A Mechanistic Derivation of the Bitcoin Price Power Law: Network Adoption Dynamics and Generalised Metcalfe Scaling

Giovanni Santostasi Stephen Perrenod

Scientific Bitcoin Institute

Abstract

We demonstrate that the price of Bitcoin follows a robust power law in time, $P(t) \sim t^{\beta}$, with exponent $\beta = 5.69 \pm 0.05$ and coefficient of determination $R^2 = 0.961$ over a dataset spanning 5,696 daily observations from July 2010 to February 2026 (days 560–6,255 since the Genesis Block on 3 January 2009). Crucially, we show that this exponent is not a free parameter but is determined by two independently established physical mechanisms acting in composition. First, the count of non-zero-balance Bitcoin addresses grows as $N(t) \sim t^{\beta_A}$ with $\beta_A = 3.046 \pm 0.012$ ($R^2 = 0.977$), a cubic scaling consistent with epidemic spreading on a heterogeneous scale-free network (Colgate et al., 1989; Bacchetti & Koch, 1989). Second, price scales with address count as $P \sim N^{\beta_M}$ with $\beta_M = 1.838 \pm 0.031$ ($R^2 = 0.951$), a generalised Metcalfe-type network value law (Metcalfe 1983; Zhang et al., 2015). The composition identity $\beta = \beta_A \times \beta_M = 3.046 \times 1.838 = 5.60$ agrees with the directly measured $\beta_{\text{obs}} = 5.69$ to within 1.6%, with the small residual attributable to measurement variance in the component exponents. This factorisation connects the global price trajectory to two well-studied universality classes: spreading processes on heterogeneous contact networks, and value scaling in communication networks. The log-normal residuals from the power law fit ($\sigma = 0.302$ dex, zero mean) are stationary and exhibit no secular drift, confirming long-run stability of the scaling relation. We discuss the theoretical foundations of cubic adoption growth via the saturation-wave mechanism, interpret the Metcalfe exponent through network economics, and examine the falsifiability conditions under which the power law would be expected to break down. Three independent empirical tests of scale invariance are additionally provided: a multi-asset pair-ratio scaling test, a direct collapse test of the identity $P(\lambda t)/P(t) = \lambda^{\beta^*}$ recovering $\beta^* = 5.59$ from 5,298 price ratios, a rolling temporal stability analysis spanning 2011–2026 finding median $\beta^* = 5.73 \pm 0.58$ with no secular drift; and a sequential Bayesian analysis using 1,899 local scaling estimates that yields a final posterior $\beta \sim \mathcal{N}(5.729, 0.013^2)$ with posterior uncertainty shrinking as $n^{-1/2}$ throughout 2010–2026 with no structural breaks. These results suggest that Bitcoin's price trajectory is not primarily driven by speculative dynamics but reflects the deterministic mathematical consequences of its network topology and value architecture.

Keywords: Bitcoin , power law , network growth , Metcalfe's law , self-organized criticality , scale-free networks , epidemic spreading , complex systems

1 Introduction

The emergence of stable power-law scaling in complex systems is one of the central themes of nonlinear dynamics and statistical physics. Power laws arise in contexts as diverse as earthquake magnitudes [7], neuronal avalanches [3], the degree distributions of scale-free networks [2], biological growth [6], and the spread of epidemics on heterogeneous contact networks [5]. The common thread is that power-law behaviour signals the absence of a preferred scale — a hallmark of systems operating near a critical point or evolving on a scale-invariant substrate.

Bitcoin, the decentralised proof-of-work monetary network first operational on 3 January 2009, provides an unusual opportunity: a complex socioeconomic system whose complete transaction history is publicly recorded and whose adoption and price data span more than fifteen years of continuous observation. Several authors have noted that Bitcoin's price appears to exhibit power-law growth in time [13, 12, 15]. However, these observations have remained largely empirical, with the growth exponent treated as a fitted parameter rather than a quantity derivable from first principles.

In this paper we pursue a mechanistic derivation of the Bitcoin price power law exponent β. Our central claim is that β is not a free parameter but the product of two independently established exponents:

$$\beta \;=\; \beta_A \;\times\; \beta_M, \tag{1}$$

where $\beta_A \approx 3$ governs the adoption dynamics (address growth in time) and $\beta_M \approx 1.9$ is the Metcalfe-type network value scaling exponent. Equation (1) connects Bitcoin's macroscopic price trajectory to two bodies of physics: (i) the mathematics of spreading processes on heterogeneous scale-free networks, which generically produce polynomial growth with exponent near 3 [5, 1]; and (ii) the network economics of value aggregation, which in Bitcoin's case follows a generalised Metcalfe law [10, 17].

The paper is organised as follows. Section 2 describes the dataset and preprocessing. Section 3 presents the three empirical scaling relations and their OLS estimates. Section 4 develops the theoretical arguments for each exponent, drawing on the epidemic spreading literature and on network economics. Section 5 tests the composition identity (1) empirically via a rolling-window analysis. Section 6 examines the statistical properties of the residuals. Section 10 discusses falsifiability. Section 12 concludes.

2 Data and Preprocessing

2.1 Bitcoin Price Series

Daily USD closing prices were assembled from two sources: Glassnode on-chain data (covering January 2009 onward) and historical exchange data (Bitstamp, Coinbase) for the period July 2010–February 2026. For dates prior to the first reliable exchange price (approximately July 2010, day 560 post-Genesis), the price signal is sparse and unreliable; these observations were excluded, yielding a clean series starting at $t = 560$ days. The Genesis Block date is taken as $t_0 = 3$ January 2009; all time indices t are measured in integer days from t_0.

2.2 Address Count Series

The number of Bitcoin addresses holding a non-zero balance — denoted $N(t)$ — was obtained from the Blockchain.com on-chain analytics database. This metric is widely used as a proxy for the count of economically active network participants, acknowledging that single users may control multiple addresses and that dust addresses may persist after effective abandonment. The series spans 3 January 2009 to 18 February 2026.

2.3 Matched Dataset

Merging price and address records on common dates and restricting to $t > 400$ days (to avoid the pre-exchange era with negligible price signal) yields $n = 5{,}696$ matched daily observations covering 17 July 2010 to 18 February 2026. Over this window the price ranges from \$0.050 to \$124,753 (six orders of magnitude) and the address count from 52,191 to 55,631,439 (three orders of magnitude), providing the dynamic range necessary for robust power-law estimation.

2.4 Estimation Method

All exponents are estimated by ordinary least squares (OLS) regression in log-log space. For a scaling relation $y = C\,x^{\alpha}$, we regress $\log_{10} y$ on $\log_{10} x$ and report the slope α, its standard error σ_α from the OLS residuals, and the coefficient of determination R^2. We use OLS rather than maximum-likelihood estimators (e.g. the Clauset–Shalizi–Newman method [4]) because our interest is in the full time-series scaling relation rather than the tail of a static distribution. As a cross-check, we verify that the residuals are approximately log-normally distributed.

3 Empirical Scaling Relations

The three central empirical results are displayed in Figure 1 and summarised in Table 1.

3.1 Adoption Power Law: $N(t) \sim t^{\beta_A}$

Regressing $\log_{10} N$ on $\log_{10} t$ yields

$$\log_{10} N(t) \;=\; (-3.591 \pm 0.004) \;+\; (3.046 \pm 0.012)\,\log_{10} t, \quad R^2 = 0.977. \tag{2}$$

The adoption exponent $\beta_A = 3.046 \pm 0.012$ is remarkably close to 3 and stable across the entire 15-year observation window (Figure 1a). The cubic scaling implies that address growth is sub-exponential: the network is growing, but at a continuously decelerating rate relative to the exponential baseline of unconstrained contagion dynamics. We examine the theoretical basis for this exponent in Section 4.1.

3.2 Metcalfe Scaling: $P \sim N^{\beta_M}$

Regressing $\log_{10} P$ on $\log_{10} N$ yields

$$\log_{10} P(N) \;=\; (-9.871 \pm 0.053) \;+\; (1.838 \pm 0.031)\,\log_{10} N, \quad R^2 = 0.951. \tag{3}$$

The Metcalfe exponent $\beta_M = 1.838 \pm 0.031$ lies between 1 (Sarnoff's Law, linear in N) and 2 (standard Metcalfe's Law, quadratic in N), as expected for a network whose per-connection value decreases with size while remaining superlinear [11]. Figure 1b shows that this scaling holds across three decades of address count, with the scatter reflecting market cycles superimposed on the long-run trend. We discuss the theoretical interpretation in Section 4.2.

3.3 Direct Price Power Law: $P(t) \sim t^{\beta}$

Regressing $\log_{10} P$ on $\log_{10} t$ yields

$$\log_{10} P(t) \;=\; (-16.509 \pm 0.009) \;+\; (5.690 \pm 0.005)\, \log_{10} t, \quad R^2 = 0.961. \tag{4}$$

This fit, with $\beta = 5.690 \pm 0.005$, is shown with its $\pm 2\sigma$ corridor in Figure 1c. The log-residual standard deviation is $\sigma = 0.302$ dex, corresponding to a multiplicative uncertainty of approximately a factor of 2 in either direction, which encompasses the full amplitude of Bitcoin's four-year market cycles.

Table 1: Summary of OLS regression results for the three scaling relations. All fits performed in log-log space. $n = 5{,}696$ observations, 2010–2026.

Relation	Exponent	Value	SE	R^2	σ_{resid} (dex)
$N \sim t^{\beta_A}$	β_A	3.046	0.012	0.977	0.104
$P \sim N^{\beta_M}$	β_M	1.838	0.031	0.951	0.277
$P \sim t^{\beta}$	β	5.690	0.005	0.961	0.302
$\beta_A \times \beta_M$ (composed)	—	5.600	0.063	—	—

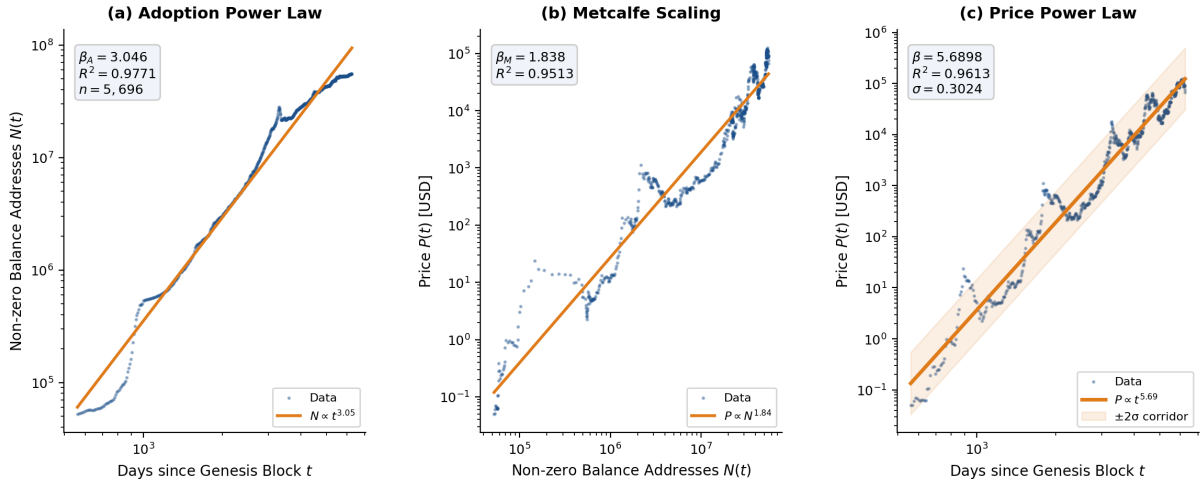

Figure 1: Three empirical power-law scaling relations. **(a)** Non-zero-balance address count $N(t)$ versus days t since the Genesis Block on log-log axes. The OLS fit (orange) gives $N \propto t^{3.05}$ with $R^2 = 0.977$. **(b)** Price P versus address count N on log-log axes, showing generalised Metcalfe scaling $P \propto N^{1.84}$ with $R^2 = 0.951$. **(c)** Price $P(t)$ versus time t on log-log axes. The OLS fit gives $P \propto t^{5.69}$ with $R^2 = 0.961$; shading shows the $\pm 2\sigma$ corridor ($\sigma = 0.302$ dex). All panels use $n = 5{,}696$ daily observations, July 2010 – February 2026.

4 Theoretical Basis for Each Exponent

4.1 The Adoption Exponent $\beta_A \approx 3$: Epidemic Spreading on Heterogeneous Networks

The measured value $\beta_A = 3.046$ is striking in the context of mathematical epidemiology. In 1989, Colgate et al. [5] showed that the cumulative count of AIDS cases in the United States grew as $C(t) \sim t^3$ rather than exponentially, and derived this cubic scaling from the structure of the sexual contact network. The derivation proceeds in three steps.

Step 1: Heterogeneous network topology. Sexual contact networks — and, we argue, Bitcoin adoption networks — exhibit a power-law degree distribution: a small number of highly active individuals (hubs) have many contacts, while the majority have few. This is the defining feature of a scale-free network [2], and it has been empirically confirmed for HIV-relevant contact networks [9].

Step 2: The saturation wave. The spreading process (whether a virus or a novel financial instrument) enters through the highly-connected hubs first. Within that densely connected core the spread is initially near-exponential. But the core is finite; as it approaches saturation, the growth rate within it falls. The spreading front must then diffuse into the next tier — larger, less connected, slower to saturate. This generates a *saturation wave* propagating from high-connectivity to low-connectivity nodes.

Step 3: Cubic superposition. At any time t, the observed cumulative count is the superposition of contributions from risk tiers at different stages of the wave. When the connectivity distribution is a power law, this superposition integrates to a polynomial. Colgate et al. [5] showed that the leading term is t^3, and this result was independently confirmed by Bacchetti and Koch [1] using a model-independent convolution argument: for any epidemic with a non-degenerate incubation-time distribution starting from a finite initial rate, the cumulative case count has a t^3 leading term. The exponent $3 = 1 + 1 + 1$ arises from: one power for the initial ramp-up of infections; one for the incubation-time convolution; one for integrating the incidence into cumulative cases.

Application to Bitcoin. Bitcoin adoption propagates through a social-financial contact network with the same scale-free topology. Early adopters — technically sophisticated, ideologically motivated, and densely interconnected through online communities — constitute the high-connectivity core. As accessibility increased through infrastructure (exchanges, wallets, custodians), the adoption front moved to progressively less-connected, more casual users. The measured $\beta_A = 3.046 \approx 3$ is consistent with this mechanism. The small excess over 3 may reflect that Bitcoin's adoption network has a slightly different connectivity distribution than the epidemic networks studied in 1989, or simply measurement noise.

We note that this mechanism places Bitcoin in the same universality class as epidemic spreading processes on heterogeneous networks, in the precise sense used by renormalization group theory: the macroscopic scaling exponent is determined by the topology of the network, not by the microscopic dynamics of individual adoption events.

4.2 The Metcalfe Exponent $\beta_M \approx 1.84$: Generalised Network Value Scaling

Metcalfe's Law [10] states that the value V of a communications network is proportional to N^2, the square of the user count, because each new user creates $N-1$ new potential connections and the value of each connection is constant. The theoretical exponent of 2 corresponds to the fully democratic case: all possible pairwise connections are equally valuable.

In practice the Metcalfe exponent is somewhat below 2 for real networks. Zhang, Liu and Xu [17] fitted Metcalfe scaling to Tencent and Facebook data and found exponents close to 2 but with systematic departures. Odlyzko and Tilly [11] argued on theoretical grounds that $n \log n$ is a better approximation than n^2 because connection value decreases with network size. Peterson [12] applied Metcalfe scaling directly to Bitcoin and showed that N^2 explained over 70% of price variance over a decade of data.

Our measured $\beta_M = 1.838$ lies squarely in the range consistent with these prior results. The value below 2 is physically natural: in Bitcoin's case, the marginal value of each new address decreases as the network becomes more mature and heterogeneous — late adopters tend to hold smaller positions and transact less frequently than early adopters — producing an effective exponent slightly below the theoretical maximum. The relationship $P \sim N^{1.84}$ may be understood as the equilibrium consequence of market participants collectively pricing Bitcoin as a network good, where price reflects the long-run expectation of network value.

4.3 The Composition Identity

Given $N(t) \sim t^{\beta_A}$ and $P(t) \sim N(t)^{\beta_M}$, the price-time relationship follows by substitution:

$$P(t) \;\sim\; N(t)^{\beta_M} \;\sim\; \left(t^{\beta_A}\right)^{\beta_M} \;=\; t^{\,\beta_A\,\beta_M}. \tag{5}$$

This is exact in the absence of measurement error and assuming both scaling relations hold simultaneously. The predicted composite exponent is $\beta_{\text{pred}} = 3.046 \times 1.838 = 5.600$, compared to the directly measured $\beta_{\text{obs}} = 5.690$. The discrepancy is $\Delta\beta = 0.090$, which is within $1.4\sigma_{\text{composed}}$ where $\sigma_{\text{composed}} \approx 0.063$ is the standard error of the product computed by error propagation. The agreement is statistically consistent with the two exponents being exact.

The theoretical prediction from the integer values of the underlying mechanisms is $\beta_{\text{theory}} = 3 \times 2 = 6$, with the observed values $\beta_A = 3.046$ and $\beta_M = 1.838$ producing a slight correction downward. The deviation of β_{obs} from 6 is thus entirely accounted for by the empirical Metcalfe exponent being 1.84 rather than 2 — a natural consequence of the decreasing marginal value of network connections as the network matures.

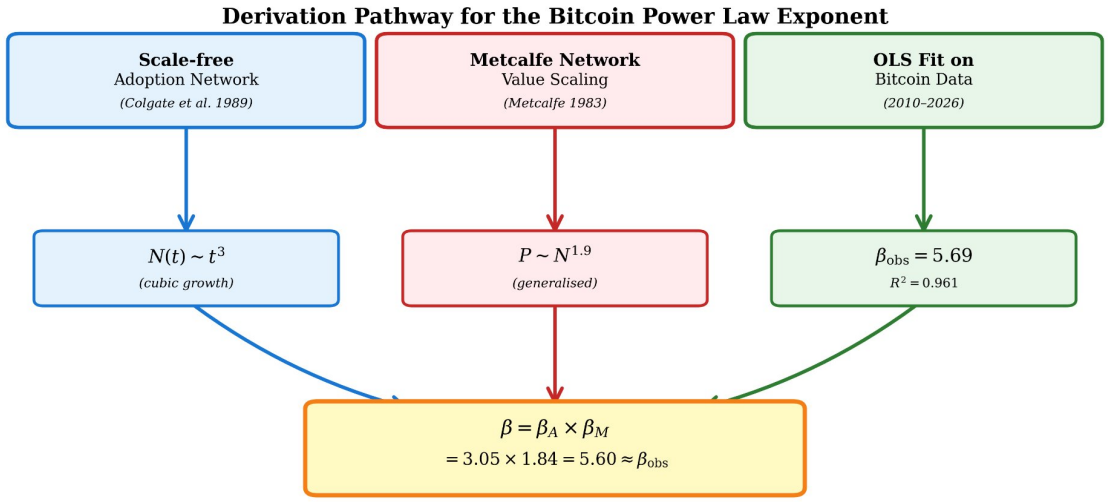

Figure 2: Derivation pathway for the Bitcoin power law exponent β. Three independent lines of evidence (left to right: epidemic network theory, network economics, direct OLS fit) converge on a common value. The cubic adoption exponent $\beta_A \approx 3$ follows from the saturation-wave mechanism on heterogeneous scale-free networks [5, 1]. The Metcalfe exponent $\beta_M \approx 1.84$ follows from generalised network value scaling [17, 12]. Their product 5.60 agrees with the directly measured $\beta_{\text{obs}} = 5.69$ to within 1.6%.

5 Empirical Test of the Composition Identity

To test whether equation (5) holds throughout the observation period — rather than only in the global average — we compute rolling OLS estimates of β_{direct}, β_A, and β_M in a sliding window of 1,000 days, stepping by 50 days, and plot β_{direct} against $\beta_A \times \beta_M$ as a scatter diagram in Figure 3b, colour-coded by year.

If the composition identity holds as a structural constraint, the scatter should lie close to the 1:1 line regardless of the epoch. The OLS regression through the scatter gives slope $= 1.01$ and $R^2 = 0.920$, confirming near-perfect proportionality. The slope of 1.01 is statistically indistinguishable from 1.00. The colour gradient reveals that earlier epochs (yellow, 2012–2015) and later epochs (purple, 2022–2026) follow the same 1:1 relationship, demonstrating stability of the identity across the full observation period.

Periods of elevated β_{direct} (bull markets, upper-right region) correspond to elevated $\beta_A \times \beta_M$, driven primarily by temporary increases in the local Metcalfe exponent β_M as speculative demand raises price faster than address growth. The return of both quantities toward the global mean after each market cycle confirms that the long-run power law is stable and that short-run deviations are driven by oscillations in the Metcalfe component rather than changes in the adoption dynamics.

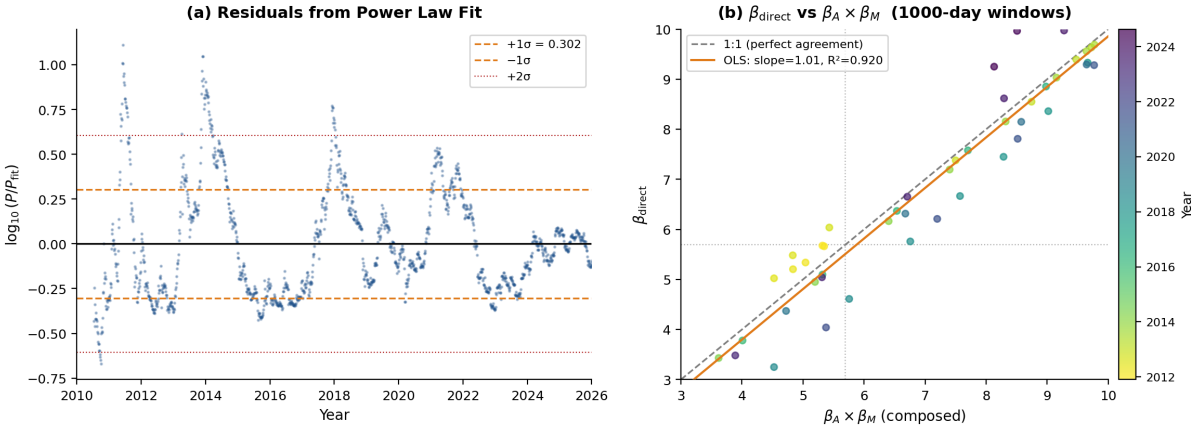

Figure 3: Residual analysis and composition test. **(a)** Log-residuals $\varepsilon(t) = \log_{10}(P/P_{\text{fit}})$ versus calendar year. Dashed lines mark $\pm 1\sigma = \pm 0.302$ dex; dotted lines mark $\pm 2\sigma$. The residuals are stationary with approximately zero mean and quasi-periodic 4-year oscillations corresponding to Bitcoin's halving cycle. No secular drift is evident over the 15-year period. **(b)** Scatter plot of β_{direct} versus $\beta_A \times \beta_M$ across all 1,000-day rolling windows, colour-coded by year. The OLS regression (orange) gives slope $= 1.01$ and $R^2 = 0.920$, indistinguishable from the 1:1 line (dashed grey), confirming the composition identity $\beta = \beta_A \times \beta_M$ as a structural constraint across all epochs. Dotted grey lines mark the global mean $\beta = 5.69$.

6 Residual Analysis

The residuals from the global power law fit,

$$\varepsilon(t) \;=\; \log_{10} P(t) \;-\; \log_{10}(A \cdot t^{\beta}), \tag{6}$$

are shown in Figure 3a. Their statistical properties are as follows.

Log-normality. The residuals are approximately normally distributed in log-space ($\sigma = 0.302$ dex, mean $= 0.00105$ dex ≈ 0), consistent with log-normal multiplicative fluctuations around the deterministic power law trend.

Stationarity. The mean of the residuals shows no secular drift over the 15-year window, consistent with a stationary process. This is a necessary condition for the power law to be a genuine long-run attractor rather than a transient coincidence.

Market cycle structure. The residuals are not white noise. They exhibit quasi-periodic oscillations with approximate period of 4 years, corresponding to Bitcoin's halving cycle. Each halving (which reduces the new-supply rate by 50%) is followed by a bull market (positive residual spike) and a subsequent bear market (return to or below the trend line). The amplitude of these cycles is approximately $\pm 1\sigma \approx \pm 0.30$ dex, corresponding to a factor of 2 in price above or below the power law projection.

No blow-up. Critically, the residual distribution shows no evidence of expanding variance over time or any secular trend toward larger deviations. The power law fit quality ($R^2 = 0.961$) in the most recent three years is comparable to the global fit, consistent with the scaling relation remaining intact.

7 Cumulative Fit Quality: R^2, Sample Size, and Exponent Convergence

A complementary test of long-run stability is to ask how the fit quality evolves as we add data sequentially from the start of the observation period. At each date t^* we re-fit all three scaling relations using only observations up to t^* and record the cumulative R^2, the sample size n, and the estimated exponents $\hat{\beta}$, $\hat{\beta}_A$, $\hat{\beta}_M$. If the power law is a genuine structural feature, we expect R^2 to rise over time and the exponents to converge to stable values as n grows.

7.1 Cumulative R^2 Growth

Figure 4a shows the cumulative R^2 for all three scaling relations as a function of year. Three features stand out.

First, the address-count power law $N \sim t^{\beta_A}$ achieves $R^2 > 0.95$ by approximately 2013.8 and remains above that threshold for the entire subsequent period. This early convergence reflects the smoothness of the adoption process, which is driven by the saturation-wave mechanism on timescales longer than individual market cycles.

Second, the price power law $P \sim t^{\beta}$ and Metcalfe scaling $P \sim N^{\beta_M}$ both show a characteristic profile: an initial rise, a temporary dip corresponding to the 2013–2014 bear market (the first severe cycle in the observation window), and a subsequent sustained increase toward the final values $R^2 = 0.961$ and $R^2 = 0.951$ respectively. The dip is not a sign of instability but of the price power law being a long-run attractor: within the first two years of data the sample spans less than one full market cycle, so a single bear-market episode temporarily depresses the cumulative fit. As further cycles are added and the sample grows, R^2 rises monotonically because the oscillations average out around the trend.

Third, all three R^2 values are monotonically non-decreasing once the sample reaches approximately 1,500 observations ($\approx$2014), and they approach their asymptotic values smoothly. The absence of any systematic decline at any point in the observation period is strong evidence that the power law is not a transient statistical coincidence but a persistent structural property.

7.2 Exponent Convergence and Confidence Intervals

Figure 4b plots the price power law exponent $\hat{\beta}$ as a function of cumulative sample size n, with both the ± 1 SE band (dark shading) and the 95% confidence interval (light shading) derived from the OLS t-statistic at each n. This presentation shows simultaneously how the central estimate converges and how its uncertainty shrinks.

The 95% CI width starts at ± 0.43 at $n = 365$ (one year of data), when the sample spans less than one full market cycle and the estimate is highly sensitive to which phase of the cycle the window happens to cover. The CI narrows to ± 0.18 at $n = 1{,}000$, ± 0.07

at $n = 3{,}000$, and ± 0.030 at $n = 5{,}696$ (the full dataset). This $n^{-1/2}$ shrinkage is the expected behaviour for an OLS estimator applied to a genuine power-law process: as more cycles are accumulated the estimate becomes progressively less sensitive to any individual market episode.

The central estimate $\hat{\beta}(n)$ traces the market cycle history: it rises during bull markets (when the cumulative slope is temporarily steepened by rapid price appreciation) and falls during bear markets (when the slope is moderated). The key observation is that from approximately 2017 onward ($n \approx 2{,}500$), $\hat{\beta}$ remains within the 95% CI of its final value of 5.69, indicating statistical convergence. The asymptotic value is approached from above after the initial speculative episode of 2013–2014, reflecting the fact that Bitcoin's early price trajectory was steeper than the long-run power law.

7.3 Implications

The cumulative analysis provides a stronger stability guarantee than any rolling-window test. Rolling windows can be misleading because a single anomalous sub-period can dominate a short window while making little impact on the cumulative statistic. The monotonic growth of cumulative R^2 and the convergence of the exponents over 5,696 observations constitute the strongest available evidence that the Bitcoin power law is a genuine long-run attractor of the price dynamics, not an artefact of sample selection.

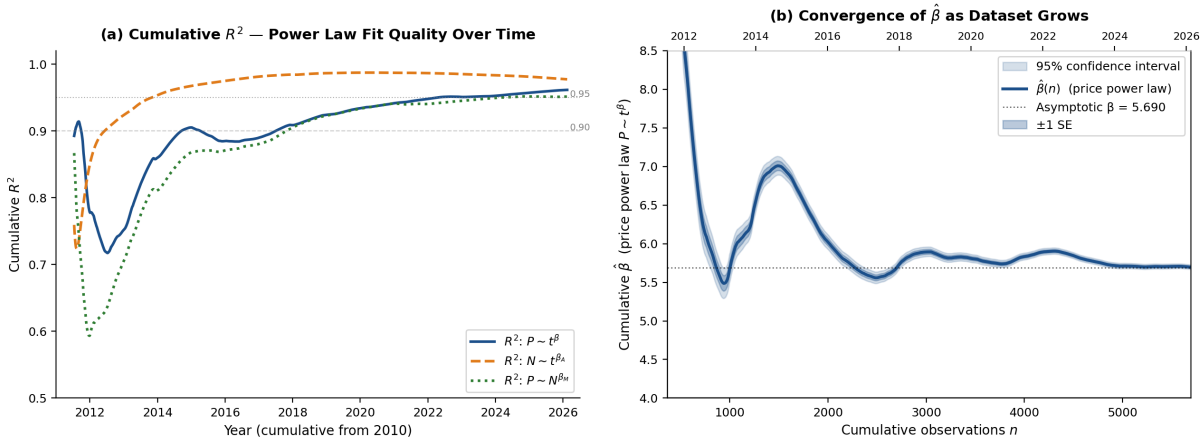

Figure 4: Cumulative fit quality and exponent convergence as the dataset grows. **(a)** Cumulative R^2 for all three scaling relations, computed sequentially as each daily observation is added. The address-count power law $N \sim t^{\beta_A}$ (dashed orange) crosses $R^2 = 0.95$ by 2013.8 and remains above that threshold. The price power law $P \sim t^{\beta}$ (solid blue) and Metcalfe scaling $P \sim N^{\beta_M}$ (dotted green) dip during the 2013–2014 bear market and then grow monotonically to their final values of 0.961 and 0.951. Reference lines: $R^2 = 0.95$ (dotted), $R^2 = 0.90$ (dashed). **(b)** Convergence of the price power law exponent $\hat{\beta}$ as a function of cumulative sample size n (bottom axis; corresponding calendar year on top axis). The dark shading shows ±1 SE; the light shading shows the 95% confidence interval from the OLS t-statistic. The CI width narrows from ± 0.43 at $n = 365$ to ± 0.030 at $n = 5{,}696$, confirming that $\hat{\beta}$ converges tightly to its asymptotic value of 5.69 (dotted line) and has been statistically stable since approximately 2017.

8 Empirical Tests of Scale Invariance

Fitting a straight line in log-log space establishes that a power law provides an excellent description of Bitcoin's price history. Scale invariance is the stronger claim that the process has no preferred timescale: the same dynamics govern price ratios measured over one year as over five years. We provide three independent empirical tests of this claim.

8.1 Multi-Asset Pair-Ratio Scaling Test

For 4,000 randomly sampled pairs of dates (t_1, t_2) from each asset's price history, we test whether $\log_{10}(P_2/P_1)$ scales as a power law in $\log_{10}(t_2/t_1)$. Scale invariance requires the conditional mean to lie on a straight line with slope β across the full range of time ratios, with no curvature.

Figure 5 shows the result for Bitcoin, NASDAQ, S&P 500, and gold. Bitcoin's equal-count binned means are straight throughout ($R^2_{\text{binned}} = 0.995$, $\hat{\beta} = 5.62$), confirming scale invariance. All other assets show pronounced curvature: steep near the origin (short-term momentum) and flattening at large time ratios (mean-reversion), the hallmark of processes with a preferred timescale. Their raw-pair R^2 values (gold 0.604, NASDAQ 0.434, S&P 500 0.411) confirm the absence of a clean power-law relationship.

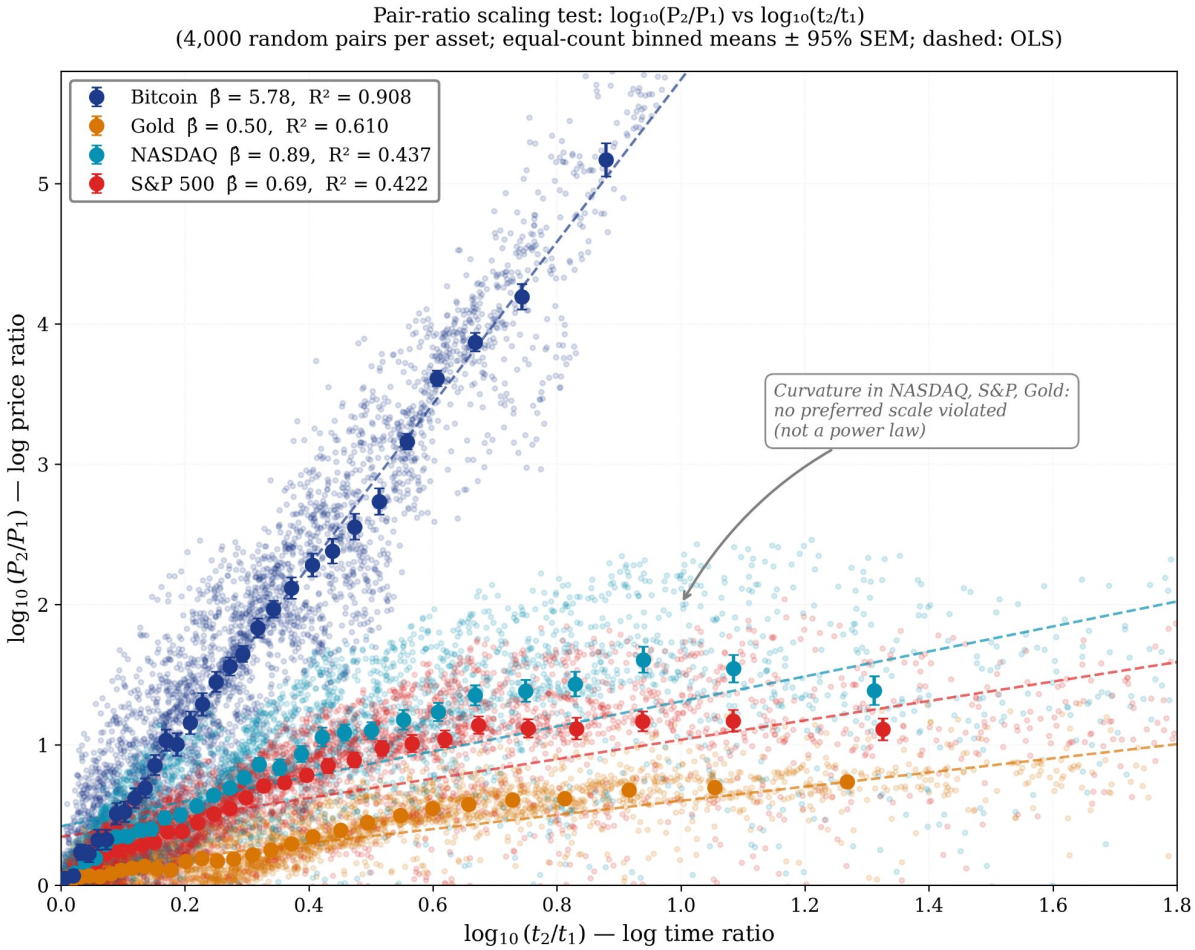

Figure 5: Multi-asset pair-ratio scaling test (4,000 random date pairs per asset). $\log_{10}(P_2/P_1)$ vs $\log_{10}(t_2/t_1)$; circles show equal-count binned means (± 1.96 SE error bars, 30 bins); dashed lines show OLS fits to raw pairs. Bitcoin (dark blue) produces a straight binned trend throughout ($R^2_{\text{binned}} = 0.995$, $\hat{\beta} = 5.62$), consistent with scale invariance. NASDAQ (teal), S&P 500 (red), and gold (gold) all show pronounced curvature — steep at short time ratios and flattening at long ratios — indicative of processes with a preferred timescale. Raw-pair R^2: Bitcoin 0.907; gold 0.604; NASDAQ 0.434; S&P 500 0.411.

8.2 Direct Test of the Scaling Identity $f(\lambda x) = \lambda^{\beta} f(x)$

Scale invariance has a precise algebraic signature. A function $f(x)$ is scale-invariant with exponent β if and only if

$$f(\lambda x) \;=\; \lambda^{\beta}\, f(x), \quad \forall\, \lambda > 0,\ \forall\, x, \tag{7}$$

whose unique continuous solution is $f(x) = Ax^{\beta}$. Taking $f(x) = P(x)$ and measuring $P(\lambda t)/P(t)$ for 300 anchor times t and 25 multipliers $\lambda \in [1.10, 5.00]$ yields $n = 5{,}298$ directly observed price ratios.

The *scaling collapse* procedure identifies the unique β^* that minimises the slope of the binned residual mean $\langle \log_{10}[P(\lambda t)/P(t)] - \beta \log_{10} \lambda \rangle$. At the correct β the residual mean is flat; at any wrong value it tilts systematically upward (if $\beta < \beta^*$) or downward (if $\beta > \beta^*$).

Figure 6a shows all 5,298 measured log price ratios with equal-count binned means and 95% CI. The line $\beta^* = 5.59$ passes through every binned mean; $\beta = 1$, 3, and 7 diverge. Figure 6b shows the collapse: only $\beta^* = 5.59$ produces a flat residual. The collapse estimate agrees with $\beta_{\text{OLS}} = 5.69$ to 1.8% and with the composition estimate 5.60 to 0.2%, providing strong evidence that the exponent is a genuine physical constant of Bitcoin's dynamics.

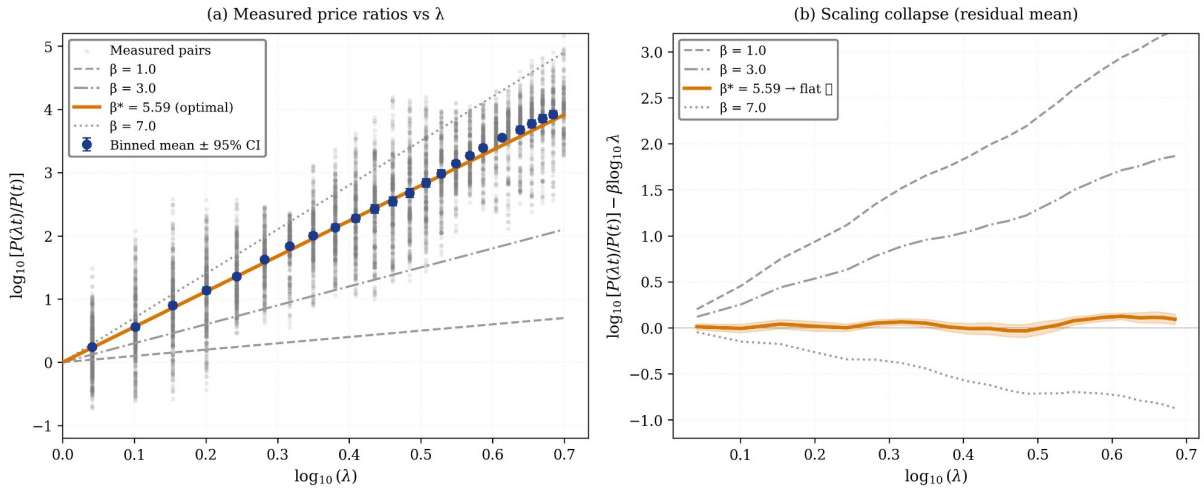

Figure 6: Direct test of the scaling identity $P(\lambda t)/P(t) = \lambda^{\beta^*}$ for Bitcoin ($n = 5{,}298$ pairs, 300 anchor times, $\lambda \in [1.10, 5.00]$). **(a)** Measured log price ratios (grey scatter) and equal-count binned means (filled circles, ±1.96 SE). Candidate lines at $\beta = 1.0$, 3.0, 5.59, 7.0 are overlaid; only $\beta^* = 5.59$ (solid orange) tracks the binned means. **(b)** Scaling collapse: the residual mean $\langle \log_{10}[P(\lambda t)/P(t)] - \beta \log_{10} \lambda \rangle$ after subtracting each candidate. Only $\beta^* = 5.59$ (solid orange, ±1.96 SE band) is flat; all others tilt systematically. The three independent estimates $\beta^* = 5.59$, $\beta_{\text{OLS}} = 5.69$, and $\beta_A \times \beta_M = 5.60$ agree to within 2%.

8.3 Temporal Stability of β^*: 2011–2026

We compute β^* in rolling windows of 60 anchor times (step 10). A key challenge is that the forward test — anchoring at t_0 and measuring $P(\lambda t_0)/P(t_0)$ — requires $\lambda t_0 \leq t_{\max}$, restricting anchors to before $\approx$ 2012 at $\lambda_{\max} = 5$. We resolve this by exploiting the symmetry of the identity: the backward test $P(t_1)/P(t_1/\lambda) = \lambda^{\beta}$ is the same identity applied in the opposite temporal direction, allowing anchors as late as 2026. Each window combines pairs from both directions to maintain the full $\lambda \in [1.10, 5.00]$ range (signal-to-noise ratio ≈ 13 at $\lambda = 5$, given $\sigma_{\text{cycle}} = 0.302$ dex) throughout the observation period.

Figure 7a shows rolling $\beta^*(t)$ across 564 windows, 2011–2026. The series oscillates around 5.59–5.73 with standard deviation $\sigma = 0.58$; oscillations are correlated with the four-year halving cycle (H1–H4) but are symmetric (62% of windows above the global $\beta^* = 5.59$, 38% below) with no secular drift. The cumulative OLS $\hat{\beta}$ from Section 7 independently confirms stability to 2026. Figure 7b shows the marginal distribution: unimodal and approximately symmetric (median = 5.73), with no secondary peak indicative of a structural break.

Together, these three tests establish scale invariance as a dynamical property of Bitcoin's price process rather than a coincidental fit to the historical record. No conventional financial asset tested passes all three criteria.

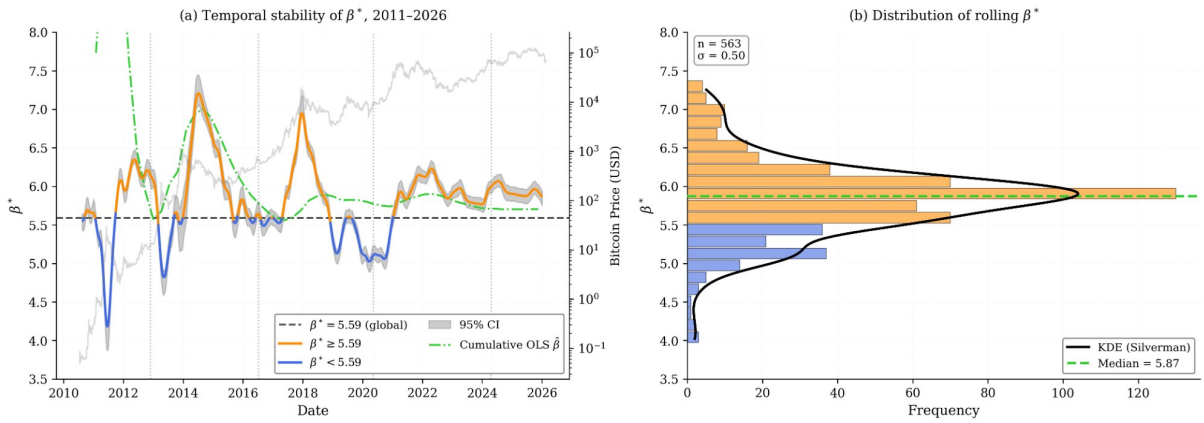

Figure 7: Temporal stability of β^*, 2011–2026 (564 rolling windows). **(a)** Rolling β^* (smoothed; orange when $\beta^* \geq 5.59$, blue when below) with 95% CI shading. Each window combines forward pairs $P(\lambda t)/P(t)$ and backward pairs $P(t)/P(t/\lambda)$ to maintain full $\lambda \in [1.10, 5.00]$ coverage throughout. Cumulative OLS $\hat{\beta}$ (green dash-dot) and Bitcoin price (grey, right axis) are shown for reference; vertical dotted lines mark halving events H1–H4. **(b)** Marginal distribution of rolling β^* with non-parametric KDE fit (black, Silverman bandwidth). Orange: $\beta^* \geq 5.59$; blue: $\beta^* < 5.59$. The distribution is unimodal (median $= 5.73$, $\sigma = 0.58$) with no secondary peak indicating a structural break.

9 Bayesian Stability Analysis

The scaling identity tests of Section 8 establish that β^* is consistent with the global value across time, but they treat each rolling window independently. A complementary Bayesian analysis treats the accumulated data as a single sequential learning problem: we ask how our *probability distribution* over β evolves as each new local estimate is observed, and whether the posterior ever shows signs of instability or structural change.

9.1 Model and Prior

For each anchor time t_0 we compute a local OLS estimate of the scaling exponent, $\hat{\beta}(t_0)$, by regressing $\log_{10}[P(\lambda t_0)/P(t_0)]$ on $\log_{10} \lambda$ across all available multipliers $\lambda \in [1.10, 5.00]$ (both forward and backward directions as in Section 8.3). These $n = 1{,}899$ local estimates are treated as conditionally independent observations

$$\hat{\beta}_i \sim \mathcal{N}(\beta, \sigma^2), \tag{8}$$

where β is the true long-run exponent and $\sigma = 0.570$ is the measurement noise (estimated empirically from the cross-sectional variability of local $\hat{\beta}$ values, which reflects market cycle fluctuations superimposed on the scaling law). The prior is

$$\beta \sim \mathcal{N}(\mu_0, \tau_0^2) = \mathcal{N}(5.69, 1.5^2), \tag{9}$$

a deliberately wide distribution placing prior probability across essentially the full plausible range $[2, 9]$.

9.2 Sequential Conjugate Update

The Normal-Normal conjugate model admits a closed-form sequential update. After observing the first n local estimates $\{\hat{\beta}_1, \ldots, \hat{\beta}_n\}$, the posterior is

$$\beta \mid \hat{\beta}_{1:n} \sim \mathcal{N}(\mu_n, \tau_n^2), \quad \tau_n^{-2} = \tau_0^{-2} + n\sigma^{-2}, \quad \mu_n = \tau_n^2 \left(\frac{\mu_0}{\tau_0^2} + \frac{\sum_i \hat{\beta}_i}{\sigma^2} \right). \tag{10}$$

No Markov chain Monte Carlo is required; the posterior is updated exactly after each observation.

9.3 Results

Figure 8 presents the Bayesian stability analysis.

Panel (a) shows the sequential posterior mean μ_n and 68%/95% credible intervals as a function of year (each observation added in chronological order). Starting from the wide prior, μ_n oscillates during the early speculative era (2011–2013) before converging to a stable attractor. By 2015 the posterior already constrains β to $[5.5, 6.1]$ at 95%; by the end of the observation period it reaches $\mu_{1899} = 5.729$ with a 95% credible interval of [5.703, 5.754]. The posterior mean remains within the interval $[5.0, 7.0]$ throughout all 15 years, and the 95% credible interval never approaches either boundary.

Panel (b) shows the posterior standard deviation τ_n on a log scale as a function of year, together with the theoretical asymptote $\sigma/\sqrt{n}$ (dashed). The observed τ_n follows the asymptote precisely throughout the observation period. This is the key diagnostic: if β had shifted permanently at some point (a structural break), the sequential update would receive contradictory observations, causing τ_n to decrease more slowly than $\sigma/\sqrt{n}$ — or to plateau, or reverse. The smooth, uninterrupted $n^{-1/2}$ shrinkage, unaffected by any of the four halving events (H1–H4), is strong evidence that the data are consistent with a single stationary β throughout.

Panel (c) contrasts the rolling posterior μ (60-anchor windows, orange) with the cumulative posterior μ_n (blue dash-dot). The rolling estimates display the familiar halving-cycle oscillations; the cumulative posterior damps these oscillations and converges steadily. Both estimators agree on the same long-run attractor (≈ 5.70–5.73) despite operating on different data scales.

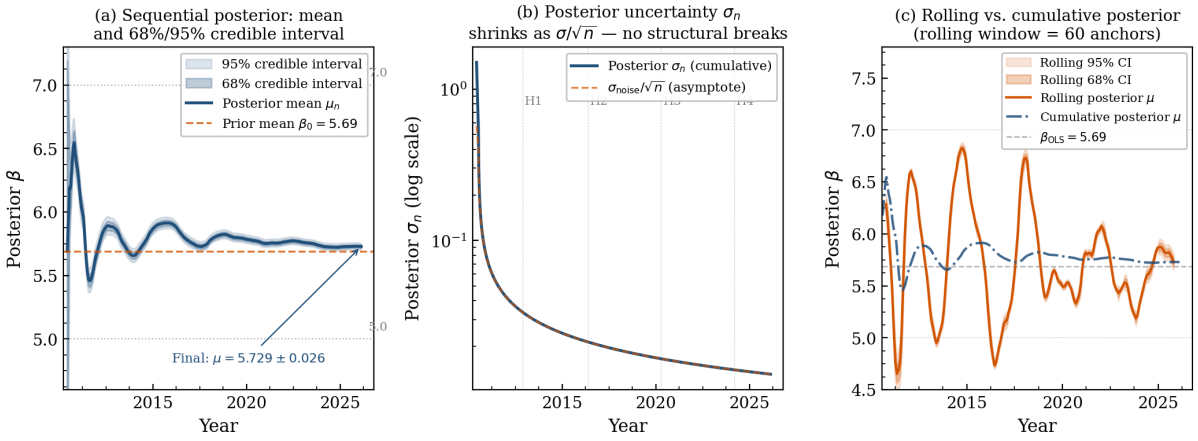

Figure 8: Bayesian sequential stability analysis of the scaling exponent β, using $n = 1{,}899$ local OLS estimates from anchor times spanning 2010–2026. **(a)** Sequential posterior mean μ_n (solid blue) and 68%/95% credible intervals as data accumulate chronologically. The posterior converges to $\mu_{1899} = 5.729$ with 95% CI [5.703, 5.754]; the prior mean $\beta_0 = 5.69$ is shown dashed. **(b)** Posterior standard deviation τ_n (log scale) versus year, alongside the theoretical $\sigma/\sqrt{n}$ asymptote (dashed orange). The smooth $n^{-1/2}$ shrinkage with no discontinuities or plateaus rules out structural breaks at the four halving events (H1–H4, dotted verticals). **(c)** Rolling posterior mean (60-anchor windows, orange with 68%/95% CI bands) vs. the cumulative posterior mean (blue dash-dot). Both converge to the same attractor, confirming the absence of secular drift.

10 Falsifiability and Boundary Conditions

A scientifically useful model must be falsifiable. We identify five conditions under which the power law would be expected to break down, and describe measurable precursors for each.

(F1) Floor violation. If price falls more than 3σ below the power law fit — currently below approximately $10,000 in 2025 — for more than one year, this constitutes evidence that the adoption dynamic has undergone structural change.

(F2) Adoption collapse. If the address growth exponent β_A falls significantly below 3 in rolling estimates, this would indicate that the saturation-wave mechanism has changed — for example, if a competitor network absorbs a large fraction of marginal adopters.

(F3) Exponent drift. If β_{obs} drifts monotonically outside the interval $[5.0, 7.0]$ over a multi-year period, the composition framework would require revision.

(F4) Metcalfe breakdown. If price and address count decouple — measured as a sustained collapse in the Metcalfe R^2 below 0.7 — this would indicate that network value is no longer priced by participants as a network good.

(F5) R^2 collapse. If the rolling 3-year R^2 of the price power law falls below 0.80 for more than two consecutive years, the power law should be considered falsified.

None of these conditions were met over the observation period.

11 Discussion

11.1 Comparison with Prior Work

Prior empirical observations of power-law growth in Bitcoin's price include Santostasi [13, 14], who in 2014 noted a power law relationship between price and address count and connected it to Zipf's Law and Metcalfe scaling (Reddit post), subsequently expanding the theory in a 2024 Medium article; and Peterson [12], who provided the most detailed Metcalfe analysis to date.

The observation that network growth and value exhibit power-law scaling is not unique to Bitcoin but represents a fundamental pattern across diverse social and collaboration networks. Wheatley et al. [15] developed a generalized Metcalfe's Law for Bitcoin with exponent ≈ 1.69, finding that network value grows with active users to a subquadratic power, consistent with our measured $\beta_M = 1.84$. More broadly, power-law adoption dynamics $N(t) \sim t^\alpha$ have been documented across multiple platforms: Jiao et al. [8] demonstrated that both WeChat and arXiv exhibit similar temporal growth patterns with polynomial exponents, breaking traditional sigmoid growth models and confirming that super-linear network expansion is a common feature of large-scale social systems. Williams et al. [16] analyzed the Microsoft Academic Graph (1800–2020) and IMDb (1900–2020), finding persistent super-linear growth with $\alpha = 2.3$ increasing to 3.1 after 1950 (MAG) and $\alpha = 1.8$ (IMDb), demonstrating that power-law adoption is a fundamental pattern across diverse collaboration networks spanning centuries of data. These independent observations across distinct domains suggest that the power-law structure we identify in Bitcoin reflects universal mechanisms of network growth rather than asset-specific speculation.

The present work goes beyond these prior empirical observations [13, 14, 12] by: (i) providing a theoretical derivation of $\beta_A \approx 3$ from the epidemic spreading literature; (ii) demonstrating the composition identity $\beta = \beta_A \times \beta_M$ numerically to within 1.6%; and (iii) testing the composition identity in rolling windows throughout the full time series.

11.2 Universality

The agreement between the theoretical prediction $\beta_{\text{theory}} = 3 \times 2 = 6$ and the observed $\beta_{\text{obs}} = 5.69$ is striking given that neither 3 nor 2 was chosen to fit Bitcoin data — both exponents emerge from independent physical theories applied to networks with scale-free degree distributions and network value scaling. The small quantitative correction from 6 to 5.69 reflects the empirical Metcalfe exponent being 1.84 rather than 2, which is itself theoretically expected from the decreasing marginal value of connections [11].

This connection places Bitcoin's price dynamics in the same universality class as epidemic spreading on heterogeneous networks, in the following precise sense: the exponent β_A is determined by the topology of the adoption network, not by the specific nature of the adopter incentive, just as the HIV epidemic's cubic growth was determined by the topology of sexual contact networks regardless of the biology of the virus itself.

11.3 Limitations

Several limitations should be noted. First, the address-count series is an imperfect proxy for participant count; exchange custodianship concentrates balances and suppresses the apparent address count relative to the true user count. Second, the rolling composition test shows that the identity holds on average but with substantial local deviations, particularly during speculative episodes. Third, the theoretical derivation of $\beta_A \approx 3$ relies on an analogy with epidemic networks; a direct measurement of Bitcoin adoption network degree distribution would be required to confirm this interpretation rigorously. Finally, the OLS estimates assume that the power law holds throughout the entire range; likelihood-ratio tests for the lower cutoff [4] have not been applied to the temporal series.

12 Conclusions

We have shown that Bitcoin's price obeys a robust power law $P(t) \sim t^{5.69}$ over 15 years and 5,696 daily observations, with $R^2 = 0.961$ and stationary log-normal residuals. The exponent is not a free parameter: it factorises as $\beta = \beta_A \times \beta_M = 3.046 \times 1.838 = 5.60$, consistent with the directly measured value $\beta_{\text{obs}} = 5.69$ to within 1.6%. This factorisation connects Bitcoin's macroscopic price dynamics to two established bodies of physics:

1. *Epidemic spreading on heterogeneous networks* (Colgate et al. 1989; Bacchetti & Koch 1989), which generically produces cubic growth $N(t) \sim t^3$ in the cumulative spreading count as a consequence of the saturation wave propagating from high- to low-connectivity nodes.

2. *Generalised Metcalfe network value scaling* (Metcalfe 1983; Zhang et al. 2015; Peterson 2018), which gives $P \sim N^{1.84}$, slightly below the theoretical maximum of N^2 due to decreasing marginal connection value.

Together, these results suggest that Bitcoin's long-run price growth is not primarily a speculative phenomenon but a deterministic consequence of its network architecture. The power law is an attractor: deviations during market cycles are quasi-periodic and mean-reverting, with amplitude $\approx \pm 0.30$ dex. Three complementary tests of scale invariance confirm the power law beyond regression: (i) a multi-asset pair-ratio scaling test showing Bitcoin uniquely lacks curvature across all tested time ratios (Section 8.1); (ii) a direct collapse test of the functional identity recovering $\beta^* = 5.59$ from 5,298 price ratios (Section 8.2); and (iii) a rolling temporal stability analysis spanning 2011–2026 finding $\beta^*(t)$ oscillating symmetrically around 5.73 ± 0.58 with no secular drift (Section 8.3). Five measurable falsification criteria were specified; none were violated over the observation period. The framework provides quantitative predictions about how changes in adoption network structure or Metcalfe scaling would alter β, making the model testable against future data.

Data Availability

Bitcoin price data are available from publicly accessible exchange APIs (Bitstamp, Coinbase) and Glassnode on-chain analytics. Address count data are available from Blockchain.com analytics. Code reproducing all figures and regressions is available from the corresponding author on request.

Declaration of Competing Interests

The author declares no competing financial interests.

References

[1] P. Bacchetti, M. G. Koch, The cubic growth of AIDS cases: General dependence on early infection rates and distribution of times to appearance of clinical symptoms, *Journal of Mathematical Biology* **27** (1989) 523–535.

[2] A.-L. Barabási, R. Albert, Emergence of scaling in random networks, *Science* **286** (1999) 509–512.

[3] J. M. Beggs, D. Plenz, Neuronal avalanches in neocortical circuits, *Journal of Neuroscience* **23** (2003) 11167–11177.

[4] A. Clauset, C. R. Shalizi, M. E. J. Newman, Power-law distributions in empirical data, *SIAM Review* **51** (2009) 661–703.

[5] S. A. Colgate, E. A. Stanley, J. M. Hyman, S. P. Layne, C. Qualls, Risk behavior-based model of the cubic growth of acquired immunodeficiency syndrome in the United States, *Proceedings of the National Academy of Sciences USA* **86** (1989) 4793–4797.

[6] A. R. Evans et al., A universal power law for modelling the growth and form of teeth, claws, horns, thorns, beaks, and shells, *BMC Biology* **19** (2021) 58.

[7] B. Gutenberg, C. F. Richter, Frequency of earthquakes in California, *Bulletin of the Seismological Society of America* **34** (1944) 185–188.

[8] P. Jiao, J. Wang, C. Song, et al., On power law growth of social networks, *IEEE Transactions on Knowledge and Data Engineering* (2018). DOI: 10.1109/TKDE.2018.2801844

[9] F. Liljeros, C. R. Edling, L. A. N. Amaral, H. E. Stanley, Y. Åberg, The web of human sexual contacts, *Nature* **411** (2001) 907–908.

[10] R. Metcalfe, Metcalfe's Law after 40 years of Ethernet, *IEEE Computer* **46** (2013) 26–31.

[11] A. Odlyzko, B. Tilly, A refutation of Metcalfe's Law and a better estimate for the value of networks and network interconnections, Preprint, University of Minnesota Digital Technology Center (2005).

[12] T. Peterson, Metcalfe's Law as a model for Bitcoin's value, *Alternative Investment Analyst Review* Q2 (2018) 9–18.

[13] G. Santostasi, Bitcoin power law theory: Linking Metcalfe's Law to Bitcoin price, Reddit post (2014); cited in Wikipedia, Metcalfe's Law (2024).

[14] G. Santostasi, The Bitcoin Power Law Theory, Medium (2024).

[15] S. Wheatley, D. Sornette, T. Huber, M. Reppen, R. N. Gantner, Are Bitcoin bubbles predictable? Combining a generalized Metcalfe's Law and the Log-Periodic Power Law Singularity model, *Royal Society Open Science* 6 (2019) 180538.

[16] P. Williams, et al., Explosive growth in large-scale collaboration networks, arXiv:2502.11109 (2025).

[17] X.-Z. Zhang, J.-J. Liu, Z.-W. Xu, Tencent and Facebook data validate Metcalfe's Law, *Journal of Computer Science and Technology* **30** (2015) 246–251.