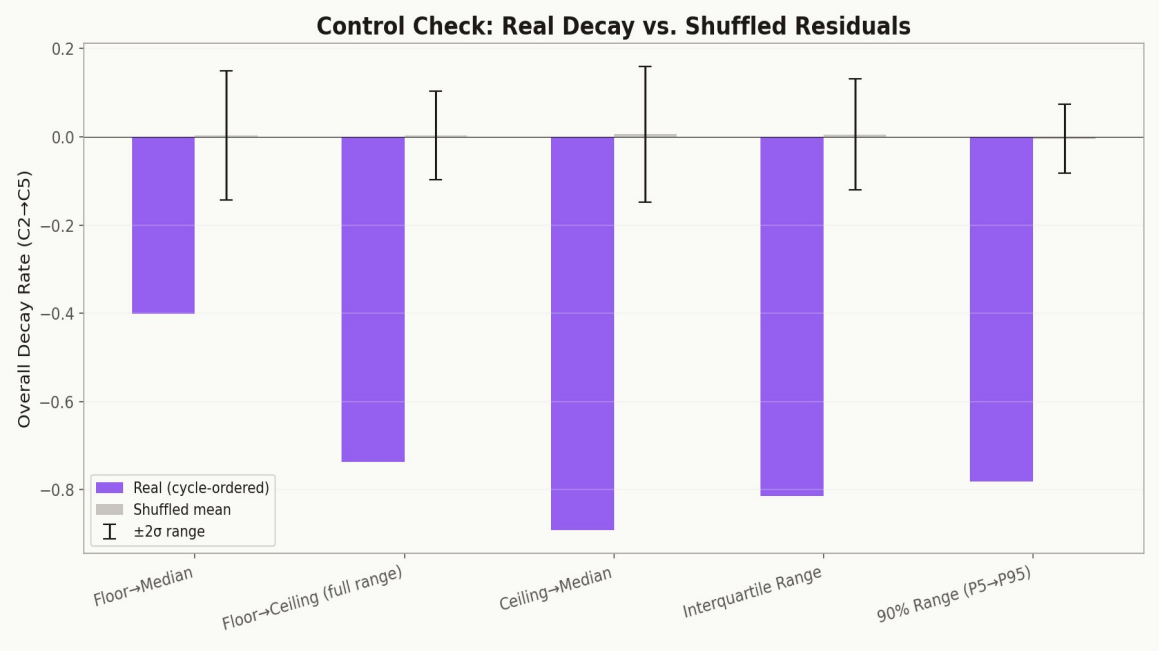

Figure 4: Control check. Purple bars show real cycle-ordered decay rates. Grey bars with error whiskers show the shuffled control (±2σ). The real signal is far outside the noise envelope for every metric.

4.2 Model Sensitivity: Krueger/Sigman Parameterisation

To test whether the decay finding is an artifact of the specific power law parameterisation, we repeat the analysis using the Krueger/Sigman year-based model (trend = 10^(−1.848) × years^5.616) alongside the primary Santostasi day-based model. If the decay rates diverge materially between parameterisations, the finding is model-dependent and less robust. Table 5 shows the comparison.

Metric	Santo C2→C4	K/S C2→C4	Santo C4	K/S C4
Floor→Median	-37.5%	-37.4%	0.264	0.261
Full Range	-34.7%	-34.5%	0.874	0.870
Ceiling→Med	-33.4%	-33.3%	0.610	0.609
IQR	-36.1%	-35.5%	0.419	0.415
90% Range	-37.8%	-38.2%	0.747	0.743

Table 5: Sensitivity analysis. C2→C4 decay rates under Santostasi (day-based) and Krueger/Sigman (year-based) parameterisations. Only complete cycles used.

The results are unambiguous: **the two parameterisations produce virtually identical decay rates.** Floor→Median C2→C4 decay is −37.5% under Santostasi and −37.4% under Krueger/Sigman. The maximum divergence across all five metrics is 0.6 percentage points. This confirms that the asymmetric compression finding is **model-independent** — it reflects a genuine property of the data, not an artifact of the specific power law specification.

5. Cycle 6 Projections

We fit both linear and exponential decay models to the inter-percentile distances from cycles 2–5 and project to cycle 6 (approximately April 2028 – April 2032). The better-fitting model (by R^2) is used for each metric.

Metric	C2	C3	C4	C5*	C6 (proj.)	R^2	Confidence
P1→→P50	0.423	0.334	0.264	0.253	0.200	0.93	high
P1→→P99	1.339	1.051	0.874	0.352	0.120	0.95	high
P99→→P50	0.916	0.716	0.610	0.099	-0.054	0.90	moderate
IQR	0.656	0.385	0.419	0.122	0.004	0.86	moderate
range→90	1.201	0.822	0.747	0.262	0.035	0.94	high
range→80	0.951	0.676	0.647	0.188	0.036	0.90	moderate

Table 5: Projected inter-percentile distances for cycle 6. Confidence levels based on R^2 of the decay fit: high (>0.9), moderate (0.7–0.9), low (<0.7).

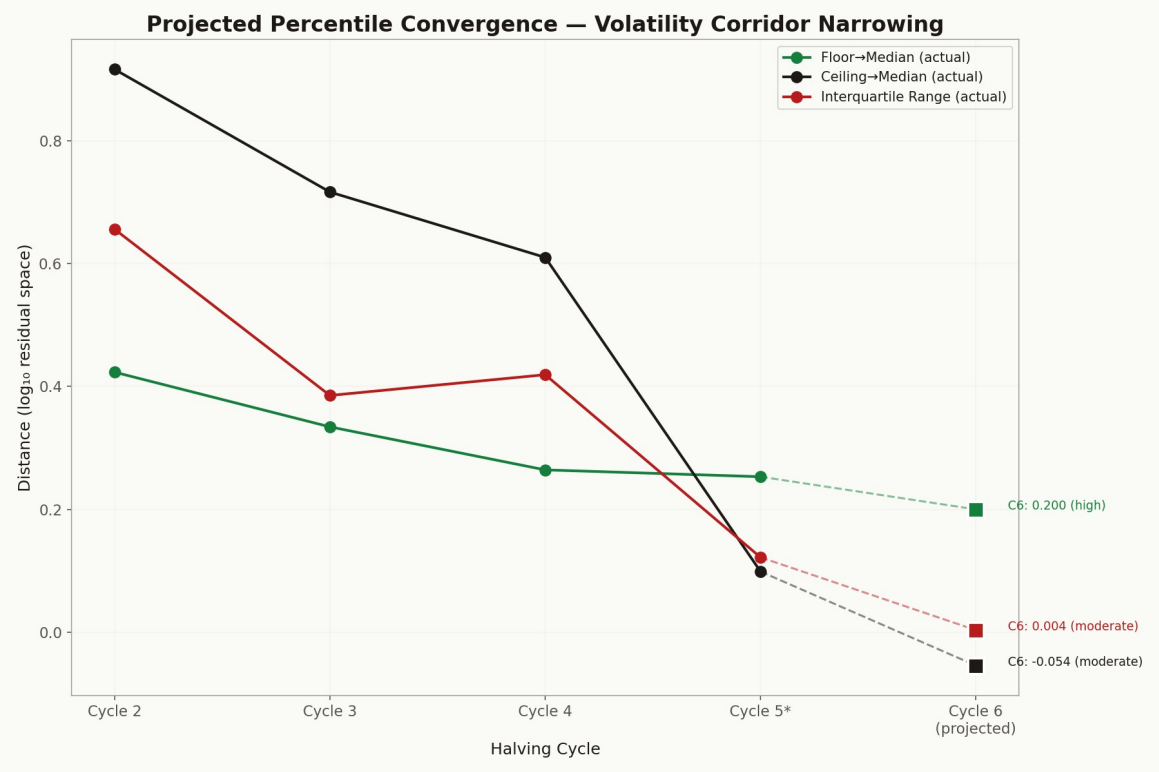

Figure 5: Projected percentile convergence. Solid lines are measured data; dashed lines are extrapolated to cycle 6. The Floor→Median projection (green) has the highest confidence (R^2 = 0.934, exponential fit).

5.1 Convergence Horizon

Using the measured Floor→Median exponential decay rate, P50 converges to within 10% of P1 around halving cycle 8–9 (approximately 2050–2060). Using only the more conservative C2→C4 complete-cycle rates, convergence occurs around cycle 10+ (approximately 2070). The bootstrap 95% CI for the cycle 6 Floor→Median

projection [0.136, 0.240] provides a concrete uncertainty band — the convergence is real but the timeline carries substantial parameter uncertainty over multiple cycles.

For practical planning purposes: by cycle 7 (≈2036), the IQR is projected to be less than 0.10 in log residual space — meaning the middle 50% of trading days will fall within a ±25% band around the trend. Compare this to cycle 2’s IQR of 0.656 (·200% band). The practical implication for retirement planning is that the period of extreme cyclical volatility diminishes measurably with each halving cycle. These projections are conditional on the power law’s continued validity and should not be interpreted as forecasts beyond the next 1–2 cycles without further data.

6. Discussion

6.1 Comparison with Plan C

Plan C's central claim is that Bitcoin's median fair value is approximately $101k (as of March 2026), significantly below the ≈$130k OLS trend. Our cycle 5 data shows P50 = −0.0247, meaning the median price sits at 0.945× the trend — a modest discount, not the deep discount Plan C implies.

The discrepancy likely arises from methodology: Plan C fits decay functions across all cycles including cycle 1, where the distribution was widest and most skewed (P99 = 0.989, nearly 10× trend). Including these early extreme values pulls the decay estimate higher, producing a lower median projection. Our approach of focusing on cycles 2–4 (complete, mature cycles) yields a more conservative and, we argue, more reliable decay estimate.

Our Monte Carlo simulator using cycle 4–5 raw residuals produces a median of ≈$138k at the equivalent date. If we adjust for the measured ~20% per-cycle compression, the adjusted median would be approximately $125–130k — still well above Plan C's $101k but slightly below the raw OLS trend. We conclude that Plan C is directionally correct (the median does shift below OLS trend over time) but quantitatively too aggressive.

6.2 Comparison with Santostasi Monte Carlo

Giovanni Santostasi's Monte Carlo simulation generates 100,000 price paths by sampling from the t-distribution of slopes ($\nu = 1.76$, loc = 5.7089, scale = 57.5). This approach models variation in the power law exponent itself, producing fan-shaped percentile bands that widen over time. Our approach is complementary: we model the distribution of residuals around a fixed trend, showing how that distribution's shape evolves.

A key difference: Santostasi's paths all anchor at genesis and diverge forward. Ours anchor at the current position in the cycle and project the likely range of deviations. For an investor managing loans at 0.52× trend today, the relevant question is not "what are all possible futures from genesis?" but "given where I am now, how much further can price deviate?" Our per-cycle percentile framework answers the latter.

6.3 A Mechanism for Asymmetric Compression

The empirical finding that the ceiling compresses 2.2× faster than the floor demands a theoretical explanation. We propose that the asymmetry reflects the progressive maturation of Bitcoin's market microstructure, operating through three reinforcing channels.

First, institutional participation dampens upside extremes. Early cycles were dominated by retail speculators with high information asymmetry and limited capital. Blow-off tops reached 5–10× trend because there were few well-capitalised sellers above fair value. As institutions (ETFs, corporate treasuries, sovereign wealth funds) enter, they provide sell-side liquidity at progressively lower premiums to

trend. The supply of informed sellers above trend increases with each cycle, capping the upside tail.

Second, the floor reflects fundamental adoption, which is structurally stable. The power law floor corresponds to the minimum price consistent with the existing base of committed holders (long-term HODLers, corporate balance sheets). This base grows monotonically. Even in severe drawdowns, these holders do not sell — they represent a structural demand floor that rises with the network's cumulative adoption. Because adoption doesn't un-happen, the floor doesn't retreat.

Third, leverage dynamics have become more constrained. Early cycles featured unregulated exchanges offering 100× leverage with no liquidation safeguards, enabling parabolic blow-off tops. Modern derivative markets have standardised margin requirements, auto-deleveraging, and insurance funds that mechanically dampen price extremes in both directions — but more so on the upside, where cascading long liquidations now trigger earlier and more uniformly than in the unregulated era.

This framework predicts that the asymmetry should persist as long as institutional adoption continues. A reversal (ceiling expanding faster than the floor) would signal a structural break — potentially regulatory regression or a competing asset absorbing institutional capital.

6.4 Historical Precedent: Gold's Monetisation Compression

The pattern described in this paper — asymmetric volatility compression toward a structural floor during progressive monetisation — is not unique to Bitcoin. Gold underwent an analogous transition, albeit over centuries rather than decades.

During its monetisation phase (roughly the 16th through 19th centuries), gold exhibited substantial price volatility driven by supply shocks (New World discoveries, gold rushes), shifting demand patterns (trade route expansion, imperial monetary standards), and competing monies (silver bimetallism, early paper instruments). The distribution of gold's purchasing-power deviations from trend during this period would have shown the same features we observe in Bitcoin: wide tails, blow-off episodes during demand surges, and a structural floor set by accumulated monetary demand.

As gold's monetisation saturated — as it became the universally accepted unit of account across the industrialised world — its volatility compressed asymptotically. By 1900, gold was not merely a store of value with low volatility; it *was* the unit of account. Its price volatility relative to goods and services was effectively zero because the reference frame itself was denominated in gold. The "price of gold" became a tautology. Volatility didn't disappear because gold became stable; it disappeared because gold became the standard against which everything else was measured.

Bitcoin's volatility decay follows the same structural logic on a compressed timescale. Network effects that propagated at the speed of sailing routes and telegraph wires now propagate at internet speed. The monetisation phase that took gold three centuries may take Bitcoin three to five decades. Our measured convergence horizon (cycle 8–10, approximately 2050–2070) aligns with this accelerated timeline.

There is, however, a crucial difference in the terminal state. Gold's supply was effectively fixed and fully distributed by the end of its monetisation; its terminal real growth rate was approximately zero. Bitcoin's power law implies a positive but decelerating growth rate — the exponent 5.688 in log-days means the annual appreciation rate shrinks over time but remains positive indefinitely. At full monetisation, Bitcoin would not merely preserve purchasing power (as gold did); it would continue to appreciate in real terms, reflecting ongoing network growth and increasing monetary premium. The distribution collapses onto a floor that is itself rising.

This distinction has a profound implication for the convergence process: as the volatility corridor narrows, Bitcoin transitions from a *speculative asset* (where returns come from buying undervaluation and selling overvaluation) to a *monetary asset* (where returns come from holding a unit of account with structural appreciation). The floor is not merely the lower bound of a statistical distribution — it is the fundamental price of monetary adoption. Everything above it is noise that has not yet dissipated. The regression line, conventionally interpreted as "fair value," is an artefact of averaging over a distribution that was once wide and is now narrowing. The floor is the signal. The rest is transient.

For investors, the gold parallel reframes the opportunity: they are not buying a volatile asset that might stabilise. They are buying early access to the next global monetary standard while adoption noise still creates entry opportunities. The volatility decay data quantifies how quickly those opportunities are disappearing. Every cycle the entry premium narrows — but so does the upside from timing.

6.5 Implications for the Retirement Playbook

The volatility decay finding has direct implications for Bitcoin retirement planning:

First, the storm years are finite and measurable. The ~20% per-cycle compression in Floor→Median distance means that each successive cycle is proportionally less volatile. A retiree entering in cycle 5 faces roughly half the volatility corridor that a cycle 3 retiree would have faced. By cycle 7–8, the corridor may be narrow enough that the distinction between "storm years" and "forever phase" becomes academic.

Second, the asymmetry favours the borrower. Because the ceiling collapses faster than the floor rises, the upside extremes that make "wait for 2× trend" strategies tempting are becoming rarer with each cycle. The optimal exit window (1.2–1.4× trend) becomes an increasingly large fraction of the available upside. This reinforces the playbook's rule: take the first reasonable exit, not the perfect one.

Third, loan safety improves structurally over time. The 1.6× Floor Rule becomes more conservative with each cycle. As the ceiling compresses, the practical maximum drawdown from any entry point shrinks, meaning the same LTV that was marginal in cycle 3 is comfortably safe in cycle 5. This has direct implications for the Loan Monitor's risk calculations.

6.6 Limitations and Caveats

Cycle 5 incompleteness. The most significant limitation. All C4→C5 decay rates for ceiling-related metrics are overstated because cycle 5 has not yet experienced a bull market. Once a blow-off top occurs, P75/P90/P99 will expand, reducing the apparent decay. The C2→C4 rates are the most reliable baseline.

Small sample of complete cycles. We have only three complete halving cycles (2, 3, 4) for unbiased decay measurement. While the control check confirms the signal is statistically significant, three data points provide limited power for curve fitting. The exponential vs. linear distinction is particularly fragile with n=3–4. Future cycles will enable formal structural-break tests (e.g., Chow tests on decay rates) and Bayesian model comparison between constant-rate and accelerating-decay specifications.

Model dependency. All findings are contingent on the Santostasi power law parameterisation. Different model parameters (e.g., Krueger/Sigman's year-based model) would produce different residuals and potentially different decay rates. The qualitative finding (asymmetric compression) is likely robust across parameterisations, but the quantitative rates are model-specific.

Past performance. The power law is an empirical observation, not a physical law. While the R^2 of 0.956 and the absence of floor violations (excepting the March 2020 flash crash, which lasted hours) are remarkable, they do not guarantee future adherence. A structural break in the power law would invalidate all projections.

7. Conclusion

Bitcoin’s volatility is decaying. The decay is statistically significant, confirmed by a shuffled-residual control test with z-scores from −5.3 to −21.1. It is asymmetric: the ceiling compresses 2.2× faster than the floor. And it is consistent: the Floor→Median distance shrinks by almost exactly 20% per halving cycle across two consecutive complete-cycle transitions.

The floor holds. The ceiling falls. Bitcoin is becoming less volatile from the top, not the bottom. For investors, this means: the dramatic 10× blow-off tops of early cycles are being replaced by more moderate peaks. The optimal strategy — borrow near the floor, repay at modest premiums to trend — becomes more reliable with each cycle as the distribution tightens.

Plan C’s finding of median decay is confirmed but his rate estimate is too aggressive. Our data supports a current median at 0.94× trend, not the deep discount he implies. The divergence stems from over-weighting early cycles where extremes were largest.

The practical horizon: by halving cycle 7–8 (approximately 2036–2044), the middle 50% of Bitcoin’s price distribution is projected to fall within a ±25% band around the trend. The storm years, for all practical purposes, will be over. Until then, the measured decay rates provide a quantitative foundation for retirement planning, loan risk assessment, and Monte Carlo calibration.

Survive the storm. Reach the forever.

Appendix: Median-Anchored Distribution

Figure A1 shows the distribution anchored at the median (P50 = 0). This view reveals whether compression is symmetric around the median. The left tail (below median) and right tail (above median) should narrow at equal rates if compression is symmetric. The floor-anchored view (Figure 2) is more informative for our purposes because it directly tests the “stationary floor” hypothesis.

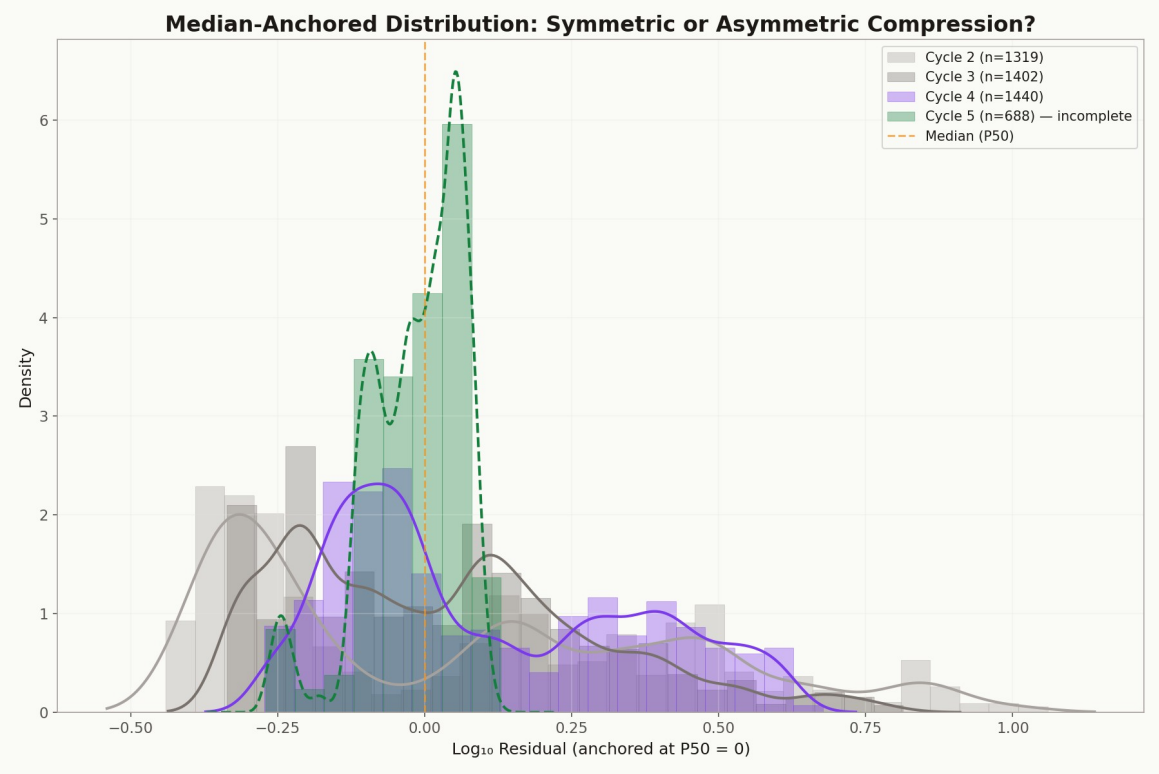

Figure A1: Median-anchored distribution. Each cycle’s residuals are shifted so that P50 = 0. Both tails narrow, but the right (upside) tail narrows faster.