THE ROGUE PROTOCOL
Forensic Financial Intelligence

THE DETERMINISM PROBLEM

A Forensic Examination of Bitcoin Power Law Modelling:From Retail Monte Carlo to Academic Physics

Paul Faulkner
The Rogue Protocol
April 2026

EXECUTIVE SUMMARY

What the Power Law Claims

Bitcoin's price, plotted on logarithmic axes against time since the Genesis Block, traces a remarkably straight line. The slope of that line – the power law exponent, approximately 5.69 – has remained stable across 15 years and four market cycles. Proponents claim this is not coincidence. It is, they argue, a deterministic consequence of Bitcoin's network structure: adoption follows epidemic spreading dynamics on a scale-free network, producing cubic growth in address count; price scales with address count via generalised Metcalfe's Law; the product of the two exponents equals the observed 5.69. The power law floor derived from this relationship is presented as a structural support level, a physical constant of Bitcoin's dynamics, and the basis for long-run price projections extending to seven figures.

The Structural Faults

This paper identifies nine structural faults across two analyses. The retail variant fails on five grounds: the regression floor's residual stationarity establishes historical consistency but does not prove the floor is a structural constraint – the S&P 500 exponential trend passes the identical test; the Monte Carlo boundary condition is circular – the floor cannot be breached because the simulation was built to prevent breaches; the volatility decay narrative is a log-scale coordinate artifact; the halving-cycle regime segmentation is entirely in-sample; and no generative mechanism is provided for why the relationship should persist.

The academic variant, Santostasi and Perrenod (2026), fails on four additional grounds: the composition identity beta = beta_A x beta_M is algebraic tautology, not independent evidence; the epidemic spreading derivation reverses the logical order, fitting the theory to the measured exponent rather than predicting it; the Bayesian stability analysis uses the OLS estimate as its own prior and treats 1,899 autocorrelated window estimates as independent draws, artificially compressing the posterior by a factor of approximately five; and the paper omits a Granger causality test between price and addresses, leaving open whether the Metcalfe coefficient measures adoption driving price or price driving adoption.

The Central Conclusion

Both analyses demonstrate that Bitcoin's price history is well-described by a power law. Neither demonstrates that it will continue to be. Residual stationarity – confirmed by independent ADF testing with p-value of approximately 0.0075 – establishes that the relationship has been consistent over the historical period. It does not establish that it will remain so. The S&P 500 exponential trend passes the same stationarity test with near-identical p-values. No one claims the S&P follows a physical law. The test does not discriminate between a structural constraint and a good historical description.

The Stock-to-Flow Parallel

The power law is the second major Bitcoin quantitative framework to follow the same five-stage trajectory: genuine empirical observation, promotion to physical law, construction of an unfalsifiability architecture, accumulation of social infrastructure that protects the model from critique, and eventual collapse when the market delivers an outcome the bands cannot contain. Stock-to-flow projected $288,000 by December 2021. Bitcoin reached $69,000 and fell to $15,500. The power law's falsification criteria are structured to prevent an equivalent reckoning for at least 25 years. Part III provides a diagnostic checklist that any practitioner can apply to identify when a quantitative model has crossed from analysis into advocacy.

PART I

The Retail Variant

Monte Carlo, Volatility Regimes, and the Floor That Cannot Be Crossed

---

1. The Analysis and Its Claims

A Monte Carlo simulation of Bitcoin price trajectories, constructed using log-residuals from the historical price series and segmented by halving-cycle volatility regimes, has been circulated on financial social media with the following central claims:

- Volatility decays toward the power law floor across successive halving cycles.
- Extrapolation of this decay trend predicts near-zero volatility by approximately 2050.
- This convergence constitutes evidence of "Hyperbitcoinisation", the absorption of Bitcoin into the role of global reserve asset.
- The power law floor functions as a structural support level below which price cannot sustainably trade.

The analysis is visually compelling. The simulation produces a fan of percentile bands that narrows over time, converging on the power law trend. The halving cycle segmentation adds apparent rigour. The comparison to physical phenomena – convergence, decay, attractors imports the language of deterministic systems into what is, in substance, a curve fit on a 12-year dataset.

Each claim is examined in turn below.

2. Fault One: The Regression Floor and What Stationarity Does Not Prove

The power law floor is estimated by regressing log price against log time. The statistical concern here is not that this regression is definitively spurious, independent stationarity testing of the residuals (Burger & Vijn, 2024) yields an ADF p-value of approximately 0.0075, indicating stationary residuals since roughly 2016. That result deserves acknowledgement.

The problem is what that result does and does not establish. Three points bear examination.

First, strict Engle-Granger cointegration requires both variables to be I(1). But log t is a deterministic trend, not a stochastic I(1) process. Cointegration in the formal sense is therefore inapplicable, a point the power law's own proponents concede. The ADF test on residuals is a necessary check but not a sufficient foundation for the causal claims the model makes.

Second, residual stationarity only emerged after approximately seven years of data, from 2010 to 2016-17. During this period the model was being fit to its own early observations. A relationship that requires seven years of history before its residuals stabilise is consistent with a good trend description, not with an underlying structural constraint.

> *Third, and most importantly: the S&P 500's long-run exponential price trend passes the identical ADF residual stationarity test with near-identical p-values. The Nasdaq composite does the same. No serious analyst claims either index follows a deterministic physical law. The ADF test does not discriminate between a law and a description. Passing it establishes historical consistency, not forward obligation.*
> *Proponents will invoke the Lindy Effect, the longer the floor has held, the more likely it is to continue holding. This confuses duration with ergodicity. The power law has been observed across one realisation of one asset. There is no ensemble of Bitcoin networks from which to derive a distributional guarantee. Ergodicity requires that the time-average of a single trajectory equals the ensemble-average across many independent realisations; Bitcoin offers the former and has no access to the latter. A single non-ergodic trajectory, however long, cannot establish a permanent boundary condition. Fifteen years of non-violation is consistent with a floor. It is equally consistent with a sample that has not yet encountered the regime that breaks it.*

Granger and Newbold's (1974) concern remains relevant: when non-stationary series are regressed against deterministic trends, the standard errors are unreliable and R-squared is systematically inflated regardless of genuine relationship. The 0.96 R-squared is consistent with both a genuine structural relationship and a very good trend fit on a short-history asset. It cannot distinguish between them.

The Monte Carlo simulation downstream inherits this ambiguity. The percentile bands, the volatility decay narrative, and the 2050 convergence are all conditioned on a floor whose causal status has not been established. The stationarity test tells us the fit has been good. It does not tell us why.

3. Fault Two: The Circular Boundary Condition

The Monte Carlo simulation generates price paths that cannot sustainably trade below the power law floor. This is presented as a finding. It is a construction choice.

The simulation is built as follows: log-residuals from the power law fit are sampled, regime-conditioned, and applied to the trend to generate forward paths. The power law floor is treated as an absorbing or reflecting barrier – paths that breach it are either eliminated or reflected upward. The resulting simulation cannot, by design, produce scenarios in which Bitcoin trades persistently below the floor.

The confidence intervals therefore do not represent probability assessments of where price will be. They represent the distribution of outcomes conditional on the model's own assumption that the floor holds. This is circular:

> *The model cannot demonstrate that the floor holds because it was built assuming the floor holds. The non-violation of the floor is not evidence. It is a boundary condition.*

To test whether the floor is a genuine structural constraint, the correct approach is to fit the model on cycles 1 and 2, then predict cycle 3 out-of-sample and compare predicted floor adherence to realised price behaviour. This has not been done.

The floor has been breached on a log-scale basis on multiple occasions during Bitcoin's history. Each breach was subsequently recovered. This is consistent with the floor being a long-run attractor, but it is equally consistent with it being a coincidental trend line that has not yet been broken permanently. Without out-of-sample testing, no evidence distinguishes between these interpretations.

4. Fault Three: The Log-Scale Volatility Illusion

The central visual claim, that volatility is decaying toward zero across halving cycles, is a log-scale coordinate artifact, not an observed property of the price series.

Volatility in log-price terms measures percentage moves: the standard deviation of the log return series. When this is plotted on a log-log chart against a growing trend, the same percentage volatility occupies a shrinking fraction of the visual space as the trend price grows. A 50% move from $1,000 appears the same size as a 50% move from $100,000 in proportional terms, but occupies a smaller absolute distance on the chart.

sigma_log = std(delta log P)   [percentage volatility, scale-independent]

The author observes that the residuals from the power law fit are narrowing in absolute log-scale terms as time progresses. This is expected: if the power law fit is improving over time (which it will appear to do as more data accumulates and the regression is re-fitted), the residuals will shrink. This is an artifact of cumulative R-squared growth, not evidence of genuine volatility compression.

Furthermore, the claim that volatility reaches "near zero" by 2050 requires extrapolating a trend in a quantity that has shown no structural decline in its own right. The 2022 bear market produced drawdowns of approximately 75% from cycle highs, comparable to previous cycles in proportional terms. If volatility were genuinely decaying toward zero, this would not have occurred.

5. Fault Four: Regime Dredging and In-Sample Validation

The halving-cycle volatility segmentation – fitting separate volatility parameters to each of the four post-Genesis halving regimes – introduces a model selection problem that the analysis does not acknowledge.

With four regimes, four sets of parameters are estimated from the same dataset against which the model is validated. This is entirely in-sample. The model has been given the answer and asked to explain it.

The correct test is out-of-sample: fit regime parameters to cycles 1 and 2, predict the volatility structure of cycle 3, report the root mean squared error, and compare to a naive benchmark. If the regime-segmented model does not materially outperform a simpler homoskedastic model out-of-sample, the regime structure is overfitting.

No such test is presented. The conclusion that volatility "decays" across regimes is driven by a secular trend in the data, Bitcoin becoming less volatile as it grows, that a single-regime model would capture equally well.

6. Fault Five: The Missing Mechanism

Every predictive model requires a generative mechanism: a causal account of why the pattern should continue, distinct from the observation that it has continued.

The retail power law analysis does not provide one. When pressed on mechanism, the response is typically: Bitcoin's store of value properties attract adoption, adoption drives price, price follows a power law. This is not a mechanism. It is a narrative that describes what the model assumes without explaining why the assumed relationship is durable.

A genuine mechanism would specify: what market structure gives rise to a power law in price rather than exponential growth or a sigmoid; why the exponent should remain stable as institutional adoption replaces retail adoption; what would cause the law to break; and how the law connects to the supply schedule, the derivatives market, and the broader macro environment in which Bitcoin trades.

None of these questions are addressed. The model is presented as a forward-looking tool. It has not earned that status.

7. Summary: Part I

> *The retail Monte Carlo analysis demonstrates that Bitcoin's price history is well-described by a power law trend. It does not demonstrate that Bitcoin's price must continue to follow one. The floor is a curve fit whose residual stationarity establishes historical consistency but not forward obligation. The boundary condition is circular. The volatility decay is a coordinate artifact. The regime segmentation is in-sample. No mechanism is provided. The analysis is an excellent historical description. It is not a forward-looking framework.*

PART II

The Academic Variant

Santostasi and Perrenod (2026): A Mechanistic Derivation That Isn't

---

8. The Paper and Its Claims

Santostasi and Perrenod (2026) – "A Mechanistic Derivation of the Bitcoin Price Power Law", is a 20-page paper presenting what it describes as a first-principles derivation of the Bitcoin price power law exponent. The paper is competently written, properly formatted, and carries academic apparatus: LaTeX typesetting, 17 references, and a multi-figure empirical section. Most readers will be intimidated into deference.

The central claim is that the price power law exponent beta = 5.69 is not a free parameter but the product of two independently established exponents:

beta = beta_A x beta_M = 3.046 x 1.838 = 5.60

Where beta_A governs adoption dynamics (address growth in time, derived from epidemic spreading theory), and beta_M governs network value scaling (price per address, derived from generalised Metcalfe's Law). The paper claims that three independent lines of evidence converge on this value: epidemic theory, network economics, and direct OLS regression.

The paper additionally provides a Bayesian stability analysis, three scale invariance tests, and a falsifiability section. Each of these is discussed in turn.

Six structural faults are identified. Any one of them would be serious. Together, they are fatal.

9. Fault One: What Residual Stationarity Does Not Establish

The entire empirical framework rests on OLS regression in log-log space. The authors acknowledge this choice explicitly, citing Clauset, Shalizi and Newman (2009), their own reference [4] but declining to use the methods it recommends on the grounds that they are "interested in the full time-series scaling relation rather than the tail of a static distribution."

Independent testing of the power law residuals has been conducted by Burger and Vijn (2024) – published under the name Quantodian – who report an ADF test p-value of approximately 0.0075 on the price-time regression residuals, with stationarity emerging in the data from approximately 2016 onward. This result deserves acknowledgement. The residuals are stationary by standard ADF criteria. The paper's claim of residual stationarity is correct.

The question is what that result establishes. Three points define the limits of the test.

First, the paper's regression is not a standard cointegrating relationship. Engle-Granger cointegration requires both variables to be I(1) stochastic processes. log t is a deterministic trend, not a stochastic process. Formal cointegration is therefore definitionally inapplicable – the paper cannot claim cointegration, and cannot use the language of cointegration to justify the relationship's structural status. This is a point the power law's more technically literate proponents concede.

Second, residual stationarity on a 15-year Bitcoin series emerged after approximately seven years of data and has strengthened as the dataset has grown. A property that requires half the observation window to stabilise is consistent with a progressively better trend description – cumulative R-squared improving as more data is added – not with an underlying physical constraint that existed from the start. The timing of stationarity emergence does not distinguish between these interpretations.

> *Third, the S&P 500's long-run exponential price trend passes the identical ADF residual stationarity test with near-identical p-values. So does the Nasdaq composite. The ADF test on price-versus-time regression residuals is not specific to structural relationships. It is passed by any asset whose price has trended smoothly for long enough. Nobody argues the S&P follows a deterministic law of nature. The test does not discriminate between a physical law and a very good long-run description. Passing it is a necessary condition for the power law to be a structural relationship. It is not a sufficient one.*

Granger and Newbold's (1974) original concern remains: when a price series is regressed against a deterministic time trend, the standard errors are asymptotically non-standard and R-squared is systematically inflated. The reported R-squared of 0.961 cannot be compared to standard distributional benchmarks. The authors cite the correct literature and do not apply it.

The adoption regression – N(t) ~ t^3.05 – and the Metcalfe regression – P ~ N^1.84 – are subject to the same stationarity concerns. The paper provides no stationarity tests on the residuals from either subsidiary regression. The empirical structure rests entirely on the price-time residual stationarity result, which establishes historical consistency and nothing more.

10. Fault Two: The Composition Identity is Algebraic Tautology

The paper's most rhetorically effective claim, and its most intellectually fraudulent, is that three independent lines of evidence converge on beta = 5.69. Figure 2 presents a derivation pathway in which epidemic theory, network economics, and direct regression all point to the same value. The authors describe this as remarkable.

It is not remarkable. It is mandatory.

The three "lines of evidence" are not independent. They share the same dataset. If you measure N(t) ~ t^beta_A from the data, and P ~ N^beta_M from the same data, then the regression P ~ t^beta will yield beta = beta_A x beta_M by algebraic substitution, within measurement error. This is not convergence. It is a consequence of the chain rule:

$$\text{If } N = C1 * t^{\beta_A} \text{ and } P = C2 * N^{\beta_M}$$

$$\text{Then } P = C2 * (C1 * t^{\beta_A})^{\beta_M} = C1^{\beta_M} * C2 * t^{(\beta_A * \beta_M)}$$

The product beta_A x beta_M must equal the direct estimate beta_obs to within measurement variance. There is no additional information in the composition identity. The three arrows in Figure 2 are an illustration of algebra, not evidence.

The authors test this identity in rolling windows and find R-squared of 0.92 between beta_direct and beta_A x beta_M. They present this as confirming "the composition identity as a structural constraint across all epochs." What it confirms is that their three regressions are mathematically consistent with each other, as they must be, given that they are three representations of the same data.

> *For the composition identity to constitute independent evidence, each exponent would need to be measured independently: beta_A from a study of Bitcoin adoption network topology using methods independent of price data; beta_M from a cross-sectional study of network value scaling at a single point in time; and beta from the direct price regression. The paper does none of this. All three quantities are estimated from the same 5,696 daily observations. The convergence is guaranteed, not discovered.*

11. Fault Three: Endogeneity and the Causal Arrow

The Metcalfe regression – P ~ N^1.84 – is presented as evidence that network value (price) is driven by network size (address count). The assumed causal direction is: more addresses means more network value means higher price.

The reverse causal arrow is at least equally plausible, and possibly dominant. Price rises generate media coverage. Media coverage drives new wallet creation. New wallet creation inflates the address count. Under this mechanism, the Metcalfe regression is not measuring adoption driving price. It is measuring price driving adoption – with price and addresses simultaneously determined by shared underlying factors including macro sentiment, halving events, and the broader risk appetite in financial markets.

When the causal arrow is bidirectional or reversed, the Metcalfe regression is endogenous. OLS on a simultaneous system produces coefficient estimates that are biased and inconsistent – the estimated beta_M of 1.84 absorbs both the genuine network value effect and the feedback from price to address creation. The true network value scaling exponent could be materially different.

The correct test is Granger causality between price and address count: does past price predict future address growth, controlling for past address growth? Does past address growth predict future price, controlling for past price? If both relationships are significant, the system is simultaneous and OLS is invalid for estimating either direction.

> *This test is not run anywhere in the paper. The assumption that causation runs addresses-to-price is never tested. It is asserted. If Granger causality runs primarily from price to addresses – which is consistent with the observable pattern of address growth accelerating sharply after each price rally*

> *– then beta_M is not a network value exponent. It is a reflexivity exponent. The entire compositional framework collapses, because the mechanism it requires (users driving value) is the reverse of what the data may show (value attracting users).*

12. Fault Four: The Epidemic Spreading Analogy is Post-Hoc Rationalisation

The derivation of beta_A = 3 from Colgate et al. (1989) is the paper's most impressive-looking section. The authors describe the cubic growth of cumulative AIDS cases in the United States, note that their measured beta_A = 3.046 is close to 3, and claim that this places Bitcoin's adoption dynamics in the same universality class as epidemic spreading on heterogeneous networks.

This is the logical inverse of derivation. It is pattern-matching.

Colgate's cubic growth result depended on three things: an empirically measured incubation period distribution specific to HIV; a sexual contact network whose degree distribution had been directly studied in the relevant population; and specific biological transmission parameters. The cubic exponent emerged from the convolution of these measured quantities. It was not assumed; it was derived from data about the actual network.

The authors provide none of this for Bitcoin. They do not measure the degree distribution of the Bitcoin adoption network. They do not measure the "incubation period" for Bitcoin adoption – the time between initial exposure to the concept and creation of a non-zero balance address. They do not measure transmission rates across network tiers.

What they do is observe that beta_A comes out close to 3 and note that Colgate also got 3. This is not derivation. If beta_A had come out at 2.6, the paper would have cited a different reference. If it had come out at 3.8, the authors would have argued for a modified saturation wave with different connectivity assumptions. The theoretical section is reverse-engineered from the fitted parameter.

The claim that Bitcoin's adoption network is scale-free is asserted, not demonstrated. **The claim that it exhibits the specific heterogeneity structure required for cubic growth is not even asserted, it is simply assumed by analogy.** The analogy may be correct. It may be a useful framework. But it does not constitute a derivation, and it does not make beta = 5.69 a consequence of physics rather than a historical observation.

13. Fault Five: The Bayesian Analysis is a Tautological Loop

Section 9 presents a sequential Bayesian stability analysis using 1,899 local OLS estimates of the scaling exponent. The analysis is presented as independent confirmation of the power law's stability. It is not independent.

The prior is:

beta ~ N(5.69, 1.5^2)

The prior mean of 5.69 is the OLS estimate from the same 5,696 daily observations that generate the 1,899 local estimates used as the data for the Bayesian update. The prior variance of 1.5 squared is described as "deliberately wide," but it is centered on the answer the paper has already established.

The posterior converges to 5.729 with a 95% credible interval of [5.703, 5.754]. The authors interpret the smooth n^(-1/2) shrinkage of posterior uncertainty as diagnostic evidence against structural breaks.

> *A Bayesian analysis with a prior centered on the OLS estimate from the same dataset, updated with noisy draws from that same dataset, cannot detect a structural break unless the break causes future observations to fall systematically outside the prior's support. The prior support here encompasses essentially the full plausible range of beta. The posterior shrinkage is a consequence of the Normal-Normal conjugate update, not evidence about the process. Entering the answer as the prior and receiving the answer as the posterior is not confirmation. It is circular.*

For a genuine Bayesian stability analysis, the prior should be set before any Bitcoin data is examined – perhaps using the theoretical value of 6 implied by the pure integer mechanism – and updated sequentially on data withheld from the prior estimation. The rolling window analysis in Section 8.3 partially addresses this, but still draws on the same historical dataset for all window estimates.

14. Fault Six: The Mechanism Argues Against the Bull Case

This fault is constructed entirely from the paper's own theoretical framework and is the most damaging of the five.

The paper claims that beta_A = 3 arises from the saturation wave mechanism: epidemic spreading moves from high-connectivity nodes (early adopters – technically sophisticated, ideologically motivated, densely interconnected) through successive tiers of lower connectivity as each tier saturates. The cubic growth of the cumulative count is a consequence of this progressive saturation.

The mechanism implies a clear prediction: as genuine market saturation is approached – as the lower-connectivity tiers are exhausted – the saturation wave slows. The adoption growth exponent beta_A must decline. And since beta = beta_A x beta_M, the price power law exponent must decline with it.

The paper acknowledges this in its falsifiability section (F2): if beta_A falls significantly below 3, the saturation wave mechanism has changed. What it does not acknowledge is that the saturation mechanism it has invoked must eventually cause this. Saturation is not a risk to the model. It is the model's own prediction for what happens next.

The testable form of this argument is the following: if rolling beta_A estimates computed on trailing 1,000-day windows have declined below 2.8 in the 2023-2026 period – which the saturation mechanism predicts they must as each lower-connectivity tier is progressively absorbed – then beta is declining with them. Published rolling estimates from btcpowerlaw.nl and the Quantodian dataset show beta_A oscillating between approximately 2.6 and 3.2 in recent windows, with the lower end of that range concentrated in post-2022 data. If the secular trend in rolling beta_A is downward, the model's own framework projects flatter price growth than its advocates claim, a narrowing of the confidence bands that cuts from below, not above.

> *The paper provides no analysis of recent rolling beta_A stability. It reports a global estimate across the full 2010-2026 window. The mechanism the paper invokes to explain beta_A = 3 is structurally incompatible with beta_A remaining at 3 indefinitely. This is not an external critique. It is an internal contradiction that the paper does not resolve.*

15. The Falsifiability Section: Five Conditions That Cannot Falsify

Section 10 of the paper lists five conditions under which the power law would be "expected to break down." This section deserves specific attention because it is presented as evidence of scientific rigour.

The five conditions are: floor violation (price below 3-sigma for over a year); adoption collapse (beta_A falling significantly below 3); exponent drift (beta_obs drifting outside [5.0, 7.0]); Metcalfe breakdown (Metcalfe R-squared below 0.7); rolling R-squared below 0.80 for two consecutive years.

Each condition is structured to be very difficult to satisfy in the short to medium term. The floor reprices upward continuously, so the 3-sigma threshold moves with the model. The exponent range [5.0, 7.0] encompasses essentially the full range of observed rolling estimates. The Metcalfe R-squared of 0.7 is far below the reported 0.951. The rolling R-squared threshold of 0.80 has never been breached.

> *A falsifiability criterion that cannot plausibly be triggered within the investment horizon of any practitioner is not falsifiability. It is inoculation against accountability. The conditions are designed to be untriggerable on any timescale shorter than 25 years. Genuine falsifiability would require specifying: a price level below the power law floor at which the model is abandoned for a defined period; a concrete test of out-of-sample predictive accuracy; and a statement of what evidence, within a 12-month window, would cause the authors to revise or withdraw the model. None of these are provided.*

16. What the Paper Actually Shows

Stripping away the apparatus, what Santostasi and Perrenod demonstrate is the following:

- **Bitcoin's price history from 2010 to 2026 is well-described by a power law with exponent approximately 5.69.**
- The same dataset shows address growth following a cubic trend and price scaling superlinearly with address count.

- These three observations are mutually consistent, as they must be by algebra.
- The power law residuals are stationary over the historical observation period.
- Bitcoin exhibits stronger scale invariance than NASDAQ, S&P 500, or gold when measured over the historical record.

These are legitimate empirical findings. They are not trivial. They tell us something real about Bitcoin's past.

What the paper does not show, and cannot show from this data, is:

- **That beta = 5.69 will remain the correct exponent in cycles 5, 6, or beyond.**
- That the adoption network has the structural properties the epidemic analogy requires.
- That the floor is a genuine causal constraint rather than a regression line through a trending dataset.
- That the Bayesian stability result is informative about future stability rather than historical consistency.
- That price at any given future date will respect the power law's confidence interval.

17. Summary: Part II

Santostasi and Perrenod (2026) is a well-constructed paper that demonstrates Bitcoin's price history is well-described by a power law. Its claims to derive the exponent from first principles rest on: a composition identity that is algebraic necessity, not independent evidence; an epidemic spreading analogy that matches a fitted parameter to an existing theoretical result without independently measuring the network's structural properties; a regression framework whose residual stationarity does not distinguish between a structural law and a good historical description; an untested assumption that causation runs from addresses to price rather than the reverse; and a Bayesian stability analysis that enters the estimate as its own prior. Its own mechanistic framework, if taken seriously, predicts declining price growth as adoption approaches saturation. Its falsifiability conditions are structured to be untriggerable within any practical investment horizon. The paper establishes that Bitcoin has followed a power law. It does not establish that it must continue to do so.

18. Fault Seven: The Bayesian Posterior is Artificially Tight

The Bayesian stability analysis in Section 9 of the paper reports a final posterior of beta ~ N(5.729, 0.013^2), with a 95% credible interval of [5.703, 5.754]. This is presented as strong evidence for the stability of the power law exponent. The posterior uncertainty is narrow enough that the authors describe it as a "physical constant" of Bitcoin's dynamics.

The width of this interval is an artifact of an invalid independence assumption.

The 1,899 local OLS estimates that constitute the Bayesian update data are computed in overlapping rolling windows across the same underlying price series. The authors use 300 anchor times with 25 multipliers each, combining forward and backward pairs. Adjacent anchor times are separated by a small number of days. The price observations used to estimate beta at anchor time t_0 and at anchor time t_0 + 10 days are largely identical, they share nearly all of their data.

The conjugate Normal-Normal update formula used by the authors is:

tau_n^-2 = tau_0^-2 + n * sigma^-2

This formula is correct only when the n observations are conditionally independent. When they are autocorrelated – as these overlapping window estimates necessarily are – the formula overcounts the information in the data. Each new estimate is not an independent draw from the posterior. It is a slightly shifted version of the previous estimate, sharing most of its underlying observations.

The correct treatment requires computing the effective sample size n_eff, which accounts for the autocorrelation structure of the local beta estimates. For a series with autocorrelation rho at lag 1, the effective sample size is approximately:

n_eff ≈ n * (1 − rho) / (1 + rho)

The rolling beta estimates produced from overlapping 1,000-day windows stepping by 50 days have an implied autocorrelation at lag 1 of approximately 0.95 or higher, the window moves by 50 days while retaining 950 days of shared observations. Under this autocorrelation structure, the effective sample size from 1,899 estimates is not 1,899. It is in the range of 40 to 100 effective observations.

Replacing n = 1,899 with n_eff = 60 in the conjugate update formula increases the posterior standard deviation by a factor of approximately sqrt(1899/60) ≈ 5.6. The reported 95% CI of ±0.051 becomes approximately ±0.29 – a range that encompasses essentially the full distribution of rolling beta estimates observed over the historical period and provides no meaningful constraint on future values.

> *The posterior CI of [5.703, 5.754] is not a physical measurement of Bitcoin's scaling exponent. It is the output of a conjugate update formula applied to autocorrelated data as though it were*

independent. The true uncertainty, properly computed, is approximately five times wider. The reported precision is an artifact of the estimation method, not a property of the data.

19. Fault Eight: No Out-of-Sample Validation

The paper provides extensive in-sample evidence that the power law fits the historical data. It provides no out-of-sample evidence that the power law predicts future data.

This is not a minor omission. For a model claiming to describe a deterministic physical process – one whose exponent is a constant of Bitcoin's dynamics – out-of-sample predictive accuracy is the primary test. A physical law that only describes the past is not a physical law.

The correct procedure is as follows. Fit the three scaling relations – adoption growth, Metcalfe scaling, and the direct price power law – using only data up to a chosen cut-off date. Then project forward from that date using the fitted model and its uncertainty bands. Compare the realised price trajectory to the model's prediction intervals. Report the fraction of realised prices falling within the 50th, 75th, and 95th percentile bands. Repeat for multiple cut-off dates.

If this test is run with a cut-off of January 2017 – roughly the midpoint of the observation window – the forward projection must navigate the 2017 bull market peak (~$19,000), the 2018-2019 bear market (~$3,200), the 2020-2021 bull run (~$69,000), and the 2022 bear market (~$15,500), before reaching the 2024-2026 cycle. A model fit to 2010-2016 data would project a smooth trend with ±0.30 dex bands. The 2022 bear market low of approximately $15,500 sits at roughly the 10th percentile of a model projected from the 2021 high. Whether this constitutes a model success or failure depends on which date the assessment is made.

The point is not that the model would necessarily fail. It might pass a well-specified out-of-sample test. The point is that the authors do not run one. Every statistical test in the paper is in-sample: the residual stationarity test, the rolling window composition identity, the Bayesian stability analysis, and the scale invariance tests all use the full 2010-2026 dataset.

A model validated only on the data used to fit it has not been validated. It has been described. The paper describes Bitcoin's price history with considerable precision. It does not demonstrate predictive validity. Until a properly specified out-of-sample test is run and reported, the claim that the power law is a forward-looking physical law is not supported by the evidence the paper presents.

20. Fault Nine: The Address Proxy Problem

The epidemic spreading mechanism requires a count of distinct economic actors in the Bitcoin network, individuals who have adopted Bitcoin as a store of value or medium of exchange. The paper uses non-zero balance addresses as its proxy for this count.

This proxy is structurally biased in multiple directions simultaneously, and the direction of bias is not neutral with respect to the paper's claims.

Exchange Consolidation and the ETF Effect

Large custodial exchanges – Coinbase, Binance, Kraken, and their equivalents – hold the Bitcoin of millions of customers in a small number of cold storage addresses. A single exchange omnibus address may represent hundreds of thousands of economic actors. The paper acknowledges this bias but does not quantify it.

The ETF approval in January 2024 sharpened the problem decisively. BlackRock's iShares Bitcoin Trust (IBIT), held approximately 570,000 BTC by early 2026, representing the investment exposure of over 500,000 individual brokerage account holders, all behind a structure that maps to a handful of on-chain addresses. Fidelity's Wise Origin Bitcoin Fund adds another 200,000-plus BTC held on behalf of a similarly large retail base. These products accelerate precisely in the period the paper treats as confirmatory of beta_A's stability – and they drive a structural wedge between address count and economic actor count that grows with each institutional inflow.

Multi-Address Users

The inverse bias also exists. Privacy-conscious users, hardware wallet holders, and technically sophisticated participants routinely operate dozens or hundreds of addresses per economic actor. Early adopters – the high-connectivity core of the epidemic spreading model – are disproportionately represented in this category. The address count overstates the user count for precisely the network cohort whose behaviour the model uses to anchor the cubic growth exponent.

The combined effect is that the address growth series N(t) does not measure what the epidemic spreading model requires it to measure. Exchange consolidation drives the proxy downward relative to true users in the recent period. Multi-address behaviour drives it upward in the early period. The exponent fitted to this series is not the exponent of user adoption across the full history, it is the exponent of address creation, which is a different and considerably messier quantity.

> *If the true user count has grown more slowly than addresses in recent years – as ETF-driven custodial consolidation directly implies – then the true beta_A is lower than 3.046 in the current cycle, and the true beta is lower than 5.69. The paper's own authors note that a direct measurement of the adoption network degree distribution is required to confirm the epidemic spreading interpretation rigorously. Given the structural changes introduced by spot ETF custody, such a measurement would yield a materially different result today than it would have in 2016.*

PART III

The Architecture of Unfalsifiability

Recognising When Empirical Observation Becomes Unfalsifiable Doctrine

---

21. A Recurring Pattern

The power law is not the first model to follow this trajectory. Bitcoin's analytical history is littered with frameworks that moved from empirical observation to physical law to investment thesis to unfalsifiable doctrine in the space of a market cycle. Understanding the pattern is more useful than refuting any individual instance of it.

The pattern has five stages.

Stage One: The Observation
A genuine empirical regularity is identified in Bitcoin's historical data. The regularity is real. The observation is legitimate. Bitcoin's price has followed a log-linear trend with remarkable persistence. The stock-to-flow ratio has correlated with price over certain periods. Halving events have preceded bull markets. These observations are not inventions.

Stage Two: The Promotion to Law
The empirical regularity is reframed as a deterministic law. The word "follows" becomes "obeys." The historical correlation becomes a causal mechanism. The curve fit acquires a physical interpretation, supply scarcity, network topology, epidemic dynamics. The fitted parameter is presented not as an estimate with uncertainty but as a constant of nature.

The transition from Stage One to Stage Two typically involves the importation of scientific vocabulary and citation architecture. Papers from epidemiology, statistical physics, or information theory are cited to provide theoretical grounding. The citations are real. The grounding is not, the imported theory is applied by analogy rather than by measurement, and the analogy is not tested.

Stage Three: The Unfalsifiability Architecture
The model is structured so that it cannot be falsified on any timescale that matters to practitioners. This is achieved through several techniques used individually or in combination:

- Wide uncertainty bands. The ±0.30 dex (factor of 2) bands of the power law model mean that any price between approximately $45,000 and $180,000 at current levels is within the model's prediction interval. A model whose bands span a factor of four cannot be wrong about direction.

- A repricing floor. The power law floor is not a fixed level. It grows with time. A price that fell below the 2022 floor is above the 2020 floor. The floor cannot be permanently breached because it retreats beneath any sustained breach.
- Distant falsification criteria. The conditions under which the model is declared falsified are set sufficiently far in the future – or require conditions sufficiently extreme – that no practitioner will live to see them triggered.
- Outcome attribution asymmetry. When price rises toward the upper band, this confirms the model. When price falls toward the lower band, this is a "buying opportunity", also confirming the model. No outcome disconfirms it.

Stage Four: The Social Infrastructure

A community forms around the model. The model's predictions become identity-forming beliefs. Dissent is reframed as intellectual failure, the sceptic "doesn't understand the model," or "is still thinking in traditional finance terms," or "hasn't done the work." The barrier to entry is raised by the model's technical apparatus: OLS regressions, Bayesian posteriors, epidemic spreading theory. Most observers cannot evaluate the mathematics and defer to those who appear to be able to.

This social infrastructure is not an accident. It is a feature. A model that requires specialist knowledge to critique is a model that is protected from critique. The academic apparatus serves a gatekeeping function as much as an analytical one.

Stage Five: The Collapse

The model eventually fails, not because it is formally falsified, but because the market delivers an outcome sufficiently extreme that the unfalsifiability architecture cannot contain it. Stock-to-flow projected a price of $288,000 by December 2021. Bitcoin reached $69,000 and then fell to $15,500. The model was not officially falsified, its creator claimed the bands were wider than commonly understood, that the methodology remained valid, that critics had misread the framework. But the practical credibility of the model as a forward-looking tool was destroyed.

The power law is currently at Stage Four. It has accumulated a substantial following. Its critics are told they don't understand the physics. Its uncertainty bands are wide enough to accommodate most near-term outcomes. Its falsification criteria are set for 2025 conditions that have not been triggered and for 2050 conditions that are untestable.

22. The Stock-to-Flow Parallel

The parallel with Plan B's stock-to-flow model is instructive and should be explicitly understood by anyone evaluating the power law framework.

Stock-to-flow regressed Bitcoin's price against the ratio of existing supply to annual new supply. The R-squared was high. The fitted relationship was presented as a law governing Bitcoin's value. The model was deployed to project prices of $100,000, $288,000, and beyond. It attracted enormous social media following, was cited in mainstream financial media, and was used by retail investors to justify allocation decisions.

The methodological critique of stock-to-flow maps almost exactly onto the critique of the power law:

- **OLS regression of an I(1) price series against a deterministic trend variable, without cointegration testing.**
- High R-squared interpreted as evidence of a causal relationship rather than a statistical artifact.
- A theoretical mechanism – supply scarcity drives value – asserted rather than derived.
- Uncertainty bands wide enough to accommodate a very large range of outcomes.
- Falsification criteria that retreated as each projected price level was missed.

The difference between stock-to-flow and the power law is one of sophistication, not structure. The power law framework borrows from epidemic theory and network physics rather than commodity valuation. The mathematical apparatus is more elaborate. The citations are more credible. The authors are more careful about their language.

The structure is identical. Both models take a genuine empirical observation – supply scarcity correlates with price over certain periods; price has followed a power law trend – and promote it to a deterministic law. Both construct unfalsifiability architectures. Both attract social infrastructure that protects them from critique.

> *The sophistication of the apparatus is not evidence of the validity of the model. It is evidence of the sophistication of the people who built it. These are not the same thing.*

23. A Diagnostic Checklist

The following questions, applied to any Bitcoin quantitative model, will identify whether it belongs to this category:

- **Does the model claim to derive a future price path from historical data alone, without reference to macro conditions, regulatory environment, competitive landscape, or market structure?**
- Are the uncertainty bands wide enough that the model cannot be wrong about direction over a 12-month period?
- Is the floor or support level defined relative to the model itself, such that it reprices upward if price rises?
- Are the falsification criteria set at price levels or time horizons that are practically untestable for a retail investor?

- When challenged, does the model's advocate claim that critics have not understood the model, rather than engaging with the specific methodological objection?
- Does the model produce a specific price target that is dramatically higher than current levels, a number sufficiently large to be motivating but sufficiently distant to be unverifiable in the short term?
- Has the model been validated out-of-sample, on data withheld from the fitting process, with pre-specified criteria for success and failure?

A model that fails any two of these tests should be treated as a narrative, not an analytical tool. A model that fails four or more should be treated as advocacy dressed as analysis.

To be clear: the checklist is not designed to fail everything. A well-constructed GARCH volatility model, for example, passes five of the seven criteria cleanly: it makes no claim to derive price paths from historical data alone (it models conditional volatility, not levels); its uncertainty bands are calibrated to realised volatility, not set to cover all outcomes; it has no repricing floor; its falsification criteria are testable on short horizons (does realised volatility fall within forecast intervals?); and it is routinely validated on withheld data with reported coverage statistics. Similarly, a cointegrated pairs trading strategy – fit on a training period, applied to a holdout period, with pre-specified entry and exit rules – passes the out-of-sample test unambiguously. The checklist identifies a specific failure mode. It does not indict quantitative finance.

The retail Monte Carlo analysis examined in Part I fails five of seven. The Santostasi-Perrenod paper examined in Part II fails six. Both are well-intentioned. Both may prove to be correct in their broad directional claims. Neither is doing what it claims to be doing.

FINAL CONCLUSION

Bitcoin may reach one million dollars. It may not. The power law may continue to describe its price trajectory for another decade. The epidemic spreading mechanism may prove to be a genuine structural feature of the adoption network. These things are possible.

What is not possible is that any of them are established by the analyses examined in this paper. The retail Monte Carlo and the Santostasi-Perrenod derivation both demonstrate that Bitcoin's price history is consistent with a power law. That is the only claim the evidence supports.

Consistency with a historical pattern is not a prediction. An R-squared of 0.96 on a non-stationary regression is not science. A composition identity is not independent evidence. A Bayesian posterior is not informative when it is built from its own prior. A falsifiability criterion that cannot be triggered before 2050 is not falsifiability.

The power law is a very good description of where Bitcoin has been. **It is not a constraint on where it is going.** Anyone telling you otherwise is either confused about the difference between description and prediction, or is hoping that you will be.

REFERENCES

[1] Santostasi, G. and Perrenod, S. (2026). A Mechanistic Derivation of the Bitcoin Price Power Law: Network Adoption Dynamics and Generalised Metcalfe Scaling. Zenodo.

[2] Burger, H.C. and Vijn, P. (2024). Bitcoin's time-based power-law and cointegration revisited. Quantodian Publications, Medium.

[3] Granger, C.W.J. and Newbold, P. (1974). Spurious regressions in econometrics. Journal of Econometrics, 2(2), 111–120.

[4] Engle, R.F. and Granger, C.W.J. (1987). Co-integration and error correction: representation, estimation, and testing. Econometrica, 55(2), 251–276.

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

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

[7] MacKinnon, J.G. (1991, updated 2010). Critical values for cointegration tests. Queen's Economics Department Working Paper No. 1227.

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

Paul Faulkner
The Rogue Protocol
April 2026
paulfaulkner.com