Bitcoin's Price Power Law Decomposed: Epidemic Speed Times Wealth Inequality

A response and extension to Santostasi & Perrenod (2026)

v2.0 — July 2026

bitcoin-trajectory.pages.dev

wineLightning Claude AI

Abstract

Bitcoin's USD price has followed a power law of its age, $P \propto t^{\beta}$, for sixteen years and six orders of magnitude. Santostasi & Perrenod establish the law and decompose its exponent as $\beta = \beta_A \times \beta_M$, deriving $\beta_A \approx 3$ from epidemic spreading but leaving $\beta_M \approx 1.84$ as an unexplained "generalised Metcalfe" exponent. We answer their open question: $\beta_M = 1 + 1/\alpha$, where $\alpha \approx 1.2$ is the Pareto exponent of global wealth, measured by economists independently of Bitcoin. The mechanism is *geometric penetration of the wealth ranks*: each e-folding of adoption reaches a constant factor deeper into the wealth distribution, and the capped supply forces each richer wave to buy from earlier holders. Simulation shows that this schedule — and not mere rank ordering, which provably fails at Bitcoin's penetration — produces the super-linearity; on-chain data show it operating. With autocorrelation-honest errors the measured scaling, $\beta_M = 1.827 \pm 0.091$, sits dead-centre in the externally predicted band $[1.80, 1.87]$. Cross-asset controls turn the claim into a prediction of *which* assets obey the law: Ethereum and Litecoin show the same generic adoption power laws but zero wealth surplus — exactly as required, since one lacks the supply cap and the other a live epidemic. Finally, the temporal exponent inherits a time-origin convention while β_M is convention-free, resolving the recent "weak structure, strong forecasts" paradox: the invariant object of the Bitcoin power law is β_M — and β_M is the shape of human wealth inequality.

1 Introduction

Since 2010, the USD price of Bitcoin has tracked a power law of its age: on doubly logarithmic axes, sixteen years of daily prices — spanning six orders of magnitude, two 80% drawdowns, successive regulatory bans, exchange failures, and four supply halvings — lie along one straight line (Fig. 1). The empirical law is due to Santostasi, refined over a decade and formalised in Santostasi & Perrenod [1], who further decompose the exponent as

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

where $N \propto t^{\beta_A}$ is the growth of the adopter base and $P \propto N^{\beta_M}$ the scaling of price with adoption. They derive $\beta_A \approx 3$ from epidemic spreading theory [3] and measure $\beta_M \approx 1.84$, which they label a generalised Metcalfe exponent. The label is a placeholder, not a mechanism: nothing in network economics selects 1.84, explains its stability, or says which other assets should exhibit it.

This paper proposes and tests the missing mechanism. Our claims, each argued and quantified in its own section:

1. **The formula.** $\beta_M = 1 + 1/\alpha$, where α is the Pareto tail exponent of global wealth — measured on tax and wealth-report data for over a century [10, 11, 12, 13], with no reference to Bitcoin. The literature band for α predicts $\beta_M \in [1.80, 1.87]$; the measured value is 1.83 (Sec. 7).

2. **The mechanism, sharpened.** The derivation requires *geometric tail penetration* — the k-th adopter near wealth rank M/k. Weaker notions of "the poor adopt first" (rank correlation, even perfect) provably fail at Bitcoin's $< 1\%$ penetration (Sec. 6).
3. **The mechanism, observed.** Capital per holder grew $\sim\times 2{,}500$ over fourteen years, as a clean power of the holder count — where order-free adoption predicts a flat line (Sec. 8).
4. **Cross-asset controls.** Adoption power laws are generic across chains; the wealth surplus exists *only* for Bitcoin, and Ethereum and Litecoin each lack exactly the ingredient the theory says they lack. Litecoin, alive since 2011, also disposes of the survivorship objection (Sec. 9).
5. **Honest statistics, and the right invariant.** Sixteen years of daily data contain only $N_{\text{eff}} \approx 6$ independent observations; we quote errors accordingly. And the temporal exponent inherits a time-origin convention while β_M involves no clock: the convention-free object of the Bitcoin power law is β_M — precisely the quantity the formula explains (Sec. 4). This resolves the "weak structure, strong forecasts" paradox of [16].

Throughout, we follow two rules adopted after an adversarial internal review: every exponent is quoted with autocorrelation-honest uncertainty, and every discarded interpretation — including two of our own earlier evidence items — is documented (Appendix C).

Because the pieces of evidence differ in kind, we state their status up front (Table 1): three results are load-bearing; the rest are consistency checks or corroboration, and the argument survives without them.

Evidence	Establishes	Status
$\beta_M = 1.827 \pm 0.091$ vs. external band $[1.80, 1.87]$ (Sec. 7)	the number match	**load-bearing**
t_0 scan: β spans 4.8–5.8, β_M clock-free (Sec. 4)	the right invariant	**load-bearing**
Surplus $+0.84$ for BTC, ≈ 0 for ETH/LTC (Sec. 9)	predicts *which* assets	**load-bearing**
Ordering simulation: rank ordering fails, geometric schedule works (Sec. 6)	mechanism necessity	theory + simulation
Capital/holder $\propto N^{0.9\text{–}1.0}$ (Sec. 8)	mechanism operating	consistency†
2024 ETF break: proxies snap, law unmoved (Sec. 10)	law is structural	consistency
$\beta_A \approx 3$ three ways; AIDS and domains $\approx t^3$; mobile phones exponential (Sec. 10)	first factor generic *and* gated	supporting
Geography of adoption; exponent zoo; lead–lag chain (Sec. 10, App. A)	corroboration	illustration

Table 1: The case at a glance. †Realized cap tracks price, so the schedule is a within-Bitcoin consistency check — its force is discriminating ordering models, not independently re-deriving β_M.

2 Related work

Santostasi & Perrenod (2026). Our foundation. They establish $P \propto t^{5.69}$ ($R^2 = 0.96$), validate scale invariance with pair-ratio, collapse, rolling-window and Bayesian tests, and derive $\beta_A \approx 3$ from epidemic dynamics on heterogeneous human networks, building on Colgate's analysis of the cubic growth of the early AIDS epidemic [3]. Their $\beta_M \approx 1.84$ is measured, not derived. We adopt Eq. (1) and their β_A result wholesale, and supply the second factor.

Metcalfe-type valuations. Metcalfe's law ($V \propto N^2$) is sound in its home domain — the value of a *communication* network to its members grows with the number of possible connections — and generalised Metcalfe fits to Bitcoin exist [14]. Three problems arise when it is used to explain β_M. Empirically, 2.0 sits $+1.9\sigma$ from the measured 1.827 ± 0.091: disfavoured, though not excluded by the numbers alone. Mechanistically, in 2009–2010 the network existed and grew while the price remained zero; value tracked belief and capital, not connectivity. Structurally, connection-value is what makes the idea spread — it is already inside β_A — so placing the network in both factors of Eq. (1) counts it twice. Our formula removes the double count: β_A lives on the network of minds, β_M on the distribution of wallets.

A fourth objection is conceptual and, we think, decisive. Metcalfe prices *connections* — the value of being able to transact with more members. But what drives Bitcoin's price today is not a transactional decision; it is a *capital-allocation* decision: how much of one's savings to store in the asset. In the monetary-good framing of Menger [8] and Ammous [9], a money monetises in stages — store of value, then medium of exchange, then unit of account — and Bitcoin is winning the first. Metcalfe values the second stage (the transaction network); β_M must instead capture the first (how much capital each new believer chooses to store). That quantity is governed by the distribution of wealth, not by connectivity — which is precisely the substrate our formula identifies.

Stock-to-Flow. A supply-side model [15] that conflated scarcity with valuation and failed out of sample in 2021–2022. In our account scarcity is a necessary ingredient (it forces each wave to buy from earlier holders) but carries no exponent of its own; demand-side structure (adoption $\times$ wealth) does.

The skeptical literature. Baquero & Menezes [16] show that the temporal power law has weak distributional support and a time-origin sensitivity, yet out-forecasts more complex models at 12–24 month horizons. We confirm both findings independently (Sec. 4) and resolve the tension: the fragile quantity is the convention-laden temporal exponent; the robust quantity is the clock-free price–adoption exponent, and that is the one with a theory behind it.

Epidemic and polynomial spreading. Sub-exponential, polynomial epidemic growth on structured human networks is generic: Colgate's t^3 [3], clustering-driven polynomial growth [4], and empirical reviews across diseases [5]. Adoption power laws with exponents ≈ 1.8–4 also appear in non-monetary networks such as academic collaboration graphs [6] — a point we exploit as a control: adoption power laws are cheap; the price surplus is not. A useful distinction: the polynomial law is *intrinsic* when transmission requires a repeated decision that confines spread to stratified sub-communities (HIV's risk groups, Colgate's original case; Bitcoin's cognitive/wealth tiers). A freely airborne contagion in a well-mixed population instead grows exponentially, turning polynomial only under intervention — COVID-19 confirmed cases in China followed $t^{2.1\pm0.3}$, which Maier & Brockmann [7] attribute to containment depleting the susceptible pool. This is why HIV, not influenza, is the right biological analogue for Bitcoin: both spread by a choice.

Wealth distribution. The Pareto tail of wealth is one of the oldest quantitative regularities in economics [10]. Modern global estimates place the tail exponent at $\alpha \approx 1.15$–1.25 [11, 12, 13]. We take α from this literature as an external input.

3 Data and methods

Data. Daily BTC/USD from blockchain.info (2010-08-18 to 2026-06-30, 5,796 observations); adopter proxy N = addresses with non-zero balance (CoinMetrics community, `AdrBalCnt`), the structurally superior proxy also used by S&P; realized capitalisation (Newhedge, to 2026-04);

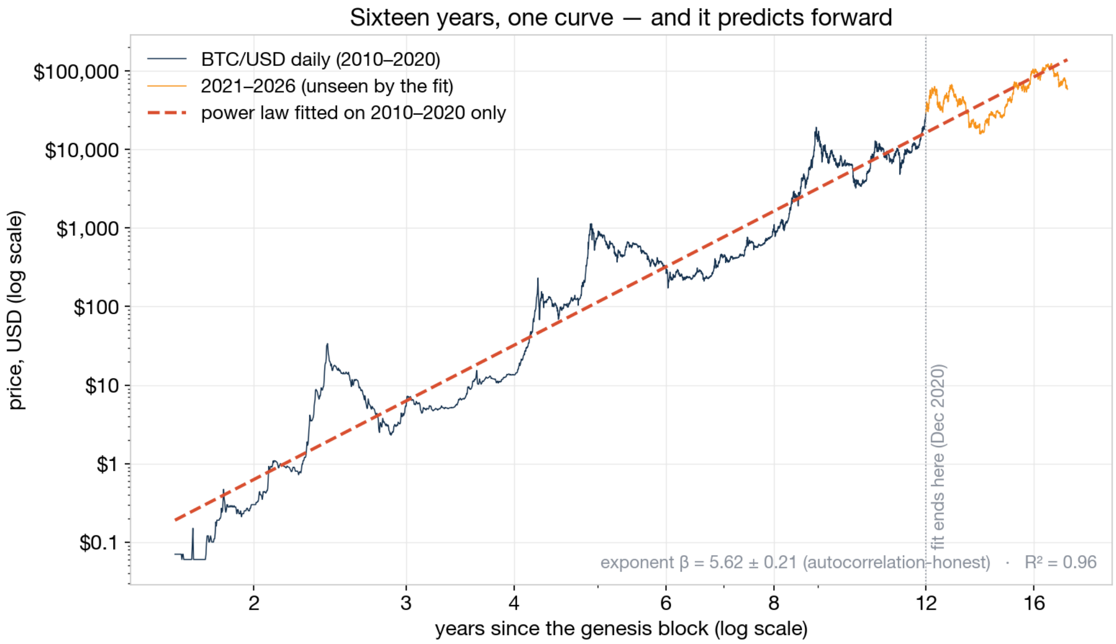

Figure 1: Bitcoin's temporal power law. Fitted on 2010–2020 only (red dashed), the law predicts 2021–2026 with a mean out-of-sample residual of -0.08 dex ($\approx -17\%$) through the 2021 mania, the 2022 collapse, and the ETF era. Full-period exponent $\beta = 5.62$, $R^2 = 0.96$.

PriceUSD and AdrBalCnt for ETH and LTC (CoinMetrics community). Wikipedia page views and Google Trends for narrative-side adoption measures. All series and scripts are public; every figure regenerates from raw data with one command (Appendix E).

Estimation. Exponents are OLS slopes in $\log_{10}$–$\log_{10}$. Time origins: each chain's genesis unless stated. Uncertainty: Newey–West (HAC) standard errors with 1–4 year bandwidths and moving-block bootstrap of residuals (blocks 0.5–2 years, $B = 2{,}000$); we quote the most conservative. Residual AR(1) autocorrelation is $\rho = 0.998$, giving $N_{\text{eff}} = N(1-\rho)/(1+\rho) \approx 6$: sixteen years of daily data are worth roughly six independent points — about the number of market cycles. Naive OLS errors (± 0.015 on β) understate uncertainty by a factor ≈ 14 and are reported only for contrast.

4 The law, measured honestly

Point estimates. Over the full period, $\beta = 5.623$, $\beta_A = 3.020$, $\beta_M = 1.827$ (all genesis-anchored where time is involved). Honest uncertainties: $\sigma(\beta) = 0.211$, $\sigma(\beta_M) = 0.091$.

Forward test. A fit restricted to 2010–2020 predicts 2021–2026 with mean out-of-sample residual -0.08 dex (Fig. 1) — consistent with the forecasting performance documented independently in [16].

The time origin is a convention; β_M is not. $P \propto (t - t_0)^\beta$ requires an origin. Genesis (2009-01-03) is principled — the epidemic's patient zero — but the data cannot certify it: the $R^2(t_0)$ profile is flat ($\Delta R^2 = 0.0014$ between genesis and the optimum), while β moves from 5.80 (whitepaper anchoring) through 5.62 (genesis) to 4.82 (first traded price) (Fig. 2). Two quantities, however, are stable: the ratio $\beta/\beta_A \in [1.843, 1.864]$ across those conventions ($\sim$1% variation, because β and β_A shift together), and the directly regressed $\beta_M = 1.827$, which involves no clock at all. *The convention-free object of the Bitcoin power law is the price–adoption exponent*

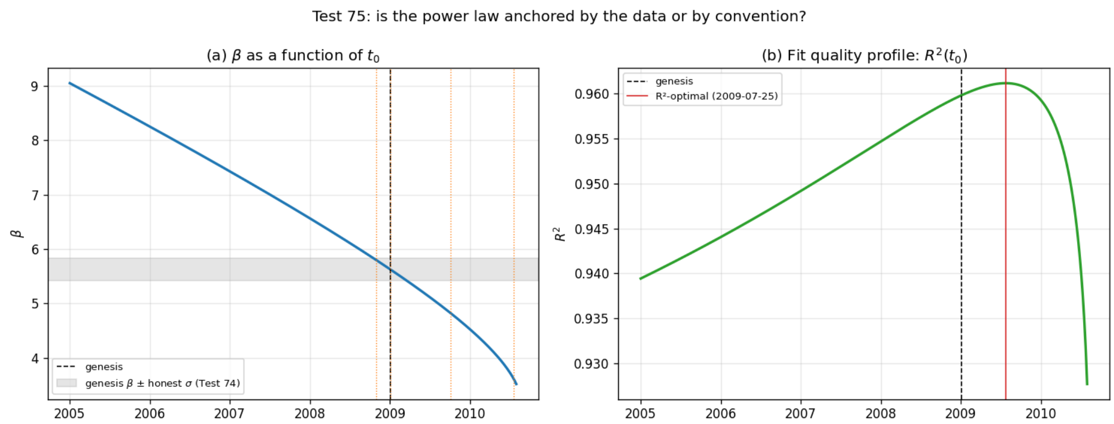

Figure 2: Time-origin sensitivity. Left: β as a function of the assumed origin t_0; the grey band is the honest $\pm\sigma(\beta)$ at genesis anchoring. Right: the R^2 profile is nearly flat across 2008–2009 — the data do not pin t_0. The ratio β/β_A moves by $\sim$1% across the same range.

β_M. The temporal law is then the composition of the clock-free relation $P \propto N^{\beta_M}$ with the genesis-anchored epidemic law $N \propto t^{\beta_A}$.

5 The decomposition and its double-counting problem

The adoption factor, three ways. The epidemic exponent is measured on three observables sharing no data source: balance-holding addresses ($\beta_A = 3.02$, $R^2 = 0.97$), cumulative views of the Bitcoin Wikipedia article (3.07, $R^2 = 0.91$), and cumulative worldwide Google searches (≈ 3.6, $R^2 = 0.99$) (Fig. 3). Awareness, curiosity and action all grow as $\approx t^3$, consistent with the Colgate class; the same class appears in Ethereum (2.5), Litecoin (2.2) and academic collaboration networks (2.3–3.1) [6]. Epidemic adoption is generic within the decision-gated class — and only there: gate-free mobile telephony grew exponentially instead (Test 79, Sec. 10) — so the first factor cannot be what distinguishes Bitcoin. Neither, in fact, is the second factor's *value*, which is set by human wealth inequality rather than by anything Bitcoin did. What is Bitcoin's own is the surplus's *existence*: it switches on only where a hard supply cap meets an epidemic still climbing the wealth pyramid (Sec. 9). As an external anchor, the same law describes a genuine biological epidemic: cumulative US AIDS cases (Colgate's original subject) fall on a log-log line ($R^2 = 0.995$) with exponent ≈ 2.7–3.0 depending on the assumed onset (Fig. 4) — the same time-origin sensitivity we disclose for Bitcoin, and squarely in the Colgate class.

Component drift is mechanical. In expanding windows the components drift in opposite directions while the product stays put; indeed $\mathrm{corr}(\beta_M, \beta/\beta_A) = 0.98$ across windows: the windowed β_M carries almost no information beyond β and β_A. A corollary worth stating because we fell for it ourselves: an expanding-window "convergence" of any quantity derived from β_M (for instance an implied α) is a mirror of the adoption proxy's settling under a quasi-constant β, not a physical trajectory. Our earlier signature plot of that kind is retracted in Appendix C and replaced by Fig. 5; only full-period quantities are used as evidence in this paper.

The itch. In S&P's reading, β_A is network spreading and β_M is network value. If both factors are "the network", the decomposition risks explaining one thing twice — the double-encoding problem. The resolution requires the second factor to live on a different object than the first. That object, we now argue, is the wealth distribution.

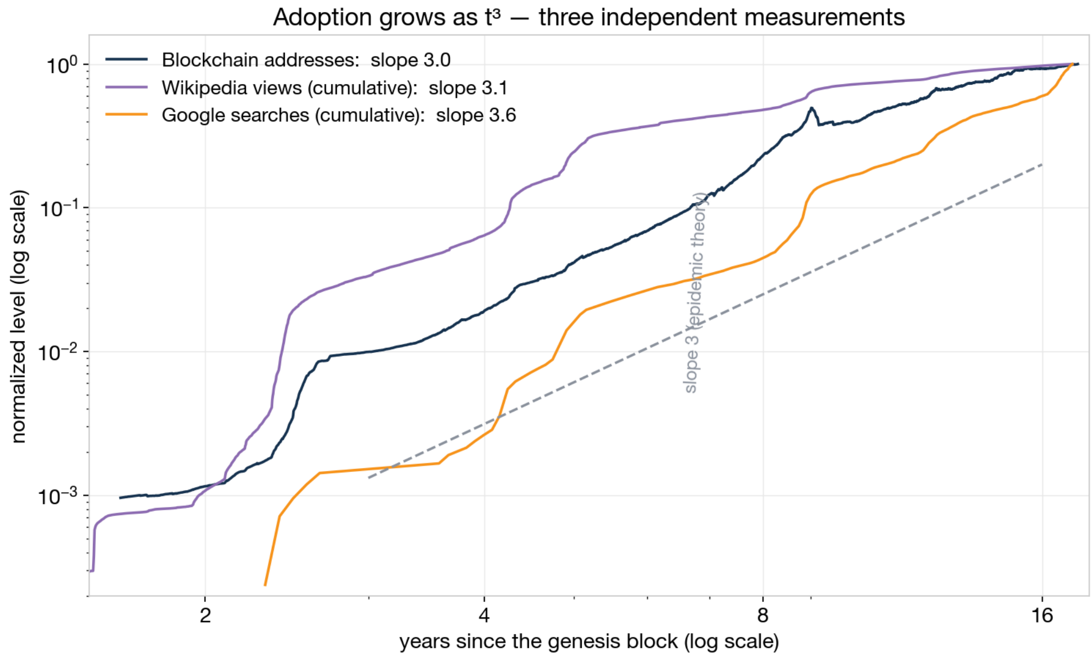

Figure 3: The adoption power law from three independent observables (each normalised to its final value). Slopes 3.0–3.6; the dashed guide has slope 3.

6 Theory: $\beta_M = 1 + 1/\alpha$

6.1 Setup and derivation

Let a population of M individuals have wealth drawn from a Pareto distribution with tail exponent α: the individual at rank r from the top holds

$$w(r) \;=\; w_{\min}\left(\frac{M}{r}\right)^{1/\alpha}. \tag{2}$$

Assume (H, *geometric tail penetration*): when N individuals have adopted, the deepest wealth rank reached is $r_{\min}(N) \simeq M/N$ — equivalently, the k-th adopter sits near rank M/k, so

$$w_k \;\simeq\; w_{\min}\, k^{1/\alpha}. \tag{3}$$

If each adopter allocates a fixed fraction of wealth, cumulative invested capital is

$$V(N) \;=\; \sum_{k\le N} w_k \;\simeq\; \frac{w_{\min}}{1+1/\alpha}\, N^{\,1+1/\alpha}, \tag{4}$$

and with a hard-capped supply S absorbing net inflow, $P \propto V/S$, hence

$$\boxed{\beta_M \;=\; 1+\frac{1}{\alpha}} \qquad \Longrightarrow \qquad \beta \;=\; \beta_A\left(1+\frac{1}{\alpha}\right). \tag{5}$$

The "1" is the equal-world baseline (price tracks adoption); the "$1/\alpha$" is the inequality surplus. For $\alpha = 1.22$, $\beta_M = 1.82$ and, with $\beta_A = 3.02$, $\beta = 5.52$ — within 2% of the directly measured 5.62, the gap being the expected inconsistency between three separate regressions and well inside $\sigma(\beta) = 0.21$.

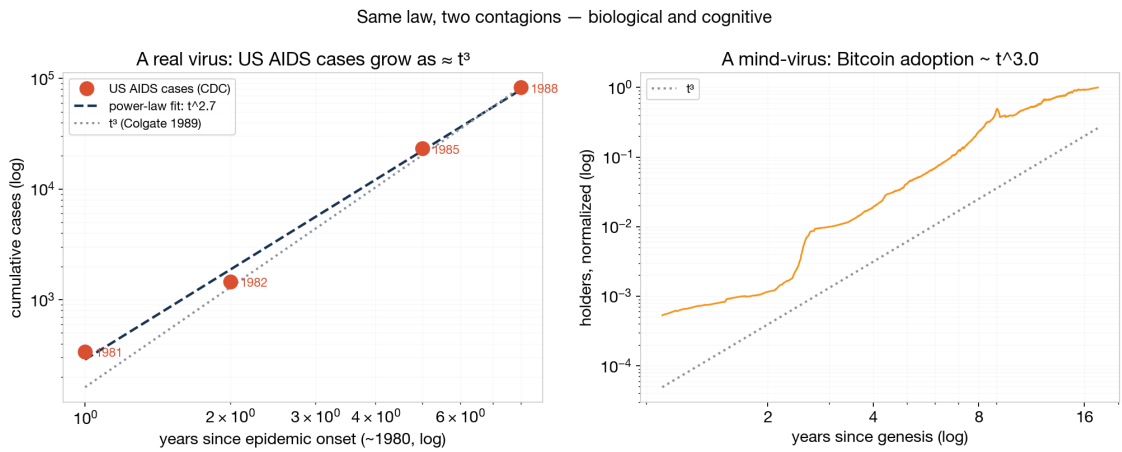

Figure 4: Same law, two contagions. Left: cumulative US AIDS cases (four CDC-sourced anchor points, 1981–1988) grow as $\approx t^3$ ($R^2 = 0.995$), not exponentially. Right: Bitcoin adoption, $t^{3.0}$. A biological and a cognitive contagion, spreading on the same clustered human network, obey the same power law of time.

6.2 What ordering the formula actually needs

"Later adopters are richer" admits two formalisations with opposite outcomes (Fig. 6):

- *Rank correlation* (a Gaussian copula of correlation ρ between wealth and adoption time): for every ρ from 0 to 1 — including a *perfect* bottom-up sweep — the fitted β_M is 1.00–1.08 at Bitcoin-like penetration ($\leq 1\%$ of the population). A sweep protects the rich tail until near saturation; the surplus of sequential ordering lives entirely in the last few percent of adoption, a regime Bitcoin has never been in.
- *The geometric schedule* (H): yields $1 + 1/\alpha$ exactly, and is robust to lognormal rank noise of about one decade ($\beta_M = 1.795$ at $\sigma = 1$ decade for $\alpha = 1.22$), degrading beyond (1.44 at $\sigma = 2$).

This sharpening matters twice over. It corrects the folk statement of the mechanism ("the poor buy before the rich" is insufficient); and it converts a vague assumption into a measurable one: under (H), the wealth depth reached — and with it the capital carried per adopter — must grow as a power of N. Section 8 measures exactly that. A sanity check on the schedule's realism: with $M \approx 5 \times 10^9$, rank M/N places the richest adopter near rank 5×10^6 at $N = 10^3$ (2010: multi-millionaire technologists), 5×10^3 at $N = 10^6$ (2013–14: centimillionaire investors), and ~ 50 at $N = 10^8$ (2024: the largest asset pools on Earth). The public adoption timeline matches.

6.3 The double count, resolved

In Eq. (5) the two factors live on different objects. β_A is the speed of an idea-epidemic on the network of *minds*: who hears, who believes — Colgate-class dynamics, generic across decision-gated contagions. β_M is a property of the distribution of *wallets*: how much capital the marginal believer commands, set by α, a quantity estimated from tax and wealth data with no reference to any network. The network appears once. As a corollary, monetary debasement cancels from the nominal exponent: rescaling all wealth $W \to \lambda W$ leaves α invariant, so Eq. (5) predicts the *nominal* exponent directly; the sizeable literature we produced on M2 deflators is thereby absorbed (Appendix D).

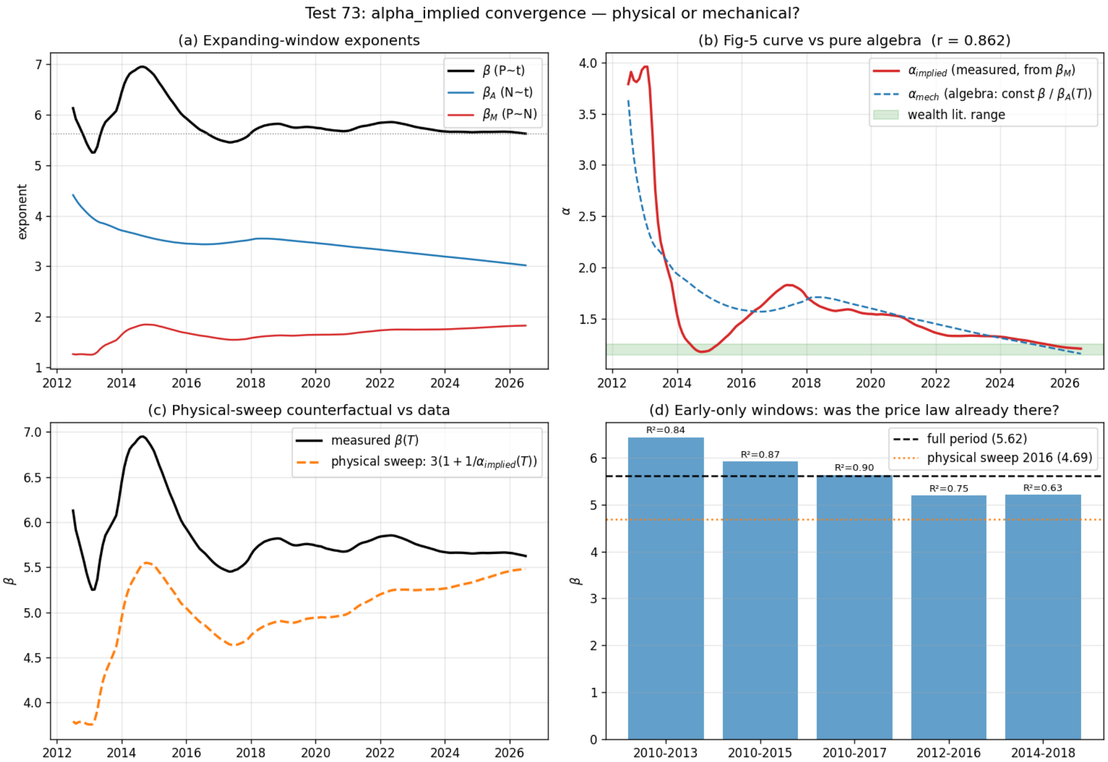

Figure 5: Expanding-window mechanics. (a) The exponents settle as the sample grows; (b) the "implied α" path is reproduced by pure algebra with constant β ($r = 0.86$; and $\mathrm{corr}(\beta_M, \beta/\beta_A) = 0.98$); (c) a physically time-resolved wealth-sweep counterfactual moves opposite to the data; (d) early windows already show the full-strength price law. Windowed component paths are measurement mechanics, not physics.

7 Confirmation I: the number

Inverting the full-period measurement, $\alpha_{\text{implied}} = 1/(\beta_M - 1) = 1.209$, dead-centre of the wealth-literature band $[1.15, 1.25]$ (Fig. 7). Equivalently, the externally predicted band $\beta_M \in [1.80, 1.87]$ contains the measured 1.827 ± 0.091 with $z = -0.08$ at the central literature value. The alternatives score poorly: pure Metcalfe ($\beta_M = 2.0$) sits at $z = +1.9$ — disfavoured by the numbers and excluded by the double-count argument of Secs. 2–5 — and Zipf (1.5) at $z = -3.6$, excluded outright. The asymmetry against Metcalfe is Bayesian rather than frequentist: the wealth theory committed ex ante to a narrow band and hit its centre, while 2.0 sits 1.9σ out — on the numbers alone 2.0 remains viable, and it is the double-count argument that removes it. We emphasise what this is and is not: a single-number coincidence of this precision has perhaps a 5% a priori probability against a diffuse alternative; it is strong support, not proof. The mechanism evidence (Secs. 8, 9) carries the rest of the case.

Which α? Three Pareto exponents appear in this paper and should not be conflated. The formula's input is the exponent of the wealth pool the epidemic *recruits from* — global fiat wealth, $\alpha \approx 1.2$, fixed by the inequality literature before looking at Bitcoin. The *on-chain* exponent ($\alpha = 1.07$, Test 62) is not an alternative input but an output: the recruited capital re-organised inside Bitcoin. The schedule slope of Sec. 8 ($N^{0.92}$–$N^{1.03}$) is a third, proxy-biased view that brackets both $1/\alpha$ candidates (0.82 fiat, 0.93 on-chain). The conclusion is not sensitive to the choice — any $\alpha \in [1.07, 1.25]$ predicts $\beta_M \in [1.80, 1.93]$, within $\sim 1.2\sigma$ of the measured 1.827 ± 0.091 — but the *a priori* input, and the one all headline numbers use, is the fiat value. We state the adverse case plainly: if the on-chain α governed recruitment, the prediction would

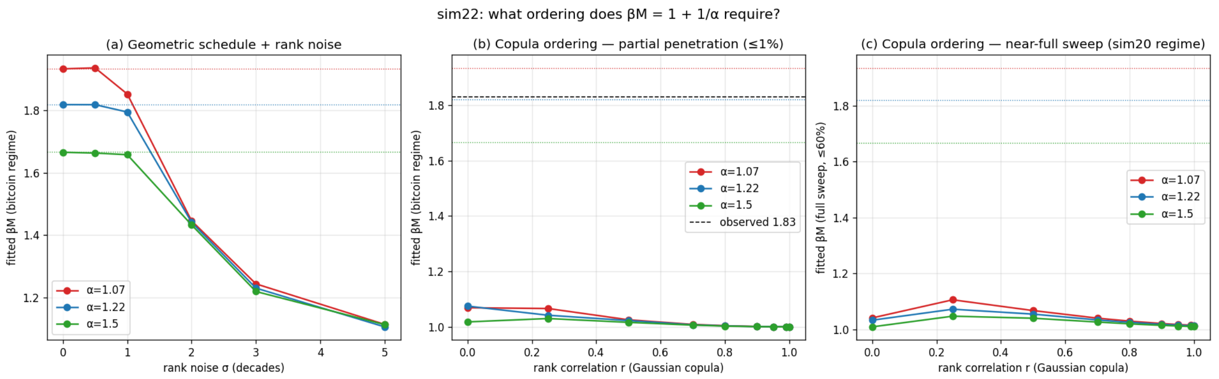

Figure 6: Simulation: which ordering produces the surplus ($M = 10^7$, Pareto wealth, capital-sum observable). (a) The geometric schedule gives $\beta_M = 1 + 1/\alpha$ with a plateau to $\sim$1 decade of rank noise. (b,c) Copula rank-ordering — even perfect — gives $\beta_M \approx 1$ at partial and near-full penetration alike.

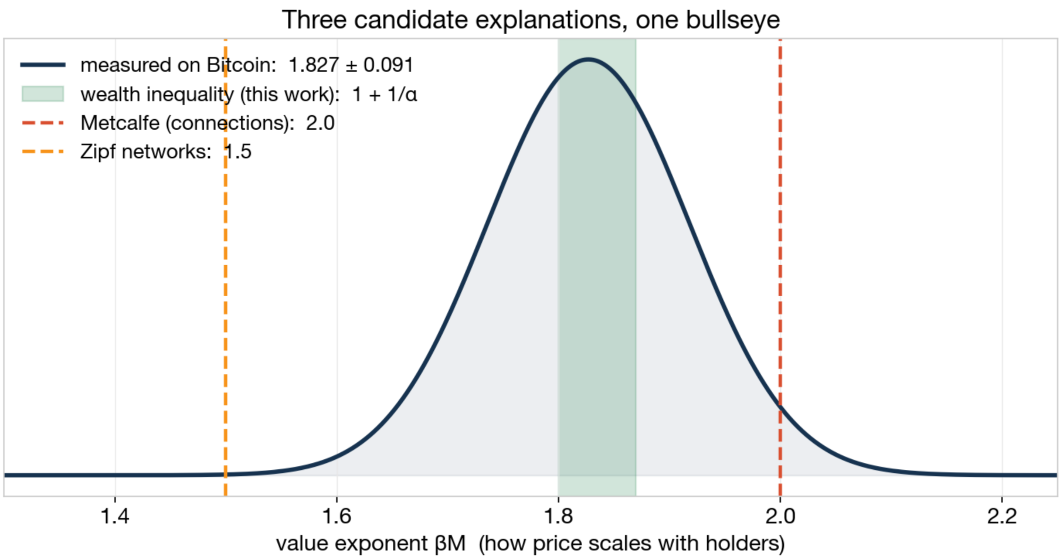

Figure 7: The measured value-scaling exponent (navy: point estimate with honest uncertainty) against three candidate theories. The green band is computed from the wealth-inequality literature with no Bitcoin data.

be 1.93 — indistinguishable from Metcalfe's 2.0 at current errors. What privileges the fiat value is not that it lands on the target but that it is the ex-ante input: the exponent of the *source* distribution being penetrated, fixed by the inequality literature before any Bitcoin data existed, whereas the on-chain exponent is itself a product of the adoption process the formula describes.

8 Consistency check: the schedule, observed

Under (H), capital per holder must grow as $V(N)/N \propto N^{1/\alpha}$; under random or rank-correlated adoption it is flat. Using realized capitalisation (coins valued at last on-chain move — a cumulative inflow proxy) per balance-holding address (Fig. 8): capital per holder rose $\sim\times$2,500 over fourteen years, scaling as $N^{0.92}$ (2010–2017), $N^{1.03}$ (pre-ETF full window). Order-free models are off by three orders of magnitude. The measured slope sits $\sim$10% above $1/\alpha_{\text{on-chain}} = 0.93$ (using the on-chain wealth Pareto $\alpha = 1.07$ that the allocated capital itself follows) and $\sim$25% above the fiat-wealth prediction 0.82; two known biases push it up (realized-cap re-marking of moved coins;

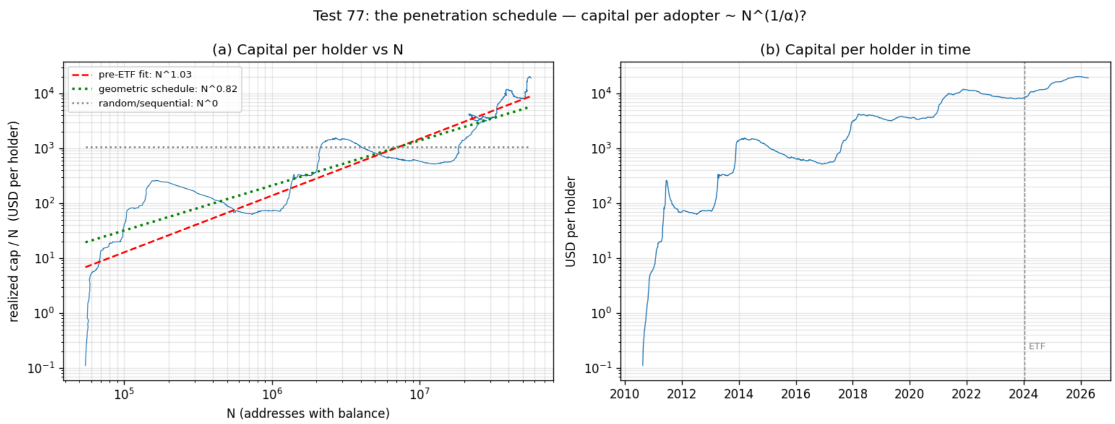

Figure 8: The penetration schedule, measured. Invested capital per holder against holder count (log–log). Flat-ordering models predict the dotted horizontal line.

address consolidation suppressing the denominator — the 2018–2026 sub-window, at $N^{1.72}$, shows the proxy-degradation era). Within proxy limitations, the schedule is real, power-law shaped, sustained, and of the predicted magnitude. We flag the honest caveat: realized price tracks price, so this is a consistency measurement within Bitcoin data — its force is discriminating ordering models (power vs. flat), not independently re-deriving 1.82.

9 Confirmation II: cross-asset controls

A theory of *why Bitcoin* must also say *why not others*. Eq. (5) requires two ingredients for the surplus $\beta_M - 1 > 0$: a hard supply cap (new capital must buy existing coins) and an epidemic still climbing the wealth tiers (the schedule H). We ran the identical pipeline — same estimator, same proxy definitions, each chain's own genesis — on Ethereum (epidemic, no cap) and Litecoin (cap, epidemic that never left the retail tier).

	β_A (R^2)	β_M (R^2)	surplus $\beta_M{-}1$	β (R^2)	CV(β), 8y	forward (oos, dex)
Bitcoin	3.03 (0.98)	1.84 (0.95)	**+0.84**	5.66 (0.96)	1.2%	−0.08
Ethereum	2.48 (0.95)	0.84 (0.88)	−0.16	2.13 (0.86)	5.6%	+0.05
Litecoin	2.21 (0.93)	0.93 (0.72)	−0.08	1.98 (0.62)	6.8%	−0.13

Table 2: Identical pipeline, three assets (data to 2026-06-30). Adoption power laws are generic; the wealth surplus exists only where the theory requires it, and matches $1/\alpha = 0.82$ within 3%.

Three facts (Table 2, Fig. 9): adoption power laws are generic; the surplus is Bitcoin-only and of the predicted size; and the two negatives separate the ingredients — scarcity alone (LTC) is insufficient, an epidemic alone (ETH) is insufficient. Ethereum's price exponent decomposes as $2.48 \times 0.84 = 2.09 \approx 2.13$: its price merely tracks its adoption, with no amplifier. Litecoin has survived since 2011 with no meaningful price law ($R^2 = 0.62$): survival does not create the curve. The claim that "ETH does not follow a power law" circulates publicly (e.g. Krueger 2024, on annual data); our decomposition makes it precise — Ethereum has a *weak* temporal power law and a *zero* wealth surplus, and it is the latter that matters.

Cross-chain caveats: address semantics differ (account model vs. UTXO), making cross-chain β_M comparisons rough at the ±0.1 level — an order of magnitude smaller than the surplus gap.

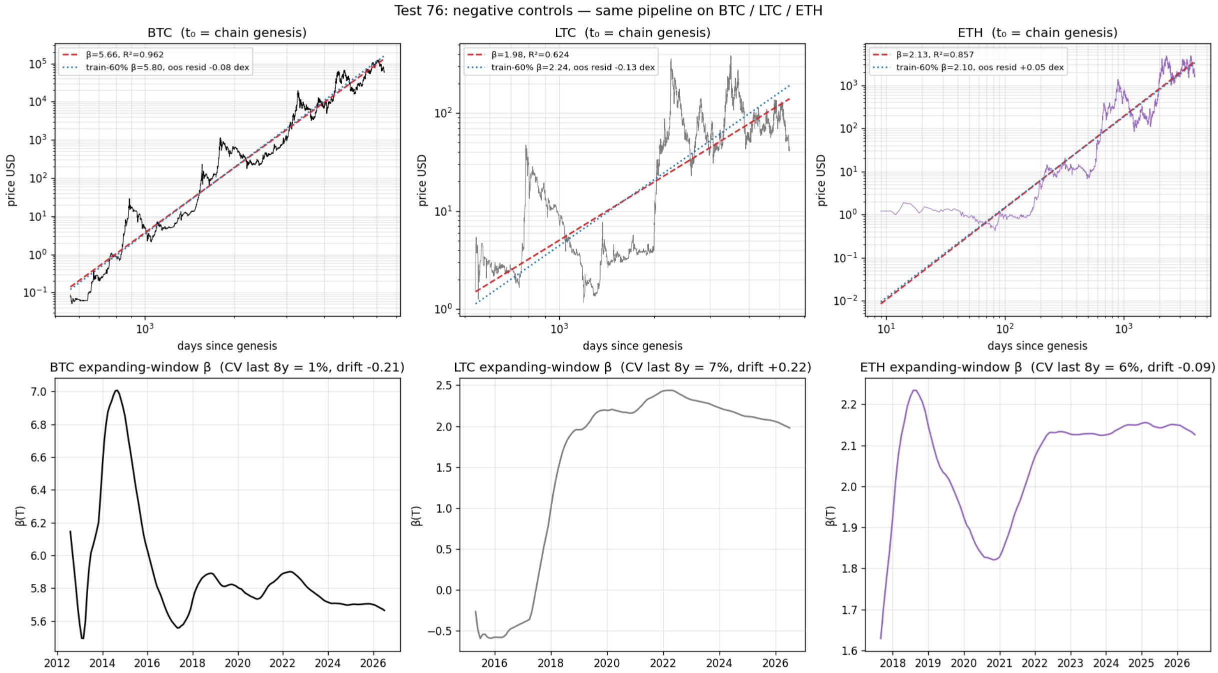

Figure 9: Same pipeline on BTC / LTC / ETH: price series with fits and forward tests (top), expanding-window β (bottom).

10 Corroborations and boundary tests

The results above carry the argument. The observations below do not — but each one had the power to embarrass the theory, and none did.

A live stress test: the 2024 ETF break. In January 2024 spot ETFs broke the adopter-counting channel: one custodial address now stands for millions of beneficial owners. Visible addresses froze near 53M, measured capital per address doubled, the balance-proxy β_A collapsed ($3.19 \to 0.55$) while the per-holder capital exponent surged ($3.27 \to 6.43$) — and the price law did not move (Fig. 10). This is not a confirmation of Eq. (5); it is evidence that the law is structural rather than an artifact of one growth channel, and the hand-over from headcount to capital-per-holder is what a wealth-climbing wave does when it moves up a layer.

A non-crypto positive control: domain names. Internet domain adoption is an independent cognitive epidemic decades older than Bitcoin. On hand-collected Verisign .com/.net registration anchors, the growth-phase adopter base (1995–2010) grows as $t^{3.0}$ (Test 78) — $\beta_A \approx 3$ on a completely different network, confirming that the epidemic factor is nothing Bitcoin-specific. The *price* side we leave open: domain sales are a unique-item rarity market with no fungible price series, so β_M is not measurable to our standard. (We initially read the post-2013 .com plateau as Bass saturation; that reading is confounded — ICANN's new-gTLD program launched exactly then, and demand migrated to 1,200+ new endings while total registrations kept growing, so the plateau is part maturation, part substitution. Bitcoin's analogue, the altcoins, do *not* absorb the monetary premium (Table 2), so its premium does not fragment the way domain demand did.) Remaining tests: domain aftermarket prices with a quality-adjusted index, and collectibles (no network value: should show no surplus).

The boundary of the class: mobile phones. If "everything grows as t^3", the epidemic factor would be unfalsifiable. It is not: world mobile subscriptions (ITU/World Bank, 1980–2025) —

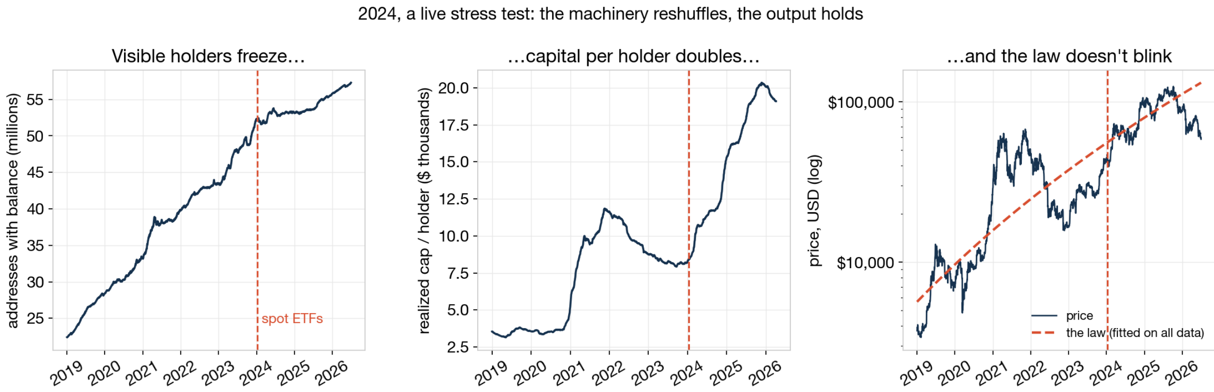

Figure 10: The 2024 ETF measurement shock: visible holders freeze, capital per holder doubles, the price law is unmoved.

the largest technology adoption on record, but with a near-zero cognitive gate and a supply-side push — grew *exponentially* through their growth phase ($\times 1.6$/yr, $R^2 = 0.977$, vs. 0.924 for the best power law, whose exponent is also unstable under the time origin, spanning 1.4–3.8), then bent into Bass saturation (Test 79). Polynomial t^3 is the signature of *decision-gated* spreading, not of technology diffusion in general — the adoption-side mirror of the airborne-vs.-HIV contrast of Sec. 2. Smartphones, the same contagion's second wave, also carry no price power law: with elastic supply there is nothing for a wealth surplus to act on.

Geography of susceptibility. A complementary check on the inverse-ordering mechanism (Sec. 6): at the country level, adoption follows two channels (Tests 59–60). Institutional adoption tracks national innovation capacity (ETF-country median innovation rank 13/133 vs. 72/133 for non-adopters, $p < 10^{-5}$), while individual ownership is also driven by monetary necessity (Turkey, Argentina, UAE over-adopt). Tellingly, stable-currency countries *under*-adopt — Switzerland sits $\sim$11 points below its innovation-predicted ownership rate — consistent with weak money lowering the cognitive threshold, i.e. the necessity channel of the susceptibility axis.

11 Predictions and domain of validity

The theory is conditional and monitorable. It holds while (i) the epidemic keeps climbing wealth tiers and (ii) wealth remains Pareto-shaped with $\alpha \approx 1.2$. Falsifiers, all observable:

- β_M leaves $[1.65, 2.0]$ (it has sat at 1.83 ± 0.09, clock-free, for sixteen years);
- the capital-per-holder schedule flattens (its $\times$2,500 power-law climb ends);
- Ethereum develops a surplus without a supply cap (scarcity not necessary), or a capped, wealth-climbing asset shows none (ingredients not sufficient — watch domain names);
- global wealth α exits $[1.0, 1.5]$ without β_M responding.

If inequality deepens ($\alpha \downarrow$), the exponent steepens; if it eases, $\beta \to \beta_A$. Saturation arrives layer by layer: at the measured pace the individual layer ($\sim$7% of world population holding) would not approach half-penetration before the late 2030s, and the corporate, institutional and sovereign layers are at first movers ($\sim$0.4%, $\sim$0.05%, $\sim$1.5% respectively); sixteen years of adoption data show no logistic bending yet. The model predicts a gradual hand-over to mature-asset growth, visible years in advance in the adoption data — not a cliff. The law says nothing about cycles: the 2011/2013/2017/2021 oscillations are residuals around it (MVRV-type reflexivity is measured trend-neutral), which is why weather-grade prediction remains impossible while climate-grade prediction has worked for a decade.

12 Limitations

(i) $\beta_A \approx 3$ is *measured* (three ways) and theoretically *motivated*, not derived: the polynomial-spreading class spans exponents $\sim$1.8–4.2 across networks, and why Bitcoin sits at 3 remains open — as does, therefore, first-principles derivation of β itself. Our contribution is the second factor. (ii) The $\sim$2% residual between $\beta_A\beta_M$ and β is regression inconsistency; second-order mechanisms (Lindy credibility, allocation deepening, leverage) are identified but deliberately unfitted. (iii) All adoption proxies degrade (consolidation, ETFs); the schedule measurement inherits realized-cap biases. A precision test — the wealth of dated adopter cohorts against $N^{1/\alpha}$ — awaits entity-level curation. (iv) Geometric tail penetration (H) is the theory's central postulate, not a derived consequence of the epidemic dynamics. The ordering simulation shows it is *necessary* — rank correlation between wealth and adoption time, even perfect, gives $\beta_M \approx 1$ (Sec. 6) — but nothing in standard epidemic models forces this schedule; to date it is supported by the adoption timeline and the consistency-grade proxy of Sec. 8, and the entity-cohort test of (iii) is what could establish it — or kill the theory. (v) The formula's independence of β_A and β_M is argued and simulated, not proven as a theorem on realistic networks. (vi) Personality-cohort evidence often cited for the cognitive filter (MBTI surveys) is illustrative only; none of our claims rest on it.

13 Conclusion

Santostasi & Perrenod measured Bitcoin's law and split its exponent; one factor had a theory, the other a label. We supplied the missing theory: the value-scaling exponent is $1 + 1/\alpha$ — one plus the inverse of the world's wealth-inequality exponent. The mechanism is a geometric climb of an idea-epidemic through the wealth ranks of a capped asset; it is simulated (rank ordering fails, the geometric schedule works), observed on-chain (capital per holder as a power of N), and cross-validated by control assets that get the adoption law but not the surplus, exactly as required. With honest errors the prediction lands within a tenth of a standard deviation.

One sentence: *Bitcoin's price exponent is the speed of an idea-epidemic multiplied by the inequality of human wealth — $\beta = \beta_A(1 + 1/\alpha)$ — and the assets missing an ingredient are measured to lack the surplus.*

A Test registry (selected)

The full suite comprises 79 numbered empirical tests and 20 simulations (numbered to sim22, with gaps); the repository carries one results file per test. Pivotal for this paper: Tests 32/74 (forward test; honest errors), 75 (t_0), 36/24/64 (β_A three ways), 21/35/73 (component mechanics), 55 (formula), sim22 (ordering), 77 (schedule), 53 (ETF break), 76 (controls), 78 (domain-name positive control), 79 (mobile phones: a weakly-gated contagion grows exponentially, not as t^3 — the boundary of the Colgate class), 59–60 (geography of susceptibility), 62 (on-chain wealth Pareto $\alpha = 1.07$), 18 (stability of the temporal law vs. multivariate alternatives), 1–48 (the monetary detour, Appendix D).

Two consistency results support the framework without being load bearing. Test 72 compiles the exponent "zoo": a dozen on-chain and attention metrics each follow a power law, and their exponents satisfy the algebraic relations the framework implies (hashrate $\approx 2\times$ price, from mining economics; volume $\approx$ price $+\beta_A/2$; addresses and Wikipedia both $\approx \beta_A$) — the law is a system of mutually consistent exponents, not one lucky fit. Test 70 measures the lead–lag structure of the reflexive loop on monthly returns: narrative (Wikipedia views) leads price ($r = 0.23$, $p = 0.004$), price leads adoption and hashrate, and volume does *not* lead price — the temporal ordering the cognitive-epidemic story assumes, observed directly.

B Pivotal tests: data and procedure

For each load-bearing or frequently cited result, the inputs and the logical steps — enough to reproduce independently, without reading the scripts. All fits are OLS in $\log_{10}$–$\log_{10}$ unless stated; "honest errors" means Newey–West (1–4 yr bandwidths) and moving-block bootstrap (0.5–2 yr blocks, $B = 2{,}000$), quoting the most conservative.

Honest error bars (Test 74). *Data:* daily BTC/USD (blockchain.info, 2010-08 to 2026-06); `AdrBalCnt` (CoinMetrics). *Procedure:* (1) fit $\log P$ on $\log t$ and $\log P$ on $\log N$; (2) measure residual AR(1) autocorrelation ($\rho = 0.998$), giving $N_{\text{eff}} = N(1-\rho)/(1+\rho) \approx 6$; (3) recompute errors by HAC and block bootstrap; (4) score candidate β_M values by z-distance. *Output:* $\sigma(\beta) = 0.21$, $\sigma(\beta_M) = 0.091$ ($\times 14$ naive); Zipf $z = -3.6$, Metcalfe $z = +1.9$.

Time-origin scan (Test 75). *Data:* as above. *Procedure:* (1) refit β with t_0 swept from 2008 to 2010 (whitepaper, genesis, first traded price, Mt.Gox); (2) record $R^2(t_0)$; (3) repeat for β_A; (4) compare β/β_A and the clock-free direct regression β_M. *Output:* β spans 4.8–5.8 with $\Delta R^2 = 0.0014$ (the data cannot pin t_0); β/β_A moves $\sim$1%; β_M is exactly t_0-free.

β_A three ways (Tests 36/24/64). *Data:* `AdrBalCnt`; Google Trends worldwide "bitcoin" (cumulative); Wikipedia article views (cumulative, Newhedge). *Procedure:* fit each observable against genesis-anchored time; none shares a data source with the others. *Output:* 3.02 ($R^2 = 0.97$), ≈ 3.6 (0.99), 3.07 (0.91).

Window mechanics (Test 73). *Data:* as Test 74. *Procedure:* (1) compute expanding-window paths of β_A, β_M, and implied α; (2) reproduce the "convergence" path with pure algebra under constant β; (3) build a time-resolved wealth-sweep counterfactual and compare directions. *Output:* $\text{corr}(\beta_M, \beta/\beta_A) = 0.98$ across windows; the convergence path is mechanical (basis for retracting our former signature plot, FL-22).

Ordering simulation (sim22). *Setup:* $M = 10^7$ agents, Pareto(α) wealth, capital-sum observable $V(N)$. *Procedure:* (1) adopt in orderings from a Gaussian copula of rank correlation $\rho \in [0, 1]$; (2) adopt on the geometric schedule (rank $\simeq M/k$) with lognormal rank noise of 0–2 decades; (3) fit β_M at Bitcoin-like penetration ($\leq 1\%$). *Output:* every copula ordering gives $\beta_M = 1.00$–1.08; the geometric schedule gives $1 + 1/\alpha$, robust to $\sim$1 decade of noise.

Penetration schedule (Test 77). *Data:* realized capitalisation (Newhedge, to 2026-04) over `AdrBalCnt`. *Procedure:* (1) form capital-per-holder $\bar{w}(N) = \text{RC}/N$; (2) fit $\bar{w} \propto N^s$ per window; (3) compare with $s = 1/\alpha$ (geometric) vs. $s = 0$ (order-free). *Output:* $s = 0.92$ (2010–17), 1.03 (pre-ETF); $\times$2,500 total — order-free models off by three orders of magnitude. Consistency-grade (realized cap tracks price).

Cross-asset controls (Test 76). *Data:* `PriceUSD` and `AdrBalCnt` for BTC/ETH/LTC (CoinMetrics), each chain's own genesis. *Procedure:* run the identical pipeline (same estimator, same proxy definitions) on the three chains; decompose $\beta = \beta_A \times \beta_M$; form the surplus $\beta_M - 1$. *Output:* Table 2 — adoption power laws generic; surplus +0.84 (BTC), −0.16 (ETH), −0.08 (LTC).

Domain names (Test 78). *Data:* hand-collected Verisign .com/.net registration anchors (DNIB/IR/SEC), 1985–2026; t_0 = registry launch (1985-03). *Procedure:* growth-phase vs. late-phase fits; t_0 scan as in Test 75. *Output:* growth phase (1995–2010) $t^{3.0}$ with $R^2 = 1.000$; post-2013 plateau confounded by the new-gTLD substitution.

Mobile phones (Test 79). *Data:* World Bank/ITU `IT.CEL.SETS`, world aggregate, annual 1980–2025. *Procedure:* (1) growth-phase (1980–2005) fits in both families — power law vs. exponential (semi-log); (2) t_0 scan over first-network candidates (1979/81/83/85); (3) registered prediction: a weakly-gated, supply-pushed contagion should *not* land in the t^3 class. *Output:* exponential wins ($R^2 = 0.977$ vs. 0.924); no stable temporal exponent (1.4–3.8 across t_0); prediction confirmed.

C False leads

Twenty-four documented dead ends, of which the reader deserves four. **FL-2**: a sign error in an early test "refuted" the Pareto model our first paper had proposed, triggering a 22-test detour that ended by re-deriving it. **FL-22**: our former signature plot — an expanding-window "convergence" of implied α interpreted as the epidemic sampling the wealth distribution — is mechanical (Fig. 5) and was retracted by our own adversarial review; only the full-period endpoint is evidence. **FL-23**: an earlier simulation cited as verifying $1 + 1/\alpha$ responds too weakly to α (sensitivity $\sim$0.3 instead of 1) and used precisely the sequential ordering that produces no surplus; it is superseded by the simulation of Sec. 6.2. **FL-24**: "the poor adopt before the rich" is not the mechanism; geometric tail penetration is.

D The monetary detour and its cancellation

Roughly a third of our test suite (Tests 1–48) explored monetary deflators: US M2, global M2 built from 35 national series, credit aggregates, Cantillon splits (US ≈ 68/32 asset/CPI; global ≈ 50/50), and forex adjustments. (One construction detail worth disclosing: the Chinese M2 series is spliced from two PBoC sources at a single 2019-08 junction with one scale factor chosen for continuity; it is the largest hand-set step in that pipeline and propagates into every global-M2 test.) The physics is real — deflating by monetary aggregates removes measurable exponent — but the Pareto tail exponent is invariant under uniform rescaling of wealth, so monetary expansion cancels from the *nominal* $\beta_M = 1+1/\alpha$: the formula predicts the nominal exponent directly, with no M2 term. Non-US monetary debasement acts on adoption (capital flight lowers the cognitive threshold), i.e. inside β_A, not as a separate factor.

E Reproducibility

Price: blockchain.info; on-chain: CoinMetrics community API (`AdrBalCnt`, `PriceUSD` for btc/eth/ltc); realized cap, Wikipedia views, Google Trends: Newhedge. All figures in this paper regenerate from cached raw data with one script each; the repository also contains the 79 numbered test scripts, 20 simulations, the epistemological map (claim-by-claim evidence typing), and the false leads register. Headline numbers quoted in the text are frozen as of the 2026-06-30 data pull; the test scripts refetch live sources, so later runs will drift at the percent level — the frozen values, not fresh runs, are the numbers this paper defends. This work was carried out in extended human–AI collaboration; the methodology note of v1.0 applies [2].

References

[1] G. Santostasi, S. Perrenod, *A Mechanistic Derivation of the Bitcoin Price Power Law: Network Adoption Dynamics and Generalised Metcalfe Scaling*, Nonlinear Science (2026) 100172. https://doi.org/10.1016/j.nls.2026.100172. Preprint: https://zenodo.org/records/19387099

[2] wineLightning, Claude AI, *Bitcoin's Price Power Law Decomposed: Epidemic Spreading, Wealth Inequality, and Money Printing*, v1.0, Zenodo (2026). https://doi.org/10.5281/zenodo.19411140

[3] 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*, PNAS 86 (1989) 4793–4797.

[4] B. Szendrői, G. Csányi, *Polynomial epidemics and clustering in contact networks*, Proc. R. Soc. Lond. B 271 (2004) S364–S366.

[5] G. Chowell, L. Sattenspiel, S. Bansal, C. Viboud, *Mathematical models to characterize early epidemic growth: A review*, Physics of Life Reviews 18 (2016) 66–97.

[6] P. Williams, Z. Chen, *Explosive Growth in Large-Scale Collaboration Networks*, arXiv:2502.11109 (2025).

[7] B. F. Maier, D. Brockmann, *Effective containment explains subexponential growth in recent confirmed COVID-19 cases in China*, Science 368 (2020) 742–746.

[8] C. Menger, *On the origin of money*, Economic Journal 2 (1892) 239–255.

[9] S. Ammous, *The Bitcoin Standard: The Decentralized Alternative to Central Banking*, Wiley (2018).

[10] V. Pareto, *Cours d'économie politique*, Lausanne (1896).

[11] J. B. Davies, S. Sandström, A. Shorrocks, E. N. Wolff, *The level and distribution of global household wealth*, Economic Journal 121 (2011) 223–254.

[12] T. Piketty, *Capital in the Twenty-First Century*, Harvard University Press (2014).

[13] Credit Suisse / UBS, *Global Wealth Report* (2023).

[14] S. Wheatley, D. Sornette, T. Huber, M. Reppen, R. N. Gantner, *Are Bitcoin bubbles predictable? Combining a generalized Metcalfe's law and the LPPLS model*, Royal Society Open Science 6 (2019) 180538.

[15] PlanB, *Modeling Bitcoin value with scarcity* (2019).

[16] C. Baquero, R. Menezes, *Bitcoin's Power Law: Weak Structure, Strong Forecasts*, arXiv:2605.21316 (2026).

[17] I. Eliazar, *Lindy's law*, Physica A 486 (2017) 797–805.

[18] E. M. Rogers, *Diffusion of Innovations*, 5th ed., Free Press (2003).

[19] M. Granovetter, *Threshold models of collective behavior*, American Journal of Sociology 83 (1978) 1420–1443.

[20] D. Kondor, M. Pósfai, I. Csabai, G. Vattay, *Do the rich get richer? An empirical analysis of the Bitcoin transaction network*, PLoS ONE 9 (2014) e86197.

[21] W. K. Newey, K. D. West, *A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix*, Econometrica 55 (1987) 703–708.

[22] H. R. Künsch, *The jackknife and the bootstrap for general stationary observations*, Annals of Statistics 17 (1989) 1217–1241.

[23] F. Krueger, public analysis of ETH price scaling, X (April 2024).