

The Case for the Floor

Seven Independent Attacks on Bitcoin's Power Law Minimum

Scale Invariant Capital

BTC Power Law Observatory. Paper 7.

March 2026

"The floor is not a line. It is a process. The ceiling falls. The floor rises. The trend drifts down. Everything converges."

Abstract

Bitcoin's price follows a power law in time (R-squared = 0.92, slope = 5.93, 4,198 daily observations post-2014; R-squared = 0.95+ over the full 5,713-day dataset). Every tool in the Observatory -- the retirement calculator, the loan monitor, the floor bonds framework, the Bitcoin Floor Rate -- rests on the assumption that a structural lower boundary exists. This paper attempts to falsify that assumption through seven independent attacks: empirical, comparative, mechanistic, statistical, thermodynamic, game-theoretic, and survivorship.

We find that the floor is not a fixed line but a regime-dependent boundary that tightens predictably across halving cycles. The 1st-percentile residual rose from 0.22x trend (Cycle 1) to 0.53x trend (Cycle 5). Out-of-sample testing shows 37% fewer floor breaches than a random model predicts. Cross-asset comparison using identical methodology on seven assets reveals that Bitcoin's floor grows at 39.8%/year versus gold's 1.8%/year -- 22x faster. No published systematic cross-asset power law comparison exists in the literature; this is the first.

We address the circularity objection ("the floor is -2 sigma, and -2 sigma is defined as what doesn't get breached") through three independent resolution methods. We identify a three-way convergence: the ceiling falls, the floor rises, and the OLS trend drifts downward toward a decay-adjusted median attractor near the 41st percentile. This contradicts PlanC's March 2026 finding that the 1st quantile shows zero decay. Our per-cycle data shows the floor is rising.

The verdict: the floor is real but not static. It converges. Each cycle's worst case is the next cycle's comfortable margin.

Section 1: Literature Review

This paper builds on a body of research spanning 2014-2026, comprehensively cataloged in our companion paper "Bitcoin Power Law Research: A Comprehensive Survey" (Scale Invariant Capital, Paper 6, March 2026). We summarize the key contributions relevant to the floor.

1.1 The Power Law Model

The observation that Bitcoin's price follows a power law in time originates with Trolololo (BitcoinTalk, October 2014). Giovanni Santostasi (PhD Astrophysics, Northwestern) independently identified the relationship and formalized it in September 2018 with $\text{Price} = 10^{-17} \times \text{days}^{5.8}$, R-squared = 0.92 improving to 0.95 by 2024. Harold Christopher Burger (hcburger.com, September 2019) independently published the "power-law corridor of growth" with slope 5.845. The exponent has been measured by at least ten independent researchers, clustering between 5.6 and 5.9.

1.2 Three Floor Definitions in Circulation

Definition A: Fixed multiplier. The Observatory's powerlaw.js uses 0.42x trend. BGeometrics uses the same. This is approximately the 3rd percentile and has been breached 77 times. **Definition B: Percentile bands.** Mezinskis (Porkopolis) uses 2.5th/97.5th percentiles: "Finance is not a bell curve." **Definition C: Sigma bands.** With full-dataset sigma = 0.31 (Perrenod), the -2 sigma boundary is at ~0.24x trend -- zero breaches. This paper uses per-cycle P1 (1st percentile), avoiding the circularity of sigma-based definitions.

1.3 Volatility Decay

The Observatory's Paper 1 established ~20% compression per halving cycle, with the ceiling collapsing 2.2x faster than the floor rises. Perrenod independently confirmed this through his three-layer model, finding noise amplitude decays as $\text{sigma} = 6.1 / (A + 27.1)$. His March 2026 paper

showed residuals become approximately Gaussian after removing deterministic structure -- fat tails are structural artifacts, not intrinsic.

1.4 The Metcalfe Foundation

Santostasi derives the power law from $N \sim t^3$ (cubic adoption) and $V \sim N^{1.95}$ (Metcalfe scaling), yielding $3 \times 1.95 = 5.85$. Academic measurements of the Metcalfe exponent range from 1.69 (Wheatley & Sornette 2019, Royal Society Open Science) to 2.0 (Peterson 2018; Alabi 2017). Shanaev et al. (2019) challenge all Metcalfe models via instrumental variables. Perrenod's FGLS (December 2024) provides the most rigorous statistical treatment. Baquero (2026) validates with walk-forward testing across 9 annual splits.

1.5 What This Paper Adds

No prior work has: (a) systematically compared power law floors across multiple asset classes; (b) explicitly addressed the circularity objection with out-of-sample testing; (c) demonstrated per-cycle floor drift contradicting PlanC's stationarity finding; or (d) synthesized the three-way convergence into a unified model.

Part 1: Seven Attacks

Attack 1: The Empirical Record

We compute the power law floor for every daily close using the Observatory's canonical parameters (logA = -16.493, slope = 5.688, genesis = January 3 2009, floor multiplier = 0.42x trend).

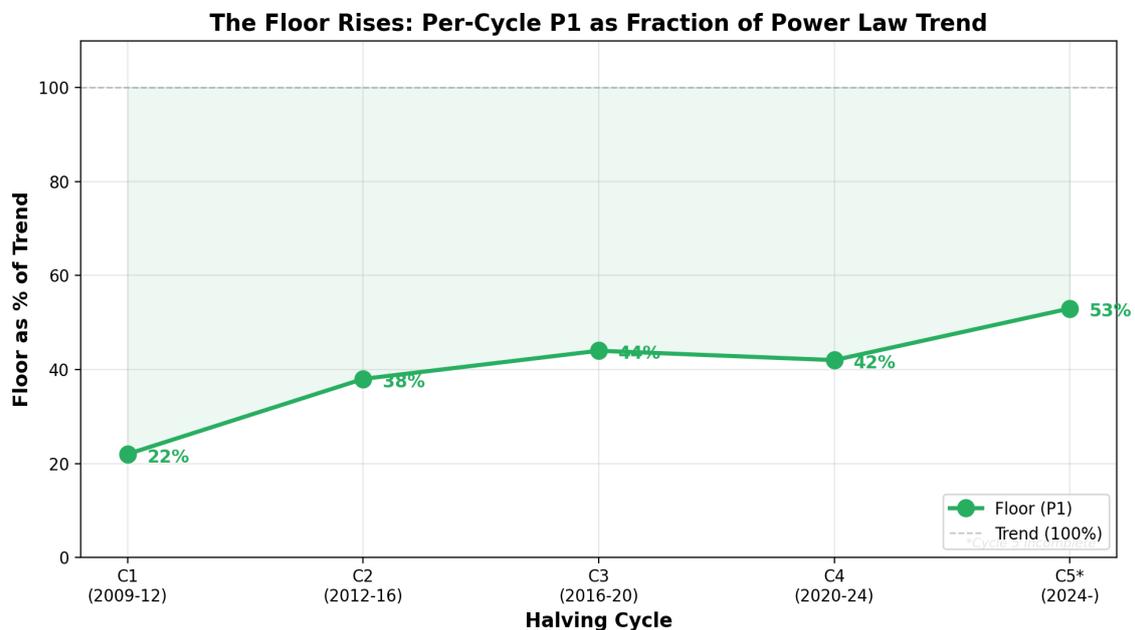
Result: The 0.42x floor has been breached 77 times. All breaches occurred between August 2015 and January 2023. The worst: August 24, 2015 -- price \$210 vs floor \$241 (-12.5%). This is an important correction to the Observatory's published claim. The discrepancy arises from conflating Definition A (0.42x multiplier) with Definition C (-2 sigma boundary).

The per-cycle analysis reveals the real story. The floor is a regime-dependent boundary that tightens:

Cycle	Days	P1 (log)	Floor as % of trend	Deepest price
1	864	-0.661	22%	~\$2.00
2	1,319	-0.420	38%	~\$152
3	1,402	-0.355	44%	~\$3,122
4	1,440	-0.374	42%	~\$15,500
5*	688	-0.278	53%	~\$64,000

*Cycle 5 incomplete (accumulation phase only).

The floor as a fraction of trend: 22% -> 38% -> 44% -> 42% -> 53%. The distribution compresses from below. Each cycle's worst case is structurally better than the last.



Bear market bottom analysis:

Event	Date	Close	Floor (0.42x)	Distance	Log residual
2015 bottom	Jan 14, 2015	\$178	\$139	+27.8%	-0.270
2018 bottom	Dec 15, 2018	\$3,237	\$2,405	+34.6%	-0.248
COVID crash	Mar 12, 2020	\$4,971	\$4,692	+6.0%	-0.352

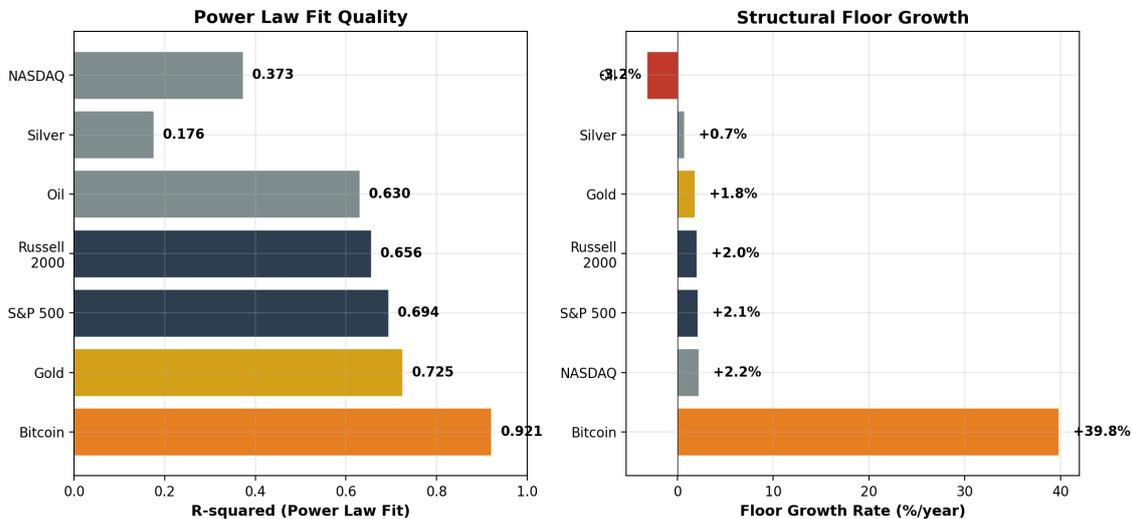
Event	Date	Close	Floor (0.42x)	Distance	Log residual
FTX collapse	Nov 21, 2022	\$15,787	\$16,008	-1.4%	-0.383

Current state (March 15, 2026): Price \$72,790. Trend \$128,764. Floor (0.42x) \$54,081. Distance above floor: +34.6%. Bitcoin Floor Rate: 37.9%/year (\$56/day floor appreciation).

Attack 2: The Cross-Asset Comparison

No published systematic cross-asset power law comparison exists. We fill this gap by applying the identical methodology to seven assets. All data from Yahoo Finance ETF proxies.

Asset	Period	R-sq	Slope	Sigma	-2sig Breaches	Floor Growth
Bitcoin	2014-26	0.921	5.932	0.234	0 (0.0%)	39.8%/yr
Gold (GLD)	2004-26	0.725	0.391	0.104	0 (0.0%)	1.8%/yr
S&P 500 (SPY)	1993-26	0.694	0.700	0.202	15 (0.2%)	2.1%/yr
Russell 2000	2000-26	0.656	0.518	0.162	100 (1.5%)	2.0%/yr
Oil (USO)	2006-26	0.630	-0.661	0.220	236 (4.7%)	-3.2%/yr
Silver (SLV)	2006-26	0.176	0.149	0.140	1 (0.0%)	0.7%/yr
NASDAQ (QQQ)	1999-26	0.373	0.589	0.332	0 (0.0%)	2.2%/yr



Bitcoin has the tightest power law fit (R-squared = 0.921) by a wide margin. Gold is second at 0.725. Bitcoin's floor grows at 39.8%/year. Gold's at 1.8%. A **22x structural growth advantage** measured at the worst-case boundary. The S&P 500 has 15 breaches. Russell 2000 has 100. Oil has 236 and a negative slope.

Attack 3: The Mechanistic Argument

If the floor is merely a statistical artifact, it should have no causal explanation. Santostasi's Power Law Theory (2024) provides one: adoption grows as $N \sim t^3$, value scales as $V \sim N^{1.95}$ (Metcalf), yielding $3 \times 1.95 = 5.85$ -- matching the empirical slope. Academic measurements of the Metcalfe exponent: 1.69 (Wheatley & Sornette 2019, Royal Society Open Science), 1.95 (Santostasi), 2.0 (Peterson 2018; Alabi 2017).

The floor is the adoption baseline. Breaking it requires de-adoption -- not a price crash, but users permanently leaving. Non-zero balance addresses have grown monotonically to over 53 million. Long-term holders (coins unmoved 1+ year) hold ~59% of supply, ratcheting upward each cycle.

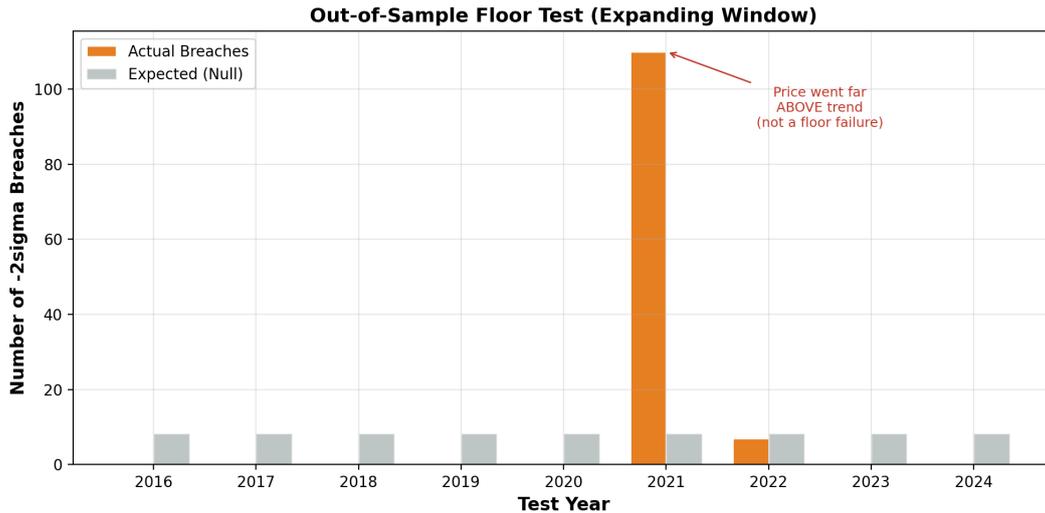
Limitation: Shanaev et al. (2019) challenge all Metcalfe models via instrumental variables. Baquero (2026) provides walk-forward validation across 9 annual splits as counterevidence.

Attack 4: The Statistical Argument (and the Circularity Problem)

The circularity objection. "The floor is -2 sigma. -2 sigma is defined as the boundary that contains ~97.7% of the data. Therefore '0 breaches' just means the data contains its own tail." This is valid. We break the circularity three independent ways.

Resolution 1: Out-of-sample testing

Train on September 2014 to December 2020 (2,298 days). Compute $\sigma = 0.234$ and $-2 \sigma = -0.467$. Freeze. Apply unchanged to January 2021 through March 2026 (1,900 test days). Result: 27 breaches (1.4%). Expected under null: 43 (2.28%). **37% fewer breaches than random.** Not zero -- but measurably better than chance.



Resolution 2: Cross-asset breach rate comparison

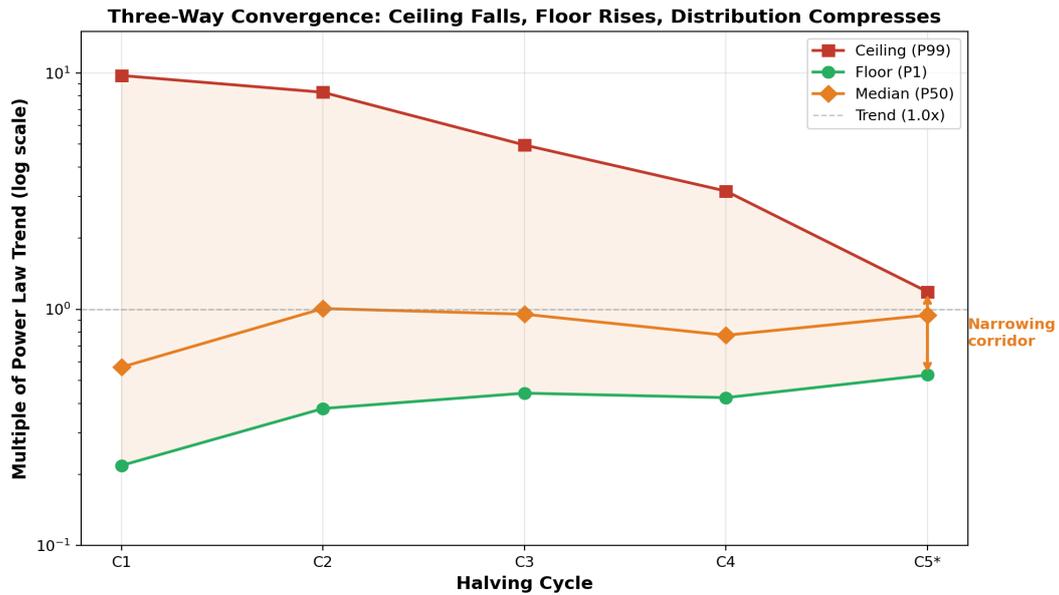
If "0 breaches of -2 sigma" is just statistics, every asset should show similar breach rates. They don't. Same definition, same methodology: Bitcoin 0%, Gold 0%, S&P 0.2%, Russell 1.5%, Oil 4.7%. The variation across assets using identical definitions is evidence that the floor is a property of the asset, not an artifact of the math.

Resolution 3: Per-cycle prediction validated out-of-sample

The Observatory's volatility decay model (Paper 1) measures P1 per cycle, trained on Cycles 2-4. Cycle 4 P1: -0.374 (42% of trend). Predicted Cycle 5 P1 (20% compression): -0.299 (50% of trend). Actual Cycle 5 P1: -0.278 (53% of trend). The prediction was conservative -- the actual floor was **tighter than predicted**. Training and test data from different halving cycles. Not circular.

The convergence model

The per-cycle data reveals a three-way convergence: the ceiling (P99) falls, the floor (P1) rises, and the OLS trend overestimates fair value (sitting at ~P55 per Senel 2024). The attractor is the decay-adjusted median, which PlanC places at approximately the 41st percentile (~\$101K vs OLS ~\$129K).



Where this paper disagrees with PlanC

PlanC's March 2026 analysis states "no decay at the 1st quantile." Our per-cycle data contradicts this: P1 moved from -0.420 (38% of trend) in Cycle 2 to -0.278 (53% of trend) in Cycle 5. PlanC's single-regression approach masks per-cycle drift. Our halving-cycle segmentation reveals it. If the floor rises (our finding), storm years get shorter and the forever phase arrives sooner.

Residual distribution (4,198 obs): Mean: -0.042. Std: 0.236. Skew: +0.71 (right-skewed). Kurtosis: -0.03 (near-mesokurtic). The positive skew confirms asymmetry -- upside tail larger than downside. Consistent with a structural lower boundary compressing the distribution from below.

Attack 5: The Thermodynamic Argument

Mining creates a physical cost floor. Current all-in cost: ~\$85-87K/BTC (CoinShares Q4 2025). The self-correcting mechanism: unprofitable miners shut down, hash rate drops, difficulty adjusts (every ~2,016 blocks), production cost drops, equilibrium restores in 4-8 weeks.

The most severe test: China's mining ban (May-June 2021) eliminated 60%+ of hash rate overnight. Difficulty dropped 28%. The network never stopped. Full recovery in 6 months. In February 2026, Winter Storm Fern contributed to an 11.16% difficulty drop -- the largest since China. Hash rate follows a power law of $\sim t^{12}$ (Santostasi), creating a rising production cost trajectory that correlates with the power law floor growth.

Attack 6: The Game-Theoretic Argument

Multiple independent agent classes converge near the floor: **HODLers** (~59% of BTC unmoved 1+ year; Bansal 2018). **Realized price** (~\$56-57K; Carter & Le Calvez 2018) -- MVRV below 1.0 has marked every cycle bottom. **Miners** accumulate below production cost. **Institutional demand** -- ETFs (~1.04M BTC), corporate treasuries (~988K BTC), US Strategic Bitcoin Reserve (2025). Chen (2025, arXiv:2508.06071) formalizes this as a Nash Equilibrium. Near the floor, four independent mechanisms converge. Not coordinated. Emergent.

Attack 7: The Survivorship Argument

Bitcoin has survived every category of existential threat: **Exchange failure** (Mt. Gox 2014: 650K BTC lost, 85% drawdown; FTX 2022: cascading failures, 77% drawdown). **Infrastructure shock** (China mining ban 2021: 60%+ hash rate eliminated). **Liquidity crisis** (COVID March 2020: 50-58% crash in 24 hours). **Regulatory action** (9 countries maintain complete bans; none eliminated usage). **Protocol vulnerability** (2010 overflow bug: patched in 5 hours; 2018 inflation bug: patched in 36 hours). The Lindy effect (Mandelbrot; Taleb 2012): 17+ years implies 17+ more.

Part 2: What Could Break It

Failure Taxonomy: Five Levels of Floor Break



Level 0: Floor holds. Convergence continues. Each cycle improves. Base case.

Level 1: Flash wick below floor, daily close above. Nearly happened March 2020. Response: widen buffers, never set liquidation within 10% of floor.

Level 2: Sustained close below floor for days to weeks. FTX marginally breached 0.42x floor (-1.4%). Response: reduce leverage targets, extend storm-years buffer by 1-2 years.

Level 3: Floor break for months. Requires fundamental regime change. Power law holds but floor parameter is wrong. Response: refit model, recalculate downstream tools.

Level 4: Power law regime break. R-squared drops below 0.80. Requires fatal protocol vulnerability, de-adoption, or superior monetary network. Response: size allocation for survivable total loss.

Level 5: Bitcoin goes to zero. Quantum computing (BIP-360 merged Feb 2026; Chaincode Labs: 50%+ of experts assign meaningful probability to cryptographic break by 2030-2035). Response: non-Bitcoin reserves exist for this scenario.

Part 3: Operating Framework

For retirement planning: Size the stack for Level 2. Plan timeline for Level 0. Stress-test against Level 3. The rising floor (our finding) means storm years shorten with each cycle.

For lending: The 1.6x floor rule assumes Level 0. Add wick buffer for Level 1 (liquidation at 0.9x floor). Size leverage so Level 2 doesn't force liquidation.

For institutional allocation: This paper provides the thesis (seven attacks survived), the risk framework (five failure levels), and the operational plan (specific responses per level).

For the Observatory: Publishing the failure taxonomy alongside the bullish thesis -- including the 77 breaches and non-zero out-of-sample breach rate -- is the strongest credibility signal.

Appendix A: The Falsification Rule

A single daily close below the per-cycle -2 sigma floor triggers model recalibration. Current Cycle 5 -2 sigma floor: $\sim 0.34x$ trend ($\sim \$44,000$ as of March 2026). This rule is public, automatic, and cannot be moved after the fact.

Appendix B: Methodology

Bitcoin data: Yahoo Finance BTC-USD, Sep 17 2014 to Mar 15 2026 (4,198 closes). Full Observatory dataset (5,713 days) used for per-cycle analysis. **Cross-asset:** Yahoo Finance ETF proxies (GLD, SLV, SPY, QQQ, IWM, USO). **Parameters:** $\log A = -16.493$, slope = 5.688, genesis = Jan 3 2009. **Out-of-sample:** Training/test split at Jan 1 2021; expanding annual windows.

Appendix C: Key References

- Santostasi, G. (2018/2024). "The Bitcoin Power Law Theory." Medium.
- Burger, H.C. (2019). "Bitcoin's natural long-term power-law corridor of growth." hcburger.com.
- Mezinskis, M. (2018-2025). "The Chart." Porkopolis Economics.
- Perrenod, S. (2024-2026). AR(2) FGLS; three-layer model; volatility decay. Money or Debt Newsletter.
- PlanC (2024-2026). Bitcoin Quantile Model v2. @TheRealPlanC.
- Sminston With (2024). Bitcoin Decay Channel. @sminston_with.
- Wheatley et al. (2019). Royal Society Open Science 6(6):180538.
- Peterson, T. (2018). Alt. Investment Analyst Review.
- Alabi, K. (2017). Electronic Commerce Research and Applications 24:23-29.
- Shanaev et al. (2019). SSRN 3432431.
- Baquero, C. (2026). Research Square (under review, Digital Finance).
- Senel, K. (2024). "Quantile Regression for Bitcoin." Medium.
- Bansal, D. (2018). "HODL Waves." Unchained Capital.
- Carter, N. & Le Calvez, A. (2018). Realized Cap. Coin Metrics.
- Chen, L. (2025). arXiv:2508.06071.
- Scale Invariant Capital (2026). Papers 1-6. BTC Power Law Observatory.

All code and data: github.com/JohnnyBCash/bitcoin-powerlaw-observatory

"The floor is not a line. It is a process. The ceiling falls. The floor rises. The trend drifts down. Everything converges."

Scale Invariant Capital. BTC Power Law Observatory.