

Asymmetric Tail Curvature in Bitcoin Price Quantiles

Benjamin Cowen

Independent Research

Correspondence: benjamincowen.com

Working Paper

This version: May 29, 2026

Abstract

We develop a rearranged asymmetric quadratic quantile regression framework for long-duration price processes, in which lower and upper distributional tails may exhibit different curvature in log-log space. The specification imposes a shared curvature within each tail group, enforces non-crossing via the rearrangement estimator of Chernozhukov, Fernández-Val, and Galichon (2010), and is motivated by distinct mechanisms plausibly governing structural support and speculative excess. We apply the framework to Bitcoin, whose price history is arguably the clearest available digital-asset record long enough to identify second-order distributional structure. The estimated quantile bands describe the conditional distribution of price level given time, not return tail risk or portfolio loss probabilities. Three prior Bitcoin models (an OLS power law, stock-to-flow, and its cross-asset extension S2FX) exhibit systematic out-of-sample optimistic bias (geometric mean price errors +32.1%, +294.5%, and +1,699%) despite strong in-sample fit. Using 5,788 daily observations from 2010 through 2026, we find an asymmetry $\Delta b = b^{\text{HI}} - b^{\text{LO}} = -0.302$ between upper-tail curvature ($b^{\text{HI}} = -0.326$) and lower-tail curvature ($b^{\text{LO}} = -0.024$), with Δb significantly negative in block-bootstrap tests (full bootstrap $p = 0.012$; concentrated bootstrap $p \leq 0.006$ across block lengths 14–90 days). The lower-tail estimate is itself not statistically distinguishable from zero ($p = 0.258$), so the apparent magnitude ratio reflects the small denominator as much as the upper-tail estimate; we treat the asymmetry as a difference, not a ratio. Under an expanding-window out-of-sample evaluation across four cutpoints, the asymmetric specification consistently reduces check-loss at the upper tail relative to the linear quantile power-law baseline (+26.7% to +29.1% at Q75%, Q95%, Q99% in cutpoint-averaged check-loss; Table 9), with the asymmetric specification favored in every upper-tail cell of the 3 quantile $\times$ 4 cutpoint grid (Table 9b; the four expanding-window cutpoints are heavily nested, so the twelve cells are not statistically independent). Lower-tail and central-quantile performance is sensitive to training-window composition: three of four cutpoints yield predictions essentially indistinguishable from the linear baseline at Q1%, Q10%, Q25%, while the 2022 cutpoint produces positive lower-tail curvature in the training fit and elevated check-loss during the 2022 drawdown, a pattern consistent with the limited-cycle identification caveat in Section 17. A reduced-form structural-plus-reflexivity model offers one mechanism consistent with the asymmetry, presented as sufficiency rather than uniqueness. The estimated curvature is a full-sample distributional summary across four halving cycles; sub-period estimates are not separately identifiable on individual sub-periods. The methodology is portable to other long-duration price processes where structural and speculative regimes plausibly differ in tail behavior.

Keywords: Bitcoin; quantile regression; power law; conditional quantiles; asymmetric curvature; rearrangement estimator; stock-to-flow; distributional modeling

JEL Classification: C21, C58, G12, G17

1. Introduction

Bitcoin's market capitalization has grown by roughly nine orders of magnitude over its history (from approximately 10^3 in late 2009, a retrospective estimate from a pre-liquidity period, to over 10^{12} at recent ATH) and has inspired a substantial literature on long-duration price modeling. Two families of models have dominated public discourse: power-law regression models, most prominently developed by Santostasi (2018–2024), and scarcity-based valuation models, exemplified by the stock-to-flow (S2F) framework and its cross-asset extension (S2FX) (PlanB, 2019; PlanB, 2020). Both achieved strong in-sample fit over their respective calibration periods, generating significant attention and shaping practitioner expectations about Bitcoin's long-run trajectory.

Several prominent implementations of these frameworks have exhibited persistent upward forecast errors since publication. An OLS power law calibrated through 2018 overpredicted on 77.2% of subsequent days (geometric mean error +32.1%, 2019–2026); the S2F model calibrated through 2019 overpredicted on 94.9% of days (+294.5%), with errors compounding to over +1,167% by 2026. The errors are systematic in direction and increase through time.

The systematic optimistic bias observed in several prominent constant-elasticity implementations motivates the present paper's central question: can a distributional framework that allows tail behavior to evolve through time better characterize Bitcoin's historical price distribution, particularly in the tails where prior models appear to struggle most? We provide evidence consistent with an affirmative answer, with three qualifications stated upfront: the curvature parameters are identified on the full sample but not on individual sub-periods (Section 10.1); the chosen tail-group partition is interpretable and competitive but not formally selected from the data (Section 10.4.1); and the quadratic specification has no literal interpretation beyond the evaluation horizon (Section 17).

Our approach builds on the practitioner quantile power-law literature, which extends Santostasi's OLS framework to the conditional distribution at multiple probability levels by estimating quantile regression lines. Prior quantile power-law specifications use linear quantile regression in log-log space, maintaining a constant log-time slope (constant log-log elasticity) at each quantile while adding distributional coverage; more recent extensions incorporate stretched-exponential decay functions to capture cycle compression. These are genuine advances that built systematically on earlier work, and the present paper builds on them in turn. We extend this research program in a different direction: rather than specifying a decay function for the full distribution, we ask whether the curvature of Bitcoin's conditional distribution differs between the upper and lower tails. The asymmetric specification introduced here estimates separate curvature parameters for the upper and lower tails. Upper-tail curvature is statistically distinguishable from zero; lower-tail curvature is not. The lower tail is therefore parsimoniously modeled as near-linear, consistent with the structural-support interpretation in prior quantile frameworks. Although the empirical application is to Bitcoin, the methodological contribution (a grouped-curvature quantile regression with rearrangement-enforced non-crossing) is portable to any long-duration price process in which

structural and speculative regimes plausibly differ in tail behavior; we return to this point in Section 18.

The principal empirical finding is a statistically significant asymmetry between upper-tail and lower-tail curvature, with the difference $\Delta b = b^{\text{HI}} - b^{\text{LO}}$ significantly different from zero and only the upper-tail curvature itself distinguishable from zero. The lower-tail curvature is compatible with near-linear power-law dynamics in the sense that linearity cannot be rejected, consistent with prior power-law characterizations of Bitcoin’s lower conditional quantiles, while speculative upper-tail regimes are better captured by a quadratic correction. Economically, this pattern is consistent with the diminishing reflexivity of speculative capital as Bitcoin’s market capitalization grows (progressively larger capital flows being required to generate comparable percentage moves), though other mechanisms generating diminishing upper-tail amplitude would produce observationally similar behavior. The present paper documents the curvature asymmetry; identifying which mechanism produces it is left for future work.

We address the quantile crossing problem, a known limitation of unconstrained quantile regression that causes quantile estimates to become non-monotonic outside the sample, using the rearrangement estimator of Chernozhukov, Fernández-Val, and Galichon (2010). This approach guarantees monotonic quantile ordering on any evaluation grid and weakly reduces empirical check-loss relative to the unconstrained estimator (Section 4 derives the check-loss inequality from a Hardy-Littlewood-Polya argument). We verify quantile monotonicity on the daily evaluation grid through 2035 (not a guarantee for arbitrary out-of-grid extrapolations). The unconstrained symmetric model would have produced quantile crossings as early as December 2026, making the rearrangement correction practically necessary, not merely theoretically desirable.

The paper proceeds as follows. Section 3 documents the systematic optimistic bias in three prior models, establishing the empirical motivation for a distributional approach. Sections 4 through 6 develop the mathematical framework, data, and model hierarchy. Sections 7 through 9 present the main empirical results: coefficient estimates, the asymmetry test, and model comparisons. Sections 10 through 13 provide robustness evidence including sub-period stability, genesis date sensitivity, expanding-window evidence for the asymmetry, tail-group partition consistency and formal BIC-based partition selection, non-crossing verification, liquidity dislocation analysis, and out-of-sample validation. Section 14 benchmarks the model against prior quantile power-law specifications. Section 15 provides the theoretical derivation showing that Proposition 1 holds under a structural-plus-reflexivity price model. Section 16 interprets results economically. Sections 17 and 18 discuss limitations and conclude.

2. Prior Literature

2.1 Power-Law and Logarithmic Regression Models

The literature spans peer-reviewed academic work (surveyed in Section 2.4) and practitioner grey literature (Sections 2.1–2.3, 2.5); grey-literature entries are marked as such with URLs in the

references. The earliest formal attempt at long-run Bitcoin price modeling appears to be the logarithmic regression published on October 22, 2014 by a pseudonymous BitcoinTalk user known as Trolololo, in a thread titled “Logarithmic (non-linear) regression — Bitcoin estimated value.” The published formula, $\log_{10}(P) = 2.9065 \cdot \ln(t) - 19.493$, where t is days since January 9, 2009 (Trolololo’s anchor, six days after the genesis block), is a log-log specification: with $\ln(t)$ on the right-hand side, it is equivalent to the power law $P \propto t^n$ with exponent $n = 2.9065 \cdot \ln(10) \approx 6.69$. It served as the analytical backbone for what later became the Bitcoin Rainbow Chart, widely used in the practitioner community, and is the same functional class as the Bitcoin Power Law that Santostasi (2018–2024) later formalized and motivated theoretically, differing only in the magnitude of the estimated exponent and the choice of time anchor. The observation that Bitcoin’s price follows an approximate power law in time was first introduced by Santostasi in 2018 (Santostasi, 2018–2024), building on earlier informal analyses including Trolololo’s logarithmic regression. The model posits that Bitcoin price $P(t)$ relates to time t (measured in days since the genesis block) as:

$$P(t) = C \cdot t^G$$

which in logarithmic form yields:

$$\log_{10}(P(t)) = \log_{10}(C) + \alpha \cdot \ln(t), \quad \text{where } \alpha = G / \ln(10)$$

This constant-elasticity specification is parsimonious and has shown remarkable in-sample fit, with pseudo- R^2 values exceeding 0.90 across the full Bitcoin price history. Santostasi motivates the power law through analogies to biological scaling, network effects, and the Metcalfe’s Law relationship between network value and the square of active users. The framework has been extended by various researchers including Fulgur Ventures (2024), who formalized the theoretical derivation and introduced the Dynamic Power Cycle modification, and Baquero (2026), who develops an activity-warped power law that rescales calendar time by a measure of on-chain network activity (price volatility or transaction volume), producing a tighter in-sample fit and improved walk-forward out-of-sample performance. These extensions preserve the constant-elasticity backbone of the Santostasi specification.

2.2 Scarcity and Stock-to-Flow Models

The stock-to-flow (S2F) model (PlanB, 2019) models Bitcoin price as a function of its stock-to-flow ratio $SF(t) = S(t)/F(t)$, where $S(t)$ is circulating supply and $F(t)$ is annual issuance. The model:

$$\log_{10}(P(t)) = \beta_0 + \beta_1 \cdot \log_{10}(SF(t)) + \varepsilon(t)$$

achieved strong in-sample fit through 2019, and was interpreted by practitioners as anticipating Bitcoin’s approach to \$10,000 and subsequent cycles. Section 3 of this paper documents that the S2F model has exhibited systematic optimistic bias since its publication, with errors growing over time as realized prices diverged from halving-driven scarcity predictions. Shelton (2024) tests S2F alongside Metcalfe’s Law, technical indicators, and market sentiment as predictors of Bitcoin’s monthly returns, finding that S2F and Metcalfe’s Law help explain returns in-sample but lose

predictive content out-of-sample, while technical and sentiment measures perform poorly in both settings. That work addresses a different statistical object (short-horizon return prediction) than the long-run conditional price-level distribution examined here, but reaches conclusions about S2F's out-of-sample bias that are directionally consistent with the findings in Section 3.

2.3 Quantile Regression Extensions

Quantile regression generalizes ordinary least squares by estimating multiple conditional quantiles of the response variable rather than only the conditional mean. Where OLS minimizes squared residuals to fit the average, quantile regression at level $\tau \in (0, 1)$ minimizes the "check-loss" function $\rho_{\tau}(u) = u \cdot (\tau - 1 \{u < 0\})$, an asymmetric loss that weights positive and negative residuals differently. Fitting separately at, say, $\tau = 0.01, 0.10, \dots, 0.99$ traces out the full conditional distribution of the response, not just its center. This makes quantile regression a natural tool for characterizing distributional structure that varies across the tails of an asset's price distribution. Prior to the formal development of the quantile power-law framework, Cowen (2020a) explored an informal precursor in practitioner research: fitting separate upper- and lower-envelope logarithmic regression bands to Bitcoin's price history, with the lower band fit exclusively to non-bubble observations and the upper band fit to the speculative-peak trajectory. The two fitted bands exhibited visibly different curvature, with the upper band bending inward more sharply than the lower band as the sample extended. In retrospect, this approach anticipates both the intuition underlying lower-quantile regression (isolating the structural support region while discarding upper-tail observations that inflate the fitted trend) and the asymmetric-tail observation that the upper and lower envelopes are governed by distinct mechanisms with distinct curvature. The 2020 work began as early as April 2020 (video discussion of multi-band logarithmic regression for support and resistance) and the asymmetric upper/lower-envelope visualization was published on August 17, 2020. It lacked the statistical machinery (check-loss minimization, multi-quantile estimation, distributional inference) that defines formal quantile regression, and the asymmetric-curvature observation was visual rather than tested; the present paper formalizes both elements within a single quantile regression specification with a tested asymmetry parameter.

Prior quantile power-law work (Plan C, 2025a) extended the framework to multiple regions of the conditional distribution using linear quantile regression at several probability levels in log-log space. This approach, which we term the "quantile power law," provides probabilistic bands around the central trend and has been widely cited in the Bitcoin practitioner literature. That prior v1 framework (Plan C, 2025a) uses linear quantile regression, maintaining a constant log-time slope at each quantile, and is the direct predecessor of the present paper. A subsequent v2 framework (Plan C, 2025b) uses piecewise quantile regression incorporating a stretched-exponential decay function, selected via Akaike Information Criterion scoring, with guaranteed non-crossing quantile lines. The present paper extends this research program by introducing asymmetric curvature that differs between the upper and lower tails, a structure motivated by the distinct economic mechanisms driving speculative peaks and structural support, while maintaining the parsimony that guards against overfitting in a dataset with limited cycle observations.

Several methods have been proposed to enforce quantile non-crossing. Bondell, Reich, and Wang (2010) develop a constrained optimization approach that imposes non-crossing as a linear constraint during estimation. Dette and Volgushev (2008) propose transformation-based estimators that guarantee crossing-free quantile curves by construction. Cannon (2018) reviews machine-learning approaches including monotone composite quantile regression. The rearrangement estimator of Chernozhukov, Fernández-Val, and Galichon (2010) is adopted as a post-processing step: it weakly reduces empirical check-loss, preserves the asymptotic distribution at continuity points, and improves finite-sample MSE under standard regularity conditions.

Quantile regression was introduced by Koenker and Bassett (1978) and treated comprehensively in Koenker (2005). The Chernozhukov, Fernández-Val, and Galichon (2010) rearrangement estimator used here weakly reduces the L^p distance to the true conditional quantile function while guaranteeing monotonicity; the accompanying weak reduction in empirical check-loss follows from a Hardy–Littlewood–Polya argument derived in Section 4. Quantile regression has been applied to Bitcoin return prediction by Troster et al. (2019) and to cryptocurrency connectedness analysis by Bouri et al. (2021), though neither addresses long-horizon distributional structure.

2.4 Academic Crypto Asset Pricing Literature

The upper-tail behavior documented in this paper is also related to the Bitcoin bubble-detection literature. Phillips, Wu, and Yu (2011) and Phillips, Shi, and Yu (2015) develop recursive unit-root tests for explosive price behavior, and Cheah and Fry (2015) apply these to Bitcoin and find evidence of speculative bubbles. This literature establishes that Bitcoin's upper-tail excursions have a speculative character distinct from fundamental value, consistent with the asymmetric curvature documented here. A growing academic literature examines cryptocurrency pricing from an asset pricing perspective. Liu and Tsyvinski (2021) show that cryptocurrency returns are driven by crypto-specific factors including network proxies for user adoption, with strong time-series momentum and investor-attention effects; the three-factor model of Liu, Tsyvinski, and Wu (2022) captures cross-sectional variation in expected cryptocurrency returns via market, size, and momentum factors.

Biais et al. (2023) provide a rigorous theoretical treatment of Bitcoin equilibrium pricing: an overlapping generations model in which fundamental value is the stream of net transactional benefits and the equilibrium price path can exhibit rational speculative bubbles. If Bitcoin's price trajectory combines fundamental and speculative components, an asymmetric distributional structure similar to that documented here is a plausible implication.

Where Liu and Tsyvinski (2021) and Liu et al. (2022) characterize the conditional mean and cross-section of returns, the present paper characterizes the conditional distribution and the asymmetric evolution of its tails.

2.5 Liquidity and Macro Frameworks

Alden (2020–2025), Pal (2020–2025), and Bittel (2021–2025) have developed frameworks emphasizing Bitcoin’s sensitivity to global liquidity, real rates, and macroeconomic cycles. Kristoufek (2015) provides early academic evidence that Bitcoin possesses properties of both a standard financial asset and a speculative one, with price dynamics driven by a combination of fundamental and sentiment factors that vary across time horizons. Bouri, Gupta, and Roubaud (2019) document herding behavior in cryptocurrency markets, suggesting that return co-movement partly reflects investor sentiment dynamics rather than fundamental co-variation. These perspectives are complementary: liquidity-driven deviations help explain why lower-tail dislocations sometimes breach structural quantile estimates and why upper-tail compression may accelerate as Bitcoin integrates with global capital markets. The present paper deliberately uses price history only, isolating the contribution of distributional curvature; macro-integrated modeling is left for future work.

3. Systematic Bias in Prior Models

3.1 Motivation

The power-law framework and its quantile extensions represent genuine intellectual contributions: with limited early data, researchers identified a durable empirical regularity (Bitcoin’s approximate log-log relationship with time across several orders of magnitude). The stock-to-flow framework similarly formalized scarcity–valuation in tractable form. The present paper does not challenge these contributions; it asks a narrower question: does the central-tendency characterization extend cleanly to the distributional tails, particularly the upper tail, given the longer dataset now available?

The evidence in this section is consistent with prominent implementations of these frameworks having exhibited persistent upward forecast errors out of sample, concentrated in the upper tail. Importantly, this is not a general critique of power-law modeling: a near-linear lower-tail power law is compatible with the present paper’s findings (b^{LO} not distinguishable from zero), while the bias documented for the OLS power law is specific to the upper tail, motivating a framework that allows upper-tail and lower-tail behavior to differ.

3.2 OLS Power-Law Bias

We estimate the OLS power-law model on data through December 31, 2018, the endpoint roughly contemporaneous with Santostasi’s public formalization, and evaluate forecast errors on the subsequent 2,698 observations through May 21, 2026. The estimated model is:

$$\log_{10}(\hat{P}(t)) = 2.5535 \cdot \ln(t) - 17.1156$$

Define the forecast error on day t as $e(t) = \log_{10}(\hat{P}(t)) - \log_{10}(P(t))$, where positive values indicate the model predicted too high. Results are reported in Table 1.

Table 1: OLS Power-Law Out-of-Sample Forecast Errors, 2019–2026

Year	Mean Error ($\log_{10}$)	% Days Optimistic	Geometric Mean Price Error
2019	+0.117	72.6%	+31.0%
2020	+0.167	93.4%	+46.7%
2021	−0.267	0.0%	−45.9%
2022	+0.174	66.0%	+49.2%
2023	+0.327	100.0%	+112.4%
2024	+0.140	99.5%	+38.0%
2025	+0.105	100.0%	+27.3%
2026	+0.334	100.0%	+115.8%
Full (2019–2026)	+0.121	77.2%	+32.1%

Note: Model estimated on data through December 31, 2018. Error = predicted − actual $\log_{10}$ price. Geometric mean price error = $10^{\bar{e}} - 1$, where $\bar{e}$ denotes the mean $\log_{10}$ forecast error; this is a ratio-of-geometric-means measure rather than the arithmetic mean of daily price ratios. The arithmetic mean of daily price ratios is typically larger and is reported as +51.4% for the full 2019–2026 row. 2021 shows negative mean error as Bitcoin’s realized price exceeded the model’s prediction during the bull market cycle.

3.3 Stock-to-Flow Model Bias (S2F and S2FX)

We evaluate two S2F-family models. The original S2F model (PlanB, 2019) regresses log Bitcoin price on log stock-to-flow ratio. The S2F Cross-Asset model (S2FX, PlanB, 2020) extends this by fitting a cross-asset regression across four Bitcoin “phase” cluster points representing distinct monetization epochs, yielding a market capitalization model. Both are evaluated out of sample using their authors’ published parameters.

The S2F model is estimated on data through December 31, 2019 using the published coefficients (PlanB, 2019); Morillon and Chacon (2022) provide a peer-reviewed analysis finding evidence of in-sample fit but limited practical trading value. The estimated model is:

$$\log_{10}(\hat{P}(t)) = 3.4012 \cdot \log_{10}(SF(t)) - 1.0456$$

The S2FX model (PlanB, 2020) uses four phase cluster points to estimate a cross-asset relationship between stock-to-flow and market capitalization. Using the published parameters (PlanB, 2020):

$$\ln(\text{mktcap}) = 12.7598 + 4.1167 \cdot \ln(SF(t))$$

$$\hat{P}(t) = \exp(12.7598 + 4.1167 \cdot \ln(SF(t))) / \text{supply}(t)$$

Table 2 reports out-of-sample errors. S2F bias accelerates to +1,167% by 2026; S2FX is substantially more optimistic still, with full-sample geometric mean price error +1,699% (projecting over \$5,000,000 by 2025–2026 as the post-2024 halving reduces issuance). Both models exhibit bias in the same direction with growing magnitude, consistent with a shared structural feature of stock-to-flow implementations.

Table 2: S2F Out-of-Sample Forecast Errors, 2020–2026

Year	Mean Error ($\log_{10}$)	% Days Optimistic	Geometric Mean Price Error
2020	+0.445	67.8%	+178.5%
2021	+0.194	100.0%	+56.3%
2022	+0.465	100.0%	+191.8%
2023	+0.462	100.0%	+189.7%
2024	+0.844	100.0%	+597.3%
2025	+0.970	100.0%	+833.8%
2026	+1.103	100.0%	+1,167%
Full (2020–2026)	+0.596	94.9%	+294.5%

Note: S2F estimated on data through December 31, 2019 (published coefficients, PlanB, 2019). Error = predicted – actual $\log_{10}$ price. Geometric mean price error = $10^{\bar{e}} - 1$, where $\bar{e}$ is the mean $\log_{10}$ forecast error. The arithmetic mean of daily price ratios is materially larger ($\approx +540\%$ for the full row vs +294.5% geometric). Halvings: Nov 28 2012; Jul 9 2016; May 11 2020; Apr 20 2024. S2FX errors in Table 2b.

Table 2b: S2FX Out-of-Sample Forecast Errors, 2020–2026

Year	Mean Error ($\log_{10}$)	% Days Optimistic	Geometric Mean Price Error
2020	+1.0955	100.0%	+1,146%
2021	+0.7709	100.0%	+490%
2022	+1.0394	100.0%	+995%
2023	+1.0344	100.0%	+982%
2024	+1.5644	100.0%	+3,568%
2025	+1.7542	100.0%	+5,579%
2026	+1.8862	100.0%	+7,596%
Full	+1.2550	100.0%	+1,699%

Note: S2FX evaluated from March 23, 2020 (publication date) using the published parameters (PlanB, 2020). The 2026 row reflects partial-year data (January 1 – May 21, 2026): $\ln(\text{mktcap}) = 12.7598 + 4.1167 \cdot \ln(\text{SF})$. Price derived as $\text{mktcap} / \text{circulating supply}$. Geometric mean price error = $10^{\bar{e}} - 1$ as in Tables 1 and 2. The increase in 2024–2026 forecast errors reflects the post-April 2024 halving reducing issuance to 3.125 BTC/block, causing the SF ratio to approximately double and the model to project prices exceeding \$5,000,000.

These results are not presented to discredit prior researchers; the OLS power law, S2F, and S2FX frameworks contributed substantially to public understanding of Bitcoin’s long-duration dynamics during their calibration periods. The purpose is narrower: to establish that constant-elasticity and scarcity-based frameworks may omit something systematic. S2FX, the more elaborate model with a stronger theoretical story, produces substantially larger errors than the simpler S2F, reinforcing the case for parsimony and distributional characterization over point prediction.

4. Mathematical Framework

4.1 Standard Power Law

Notation: $\ln$ denotes natural log; $\log_{10}$ denotes common log; prices enter as $\log_{10}(P)$ and time as $\ln(t)$. The standard power-law model in logarithmic form is:

$$\log_{10}(P(t)) = \beta_0 + \beta_1 \cdot \ln(t)$$

where t is days since the genesis anchor (January 1, 2009), β_0 is the intercept, and β_1 is the slope of $\log_{10} P$ with respect to $\ln t$. Equivalently, $P(t) = C \cdot t^G$ with $C = 10^{\beta_0}$ and $G = \beta_1 \cdot \ln(10)$. The constants here (β_0, β_1) differ from the per-quantile parameters ($c\tau, a\tau$) introduced in Section 4.2. This constant-elasticity assumption is the defining characteristic of both the OLS and quantile power-law frameworks.

4.2 Centered Quadratic Specification and Asymmetric Extension

To allow conditional elasticity to evolve through time, we introduce a quadratic term in log-time. The raw correlation between $\ln(t)$ and $\ln^2(t)$ in our sample is approximately 0.999, creating severe multicollinearity in an uncentered specification. We center log-time at its sample mean:

$$x = \ln(t) - \mu \quad \text{where } \mu = (1/n) \sum_i \ln(t_i) = 7.9914$$

Centering reduces the correlation between x and x^2 to approximately -0.618 , materially improving coefficient interpretability. The centered quadratic quantile model at quantile level τ is:

$$Q\tau(\log_{10}(P(t))) = c\tau + a\tau \cdot x + b\tau \cdot x^2$$

The slope of log-price with respect to log-time is:

$$d \log_{10}(P) / d \ln(t) = a\tau + 2b\tau \cdot x$$

When $b\tau < 0$, the log-time elasticity $d \log_{10} P / d \ln t$ declines over time. The second derivative $d^2 \log_{10}(P) / d(\ln t)^2 = 2b\tau$ captures the curvature of log-price in log-time; this is a statement about log-log geometry, not about calendar-time growth.

4.3 Asymmetric Curvature Specification

The symmetric quadratic estimates a free curvature parameter $b\tau$ at each quantile independently. ("Symmetric" refers to the absence of imposed asymmetry between upper and lower tails.) Section 8 formally tests whether this symmetry assumption is supported by the data and rejects it: the difference $\Delta b = b^{HI} - b^{LO} = -0.302$ is significantly negative at the 5% level ($p = 0.012$). The lower-tail estimate is not distinguishable from zero ($p = 0.258$); the asymmetry is identified by the upper-tail curvature. Table 5 reports that the concentrated bootstrap yields $p \leq 0.006$ across block lengths 14–90 days. This empirical finding motivates the asymmetric specification introduced here, which allows b to differ between tail regions to reflect the distinct mechanisms driving speculative peaks and structural support. The asymmetric specification partitions quantiles into three groups:

$$b\tau = b^{LO} \quad \text{for } \tau \in \{0.01, 0.10, 0.25\} \quad (\text{lower tail, shared})$$

$$b\tau = b^{\text{MED}} \quad \text{for } \tau = 0.50 \quad (\text{median, free})$$

$$b\tau = b^{\text{HI}} \quad \text{for } \tau \in \{0.75, 0.95, 0.99\} \quad (\text{upper tail, shared})$$

This specification has 17 free parameters: seven intercepts $c\tau$, seven slopes $a\tau$, and three curvature parameters. The lower and upper tail curvature parameters are estimated by jointly minimizing the sum of check-loss functions within each tail group. The choice of three curvature parameters is shown in Section 10.4.1 to be competitive under a BIC-style penalized check-loss criterion (rank 4 of 877 partitions; $\Delta\text{BIC} = 5.26$ vs the minimizer; Kass and Raftery, 1995). The criterion is BIC-style rather than formal BIC because the likelihood is not fully specified.

The check-loss (tick) function for quantile τ is:

$$\rho\tau(u) = u \cdot (\tau - \mathbb{1}\{u < 0\})$$

where $\mathbb{1}\{\cdot\}$ denotes the indicator function (equal to 1 when the bracketed condition holds, 0 otherwise). The lower-tail joint estimation solves:

$$\min_{\{b^{\text{LO}}, \{c\tau, a\tau\}\}} \sum_{\tau \in \Lambda^{\text{LO}}} \sum_i \rho\tau(y_i - c\tau - a\tau x_i - b^{\text{LO}} x_i^2)$$

where $\Lambda^{\text{LO}} = \{0.01, 0.10, 0.25\}$, and symmetrically for the upper tail $\Lambda^{\text{HI}} = \{0.75, 0.95, 0.99\}$.

4.4 Rearrangement Estimator

Unconstrained quantile regression does not guarantee that estimated quantile functions are monotonically ordered at all evaluation points. When quantile functions cross, the model produces the logical contradiction that a higher quantile is associated with a lower price. We address this using the rearrangement estimator of Chernozhukov, Fernández-Val, and Galichon (2010).

Let $\{\hat{Q}\tau(x)\}_{\tau \in T}$ be the set of unconstrained quantile estimates at evaluation point x , where $T = \{0.01, 0.10, 0.25, 0.50, 0.75, 0.95, 0.99\}$. The rearranged estimate $\tilde{Q}\tau(x)$ is defined as:

$$\tilde{Q}\tau(x) = \hat{F}^{-1}(\tau \mid x)$$

where $\hat{F}(\cdot \mid x)$ is the empirical distribution function of $\{\hat{Q}\tau(x) : \tau \in T\}$ evaluated at x . In practice, rearrangement sorts the vector of quantile predictions at each evaluation point x into ascending order. Chernozhukov, Fernández-Val, and Galichon (2010, Theorem 1) prove that rearrangement weakly reduces the L^p distance to the true conditional quantile function. The empirical-check-loss reduction used in this paper is a separate consequence of the Hardy-Littlewood-Polya rearrangement inequality applied to the per-observation sum: $\rho\tau(y - q)$ is submodular in (τ, q) (the mixed partial $\partial^2 \rho\tau / (\partial \tau \partial q)$ equals -1 almost everywhere), so reordering the predicted values $\{\hat{Q}\tau(x)\}_{\tau}$ at any x to match the natural ordering of the τ levels weakly reduces $\sum \rho\tau(y - \hat{Q}\tau(x))$. Summing over observations:

$$\sum_{\tau \in T} \sum_i \rho\tau(y_i - \tilde{Q}\tau(x_i)) \leq \sum_{\tau \in T} \sum_i \rho\tau(y_i - \hat{Q}\tau(x_i))$$

That is, monotone rearrangement weakly reduces the empirical check-loss while guaranteeing $\tilde{Q}\tau_1(x) \leq \tilde{Q}\tau_2(x)$ for all $\tau_1 < \tau_2$ and all x . We verify non-crossing on a dense evaluation grid covering the historical sample and extrapolation through December 31, 2035.

5. Data and Reproducibility

The analysis uses daily Bitcoin price observations from July 17, 2010 through May 21, 2026 (5,788 observations, no calendar gaps), obtained as a publicly available aggregated daily CSV. The sample begins on July 17, 2010, the first date with a non-degenerate price in the CSV (\$0.05); earlier 2009–2010 observations are exchange-specific, sparse, and pre-liquidity, and we follow the standard practitioner convention of starting in July 2010. The early portion of the sample is informative about the long-run distributional arc rather than as a high-frequency price record; we document anchor and sub-period sensitivity in Sections 10.1 and 10.2. All regression estimates use the close field; volume and market capitalization appear only descriptively in Section 1. Intraday wick lows are used only in the Section 12 dislocation analysis.

The time variable is defined as $t_i = \text{days since January 1, 2009 (the genesis anchor)}$. Section 10.2 verifies robustness to alternative anchor choices including the genesis block date (January 3, 2009) and the Bitcoin whitepaper date (October 31, 2008). The dependent variable is $y_i = \log_{10}(P_i)$, where P_i is the daily closing price.

The raw correlation between $\ln(t)$ and $\ln^2(t)$ in the sample is 0.999, confirming severe multicollinearity in the uncentered specification. Centering at $\mu = 7.9914$ reduces the correlation between x and x^2 to -0.618 , as established in Section 4.2.

6. Model Hierarchy

We compare three nested models, each building on the previous:

Model	Specification	Elasticity	Parameters
Linear quantile power law	$Q\tau = \alpha\tau \ln(t) + \beta\tau$	Constant per quantile	2 per quantile
Symmetric quadratic	$Q\tau = c\tau + a\tau x + b\tau x^2$	Evolving, symmetric	3 per quantile
Asymmetric quadratic (proposed)	$Q\tau = c\tau + a\tau x + b(\tau)x^2$	Evolving, asymmetric	17 total

Note: $x = \ln(t) - \mu$ is centered log-time. $b(\tau)$ denotes the tail-group curvature: b^{LO} for $\tau \leq 0.25$, b^{MED} for $\tau = 0.50$, b^{HI} for $\tau \geq 0.75$.

The quantile levels estimated are $\tau \in \{0.01, 0.10, 0.25, 0.50, 0.75, 0.95, 0.99\}$. Linear quantile regression baselines (Section 6, model 1) are estimated using the modified Barrodale-Roberts simplex algorithm as implemented in the R `quantreg` package (Koenker, 2005). The symmetric and asymmetric quadratic specifications impose cross-quantile constraints (shared b^{LO} across the lower tail; shared b^{HI} across the upper tail) that fall outside `quantreg`'s single-quantile interface, so those models are estimated by directly minimizing the pooled check-loss using Nelder-Mead with tolerance 10^{-10} and 100 random initializations. As verification, re-estimating each quantile of the asymmetric model with `quantreg` holding curvature fixed at the Nelder-Mead solution agrees to

within 10^{-4} at every quantile; the 27 stable expanding-window re-estimates of Section 10.3 provide additional de facto convergence evidence.

7. Full-Sample Results

7.1 Asymmetric Model Coefficients

Table 3 reports the estimated coefficients for the rearranged asymmetric quadratic model. The centering constant $\mu = 7.9914$ corresponds to a reference time of approximately $e^{7.9914} \approx 2,955$ days since genesis, or early February 2017.

Table 3: Rearranged Asymmetric Quadratic Quantile Model: Coefficient Estimates

Quantile (τ)	$c\tau$	$a\tau$	$b(\tau)$	Pseudo- R^2
1%	2.837	2.578	$-0.0241 (b^{LO})$	0.905
10%	2.933	2.552	$-0.0241 (b^{LO})$	0.891
25%	3.004	2.554	$-0.0241 (b^{LO})$	0.861
50%	3.214	2.482	$-0.1126 (b^{MED})$	0.808
75%	3.562	2.283	$-0.3259 (b^{HI})$	0.747
95%	3.897	1.964	$-0.3259 (b^{HI})$	0.654
99%	4.028	1.904	$-0.3259 (b^{HI})$	0.619

Note: b^{LO} is the shared lower-tail curvature ($\tau \leq 0.25$); b^{HI} is the shared upper-tail curvature ($\tau \geq 0.75$). Pseudo- $R^2 = 1 - L^M/L_0$ where L^M is model check-loss and L_0 is the check-loss of the unconditional quantile estimator.

7.2 Block-Bootstrap Standard Errors

Bitcoin price data exhibit strong autocorrelation, heteroskedasticity, and regime clustering. Standard asymptotic standard errors for quantile regression are therefore unreliable. We estimate standard errors using a moving-block bootstrap on the (t, y) pairs with $B = 1,000$ replications and block length $\ell = 30$ calendar days, following the approach of Koenker (2005) for dependent data. Procedure: blocks of ℓ consecutive (t_i, y_i) pairs are sampled with replacement and concatenated to the original sample length; the asymmetric model is refit and curvature estimates recorded. Block resampling distorts the strictly increasing t -grid of the original sample (the standard cost of preserving within-block dependence) and captures only the dependence structure approximated by the chosen block length. Table 5 reports sensitivity across $\ell \in \{14, 30, 60, 90\}$; the asymmetry holds throughout, with p -values understood as conditional on a short-memory approximation rather than exact under long-range dependence. Table 4 reports standard errors and 95% confidence intervals for the three curvature parameters.

Table 4: Block-Bootstrap Standard Errors for Curvature Parameters

Parameter	Estimate	SE	95% CI	p-value ($H_0: b \geq 0$)
b^{LO} (lower tail)	-0.0241	0.036	[-0.077, +0.065]	0.258
b^{MED} (median)	-0.1126	0.097	[-0.276, +0.081]	0.176
b^{HI} (upper tail)	-0.3259	0.098	[-0.455, -0.072]	0.019 **

Note: ** $p < 0.05$ (one-sided bootstrap test). $B = 1,000$ replications, block length 30 days. p-value is the proportion of bootstrap draws with $b \geq 0$. Table 5 reports block-length sensitivity. The full bootstrap (Table 4) yields $p = 0.012$ for the asymmetry vs the concentrated bootstrap's $p \leq 0.006$ across block lengths 14–90, because the full procedure re-estimates all parameters jointly. Both support the asymmetry, with the 95% CI for Δb excluding zero across all block lengths in Table 5.

Table 5: Bootstrap Block-Length Sensitivity: Asymmetry Test ($\Delta b = b^{HI} - b^{LO}$)

Block Length	SE(Δb)	p-value	95% CI	Significant?
14 days	0.0281	<0.001 ***	[-0.361, -0.258]	Yes ***
30 days (baseline)	0.0408	<0.001 ***	[-0.381, -0.226]	Yes ***
60 days	0.0596	0.006 ***	[-0.401, -0.177]	Yes ***
90 days	0.0656	<0.001 ***	[-0.443, -0.168]	Yes ***

Note: $\Delta b = b^{HI} - b^{LO} = -0.302$ in all cases (full-sample estimate). $B = 500$ bootstrap replications per block length. p-value is the proportion of bootstrap draws with $\Delta b \geq 0$. 95% CI is the bootstrap percentile interval. Baseline block length (30 days) highlighted. *** $p < 0.01$.

8. The Asymmetry Test

The central hypothesis of this paper is that upper-tail curvature is more negative than lower-tail curvature. Formally, define the asymmetry as $\Delta b = b^{HI} - b^{LO}$. Under the null hypothesis of symmetric curvature, $\Delta b = 0$. The alternative is $\Delta b < 0$ (upper tail curves more steeply).

We test this using the block-bootstrap distribution of Δb . In each of the 1,000 bootstrap replications, we re-estimate b^{LO} and b^{HI} jointly using the asymmetric specification and record $\Delta b = b^{HI} - b^{LO}$. The bootstrap p-value is the proportion of replications in which $\Delta b \geq 0$.

Results:

$$\Delta b \text{ (observed)} = -0.3259 - (-0.0241) = -0.3018$$

$$\text{Bootstrap SE}(\Delta b) = 0.095 \quad 95\% \text{ CI: } [-0.439, -0.057]$$


p-value = 0.012 (significant at the 5% level). Note: Table 5 reports $p \leq 0.006$ across block lengths 14–90 days using a concentrated bootstrap that fixes intercept and slope parameters at full-sample values and re-estimates only the curvature parameters b^{LO} and b^{HI} . The full bootstrap above re-estimates all parameters jointly, increasing sampling variation and yielding a higher (more conservative) p-value. Both procedures support the asymmetry finding.

The asymmetry $\Delta b = -0.302$ has a bootstrap 95% CI that excludes zero (Table 5: [-0.443, -0.168] at the most conservative block length). The corresponding ratio $|b^{HI}| / |b^{LO}| \approx 13.5$ is unstable

because b^{LO} is not distinguishable from zero; we report only the difference Δb , not the ratio. Together, these results support the following interpretation: the lower tail of Bitcoin’s conditional price distribution is consistent with near-linear power-law dynamics in the sense that linearity cannot be rejected, while the upper tail exhibits significant additional curvature that the linear specification misses. We emphasize that non-rejection is not acceptance: the lower-tail confidence interval $[-0.077, +0.065]$ is wide enough to accommodate economically meaningful negative curvature. The asymmetric specification captures a strictly weaker claim than “the lower tail is linear”: it imposes a parsimonious near-zero b^{LO} that the data do not reject, leaving room for future work to identify smaller lower-tail effects with longer data.

9. Model Comparison

Figure 1: Bitcoin Rearranged Asymmetric Quantile Regression Fan Structure, 2010–2026



Note: BTCUSD weekly chart on logarithmic scale. Colored bands represent the 1%, 10%, 25%, 50%, 75%, 95%, and 99% conditional quantile estimates from the rearranged asymmetric quadratic model. Dark green (Q1%) to dark red (Q99%). Dashed lines below Q1% show all four historical intraday-wick dislocation reference levels: -7.35% (November 2022, FTX collapse, wick low \$15,474), -17.4% (March 2020, COVID crash, wick low \$3,881), -22.6% (August 2015, wick low \$162), and -34.6% (August 2010, wick low \$0.032), each computed as the percentage deviation of the realized intraday low from the contemporaneous Q1% estimate. The shaded golden zone between the dislocation lines represents the full range of historical structural dislocations. Table 8 reports the same events using daily close prices (more conservative and reproducible). The progressive narrowing of the upper bands relative to lower bands reflects the estimated upper-tail curvature $b^{HI} = -0.326$.

Table 6 compares in-sample check-loss and pseudo- R^2 across the three models. The improvement metric is the percentage reduction in check-loss relative to the linear quantile power law baseline.

Table 6: Model Comparison, In-Sample Check-Loss and Pseudo-R²

Quantile	Linear R ²	Sym. Quad R ²	Asym. R ²	Sym. vs Lin.	Asym. vs Lin.
1%	0.902	0.907	0.905	+4.55%	+2.55%
10%	0.890	0.891	0.891	+0.29%	+0.29%
25%	0.861	0.861	0.861	+0.03%	+0.00%
50%	0.807	0.808	0.808	+0.73%	+0.73%
75%	0.735	0.747	0.747	+4.59%	+4.56%
95%	0.635	0.655	0.654	+5.50%	+5.02%
99%	0.590	0.620	0.619	+7.39%	+7.24%

Note: Improvement percentages represent reduction in quantile check-loss relative to the linear quantile power law. R² values are pseudo-R² based on check-loss. Asym. = rearranged asymmetric model (proposed).

The in-sample fit improvements are modest at the center quantiles and more meaningful at the tails. This pattern is consistent with the asymmetry test results: the quadratic term adds little at the center because estimated curvature is small and not statistically distinguishable from zero there, and adds more at the upper tail where curvature is large and statistically significant. The asymmetric model achieves comparable fit to the symmetric quadratic while imposing a more parsimonious structure and providing a cleaner economic interpretation.

10. Sub-Period Stability and Sensitivity Analysis

The full-sample asymmetry result in Section 8 invites four classes of objection that this section is designed to address in turn. Section 10.1 asks whether the curvature coefficients are locally identified on individual sub-periods or are full-sample distributional summaries; we report sub-period estimates and document that neither b^{HI} nor b^{LO} is reliably identified on the shorter sub-periods, and we interpret the full-sample estimates accordingly as long-run distributional summaries across cycles. Section 10.2 asks whether results are sensitive to the genesis-date anchor; we report estimates under three alternative anchors and find that the curvature estimate is stable to day-level anchor variations around January 1, 2009 but shifts materially under month-scale changes to the anchor. Section 10.3 asks whether the asymmetry is driven by any single cycle; we report an expanding-window asymmetry test across 27 evaluation windows. Section 10.4 asks whether the imposed tail-group partition is data-consistent (free-curvature estimates by quantile) and provides a formal model-selection comparison under a BIC-style penalized check-loss criterion (Section 10.4.1) over all 877 set-partitions of the seven estimated quantiles. Together with the block-length sensitivity already reported in Table 5 (Section 7), these analyses cover the principal robustness checks a reader would request before accepting the asymmetric specification as a description of Bitcoin’s historical distributional evolution.

10.1 Sub-Period Coefficient Stability

A key concern with any full-sample model is whether estimated coefficients are stable across sub-periods or merely reflect the overall arc of the data. Table 7 reports asymmetric curvature estimates across three sub-periods: Early (July 2010 – December 2017), Mid (January 2018 – May 2020), and Late (May 2020 – May 2026).

Table 7: Sub-Period Stability of Asymmetric Curvature Estimates

Period	n	b^{LO} (lower)	b^{HI} (upper)	Δb
Early (2010–2017)	2,711	−0.194	−0.701	−0.507
Mid (2018–2020)	862	+28.814	+13.867	−14.947
Late (2020–2026)	2,201	+2.025	+6.285	+4.260
Full (2010–2026)	5,788	−0.024	−0.326	−0.302

Note: b^{LO} is small and negative in the full sample and in the Early sub-period, but on the shorter Mid and Late sub-periods it takes large positive values; b^{HI} behaves similarly, strongly negative in the Early period and large and positive on the Mid and Late windows. Neither tail curvature is well identified on the shorter sub-periods. The full-sample $b^{HI} = -0.326$ is a distributional arc estimate, a parameter summarizing Bitcoin’s price evolution across multiple cycles rather than a locally stable regime coefficient.

The large positive mid- and late-period curvatures reflect weak identification on short windows, not economically meaningful curvature. A profile-likelihood-style 95% interval ($L^*(b) - L^*(\hat{b}) \leq 1.92$, the $\chi^2_{1,0.95}/2$ cutoff) for b^{HI} has width ≈ 27 in the Mid sub-period and ≈ 6 in the Late sub-period, vs ≈ 0.12 for the full sample (interval $[-0.39, -0.27]$). The profile-likelihood interval holds intercept and slope at their fitted values and is therefore narrower than the block-bootstrap Table 4 interval ($[-0.455, -0.072]$), which propagates additional sampling variation; the two are not directly comparable. Only the full sample and the longer Early sub-period ($b^{HI} \in [-0.80, -0.53]$ at the 95% profile threshold) contain enough cycle variation for reasonable precision. The full-sample $b^{HI} = -0.326$ is the reliable summary; the sub-period values document lack of local identification, not parameter instability.

The lower-tail curvature b^{LO} is small and negative on the full sample and Early sub-period, and not separately identified on Mid/Late. Full-sample lower-tail near-linearity and upper-tail curvature should be read as long-run distributional summaries across cycles, not regime-stable parameters. Macro liquidity variables may explain the expanding-window b^{HI} trajectory (Section 13.2).

10.2 Genesis Date Sensitivity

The choice of genesis anchor affects the time variable and hence all coefficient estimates.

Results are robust to day-level variations in the genesis anchor around January 1, 2009: the genesis block (January 3, 2009) yields $b^{HI} = -0.325$ versus baseline -0.326 . Larger, month-scale anchor shifts produce more material differences: the whitepaper date (October 31, 2008, approximately two months before the baseline) yields $b^{HI} \approx -0.39$, roughly a 0.06 shift in curvature relative to baseline. The curvature estimate is therefore stable to day-level anchor choices but sensitive to month-scale shifts, because earlier anchors lengthen the early-sample log-time span that drives the

quadratic term. An anchor at January 1, 2010 produces a materially different estimate ($b^H \approx 0$). The anchor change does not discard any observations (the price sample still begins July 17, 2010); it shifts the t-value assigned to each observation, reshaping the $\ln(t)$ distribution and the weight on early-Bitcoin observations. The genesis anchor is the natural choice as the start of Bitcoin's monetary history; January 1, 2010 is reported only to document sensitivity, not as a competing specification.

10.3 Expanding Window Evidence for the Asymmetry


To check whether the Section 8 asymmetry depends on any particular cycle, we repeat the concentrated bootstrap on an expanding-window basis: data through January 2013 ($n = 930$), adding six months at each step through January 2026 ($n = 5,678$). This yields 27 windows, each producing a bootstrap p-value for $H_0: b^H = b^{LO}$.

Figure 2 displays the results. The upper panel plots the bootstrap p-value at each window. All 27 windows yield $p \leq 0.015$, with 20 of 27 yielding $p < 0.005$. The evidence strengthens as data accumulate; no window fails to reject symmetric curvature at the 5% level. Because expanding windows share substantial overlap, the 27 p-values are not independent tests; the persistence indicates no single time period is responsible for the finding. The lower panel shows that the estimated b^H (upper-tail curvature, solid red) remains consistently in the range -0.29 to -0.35 throughout, while b^{LO} (lower-tail curvature, dashed green) stays near zero. This expanding-window behavior differs from the unstable sub-period estimates in Table 7 because expanding windows preserve information from earlier speculative episodes, whereas isolated sub-period estimation weakens identification of second-order terms.

These results partially address the limited-cycle criticism. Rather than treating the four-cycle span as a constraint on the finding, the expanding window analysis shows that the asymmetry is not explained by any single cycle. It was present before the 2017 peak, before the 2020 halving, and before the 2021 cycle, and it has not disappeared as more data have been added. The finding appears consistent with a persistent feature of Bitcoin's conditional price distribution rather than a sample-specific artifact.

Figure 2: Expanding-Window Asymmetry Test, January 2013 – January 2026

Figure 2: Expanding Window Asymmetry Test, 2013-2026



Note: Expanding window bootstrap ($B = 200$, block length 30 days). Each point uses all data from July 2010 through the indicated date. Upper panel: bootstrap p -value for $H_0: b^H = b^L$ (symmetric curvature). All 27 windows yield $p \leq 0.015$. Lower panel: estimated curvature coefficients b^H (upper tail, solid red) and b^L (lower tail, dashed green). Dashed purple verticals mark the four Bitcoin halving dates. Both panels share the same x -axis.

10.4 Data-Consistency of the Tail-Group Partition

The asymmetric specification in Section 4.3 imposes a partition in which Q1%, Q10%, and Q25% share a common curvature b^L and Q75%, Q95%, and Q99% share a common curvature b^H . This partition is motivated by the theoretical decomposition in Section 15 (structural support vs. speculative reflexivity) rather than estimated from the data. We assess whether the partition is consistent with the data by re-estimating the centered quadratic model with a free curvature parameter at each quantile and comparing the resulting bt to the imposed tail-group values.

The free- bt estimates are: Q1% = -0.058 , Q10% = -0.026 , Q25% = -0.013 , Q50% = -0.113 , Q75% = -0.301 , Q95% = -0.429 , Q99% = -0.368 . The three lower-tail estimates (-0.058 , -0.026 , -0.013) are all small in absolute terms ($|bt| \leq 0.06$); the three upper-tail estimates (-0.301 , -0.429 , -0.368) are substantially larger ($|bt| \geq 0.30$), between roughly $5\times$ and $30\times$ the lower-tail values, depending on which pair is compared. The free- bt estimates show scatter within each tail (upper-tail span roughly -0.30 to -0.43 ; lower-tail -0.058 to -0.013), and the Q50% estimate (-0.113) lies between them. The two groups are cleanly separated: the most negative lower-tail estimate (-0.058) is well above the least negative upper-tail estimate (-0.301), a gap of roughly 0.24 with

no overlap. This separation, rather than tight within-group clustering, supports the imposed partition.

A robustness check: moving Q25% and Q75% into the median group makes both b^{LO} (-0.036) and b^{HI} (-0.425) slightly more negative, and widens rather than narrows the asymmetry. We retain the symmetric three-quantile-per-tail partition for interpretability and lower OOS check-loss. Section 10.4.1 below provides a formal model-selection comparison of the partition via BIC-style penalized check-loss enumeration over all 877 set-partitions of the seven estimated quantiles.

10.4.1 Formal Partition Comparison via BIC-Style Penalized Check-Loss

The free-curvature consistency check above asks whether the imposed partition is consistent with quantile-by-quantile estimates. A complementary question is whether it is favored over alternative partitions by a formal model-selection criterion. We address this by enumerating all 877 set-partitions of the seven estimated quantiles $\{Q1\%, Q10\%, Q25\%, Q50\%, Q75\%, Q95\%, Q99\%\}$. For each partition, we fit a shared-curvature quadratic quantile regression per block (with free per-quantile intercept and slope, shared b within each block), sum the check-loss L across blocks, and compute the BIC-style penalized check-loss $BIC = 2L + p \cdot \log(N \cdot K)$, where $p = 2K +$ (number of distinct b 's), $K = 7$ is the number of quantiles, and $N \cdot K = 40,516$. Under an asymmetric Laplace likelihood with fixed scale, $-2 \log L$ is proportional to the pooled check-loss up to constants that cancel in pairwise differences, so the form is BIC-style without committing to a fully specified parametric likelihood. The $N \cdot K$ effective-sample choice treats each (observation, quantile) pair as a distinct unit; using N alone shrinks the penalty and tightens the rank ordering further. Estimation uses an LP-exact inner solver for the quantile regression conditional on b and a one-dimensional bounded search over b , exploiting convexity of the check-loss in b for fixed (c, a) .

Table 7b: BIC-Based Selection over All 877 Tail-Group Partitions

# of b 's	BIC-minimizing partition	Loss	BIC	ΔBIC
1	[Q1%, Q10%, Q25%, Q50%, Q75%, Q95%, Q99%] (all shared)	2173.65	4506.45	+47.94
2	[Q1%, Q10%, Q25%, Q50%] [Q75%, Q95%, Q99%]	2144.38	4458.51	0.00 (best)
3	[Q1%, Q10%, Q25%] [Q50%] [Q75%, Q95%, Q99%] (paper)	2141.70	4463.77	+5.26
4	[Q1%, Q50%] [Q10%, Q25%] [Q75%] [Q95%, Q99%]	2140.59	4472.14	+13.63
5	(best 5-b partition)	2140.08	4481.74	+23.23
6	(best 6-b partition)	2140.00	4492.19	+33.68
7	Each quantile its own b (fully free)	2139.93	4502.66	+44.15

Notes: Each row shows the BIC-minimizer within its block-count class (the 2-b row is the overall BIC winner). Two further partitions, one 3-b ($[Q01, Q50] \mid [Q10, Q25] \mid [Q75, Q95, Q99]$, BIC 4463.63) and one 2-b ($[Q01, Q10, Q25, Q50, Q99] \mid [Q75, Q95]$, BIC 4463.69), fall between the row-2 winner and the paper's spec, placing

the paper's 3-b specification 4th overall of 877. ΔBIC is the absolute difference vs the BIC-minimizing partition (the 2-b partition that pools Q50% into the lower block). The paper's 3-b specification ranks 4th of 877 by BIC and 5th by AIC. Kass and Raftery (1995) characterize ΔBIC in (2, 6] as "positive but not strong" evidence; ΔBIC in (6, 10] as "strong"; $\Delta BIC \geq 10$ as "very strong."

The paper's three-block specification [Q1%, Q10%, Q25%] | [Q50%] | [Q75%, Q95%, Q99%] ranks fourth of 877 partitions by BIC and fifth by AIC. The 877 partitions are not statistically independent candidates: they are nested and overlapping rearrangements of the same seven quantile assignments, so the ordinal rank carries no exact probabilistic interpretation. The rank does, however, place the imposed specification clearly within the small upper tail of the partition lattice under either criterion. The BIC-minimizing alternative pools Q50% into the lower-tail block: [Q1%, Q10%, Q25%, Q50%] | [Q75%, Q95%, Q99%], with two distinct b's. The gap is $\Delta BIC = 5.26$, which by Kass and Raftery (1995) constitutes "positive but not strong" evidence. The pooling forces the shared lower-block curvature toward zero: the free-Q50% estimate is -0.113 , the pooled lower+Q50% block estimate is -0.036 , and the BIC saving from one fewer parameter ($\log(40,516) \approx 10.6$) just outweighs the loss increase from this distortion.

We retain the three-block specification for two reasons. First, the median is the conventional reference quantile in the prior practitioner literature (Plan C, 2025a; 2025b); pooling it into the lower-tail block forces a near-zero pooled-median curvature, complicating direct comparison with the Section 3 OLS baseline. Second, $\Delta BIC = 5.26$ is within Kass–Raftery's "positive but not strong" range and shrinks further under the alternative N-based effective sample. All partitions with four or more distinct b's are decisively rejected ($\Delta BIC \geq 13.6$), as are the fully restricted single-b model ($\Delta BIC = 47.9$) and the fully free seven-b model ($\Delta BIC = 44.2$). The imposed three-block structure is competitive among all 877 alternatives.

11. Non-Crossing Verification

The unconstrained symmetric quadratic model would produce quantile crossings as early as December 2026, less than seven months beyond the sample end. Specifically, the Q75% and Q95% quantile curves would cross at $x \approx 0.798$, corresponding to approximately December 22, 2026.

After rearrangement, we verify non-crossing on two evaluation grids:

- In-sample grid: all 5,788 historical observations. Zero crossings at all adjacent quantile pairs.
- Extrapolation grid: daily observations from January 1, 2009 through December 31, 2035 (9,861 points). Zero crossings at all adjacent quantile pairs, with strictly positive inter-quantile gaps throughout.

The rearrangement procedure is implemented by sorting the vector of seven quantile predictions at each evaluation point into ascending order before converting to price levels. This is computationally equivalent to the theoretical definition of the rearranged estimator and does not require any constrained optimization. The 2035 extrapolation grid is used only to verify quantile

monotonicity through the evaluation window; the long-horizon behavior of the quadratic specification has no literal interpretation (see Section 17).

12. Historical Liquidity Dislocations

The 1% quantile boundary should not be interpreted as an absolute price floor. It is the empirical 1% conditional quantile of the historical price distribution: a level that historical realized prices have fallen below approximately 1% of the time. To characterize historical dislocations, we define the undershoot metric:

$$U(t) = [P_{\text{wick}}(t) - Q_{0.01}(t)] / Q_{0.01}(t)$$

where $P_{\text{wick}}(t)$ denotes the realized intraday wick low and $Q_{0.01}(t)$ is the contemporaneous 1% quantile estimate. Figure 1 reports this wick-based metric. Table 8 instead reports the closing-price analogue $U(t) = [P_{\text{close}}(t) - Q_{0.01}(t)] / Q_{0.01}(t)$, where $P_{\text{close}}(t)$ is the daily close: the close-price metric is more conservative (the daily close is by construction at least as high as the intraday wick low) and is reproducible from the public dataset. Negative values indicate the relevant price (intraday wick for Figure 1, daily close for Table 8) fell below the Q1% boundary. Using daily close prices, we identify all events where the lowest close over a 30-day window fell below the contemporaneous Q1% boundary, clustering breaches within 30 days into single events.

Table 8: Historical Liquidity Dislocation Events

Event	Peak Undershoot	Days Below	Recovery
September–October 2010	−12.21%	15	< 1 week
August–September 2015	−0.48%	2	< 1 week
November 2022–January 2023 (FTX)	−5.42%	40	< 2 weeks

Note: Peak undershoot figures in this table are computed from daily close prices in the historical dataset (i.e., the close price each day, with the lowest close over the dislocation window taken as the peak undershoot), a conservative and reproducible benchmark. Intraday-wick lows observed on individual exchanges or aggregated chart platforms produce deeper dislocations and are used as reference levels in Figure 1 and in the companion TradingView indicator. Wick-vs-close resolution: Figure 1 shows −7.35% for FTX (intraday wick) vs −5.42% here (daily close). The March 2020 (COVID) and August 2010 dislocations shown in Figure 1 (intraday wicks of −17.4% and −34.6% respectively) do not appear in Table 8 because the daily close did not fall below the Q1% boundary on either occasion. Events here are clustered with a minimum 30-day gap; all were transient, with prices recovering above the Q1% boundary within two weeks of the peak undershoot.

Three daily-close dislocation events are identified by the cluster criterion. The deepest undershoot was September–October 2010 (−12.21%, 15 days below Q1%), within Bitcoin’s first year of liquid trading. The longest-duration breach was the FTX collapse, November 2022 – January 2023 (−5.42% peak undershoot, 40 days below), consistent with idiosyncratic exchange insolvency rather than systemic liquidity withdrawal. A third event, August–September 2015, was minor (peak undershoot −0.48% over two days within a 29-day cluster window). The March 2020 (COVID) and August 2010 dislocations visible in Figure 1 are intraday-wick events: at the daily-close level, the price did not fall below the Q1% boundary in either case.

These events indicate the lower-tail quantiles function as descriptive summaries rather than mechanical floors: realized closes have temporarily fallen below the fitted Q1% curve and returned within two weeks of the peak. The September 2010 dislocation provides the deepest historical daily-close stress reference at -12.21% , while intraday wick excursions (Figure 1) have reached materially deeper at -34.6% (August 2010) and -22.6% (August 2015). All identified events predate Bitcoin’s full institutional integration and a prolonged macro-driven recession (see Section 17), so the Q1% band should be read as historical, not as a guarantee under unprecedented macro scenarios.

13. Out-of-Sample Validation

13.1 Procedure

We conduct an expanding-window out-of-sample evaluation using four training cutpoints: January 1, 2022; January 1, 2023; January 1, 2024; and January 1, 2025. At each cutpoint, every model evaluated (the asymmetric specification, the linear quantile power-law baseline, the Plan C v1 specification of Section 14, and a stretched-exponential alternative) is estimated on all data strictly preceding the cutpoint and evaluated on all data strictly following. No information from the post-cutpoint sample enters the training procedure. Hindsight qualification: the asymmetric specification was developed with knowledge of the 2022 drawdown and the documented post-2021 underperformance of constant-elasticity models. The 2022–2025 cutpoints fall within that visibility window, so the results below are stress-testing under explicit train/test separation, not a true forward prediction record. We report both the cutpoint-average mean check-loss (Table 9, equal-weighted across the four cutpoints) and the per-cutpoint structure (Table 9b); both use the cutpoint-average convention throughout Section 13 (Tables 9, 9b, 10, 11). Diebold–Mariano statistics (Table 10) are computed within each cutpoint window using Newey–West HAC variance, with the average signed DM reported as a descriptive summary rather than a combined inferential statistic.

13.2 Results

Table 9: Out-of-Sample Check-Loss Comparison (Expanding Window)

Quantile	Linear OOS Loss	Asymmetric OOS Loss
1%	0.004205	0.017542
10%	0.023189	0.033326
25%	0.047485	0.054484
50%	0.072444	0.089002
75%	0.074056	0.052543
95%	0.022781	0.016465
99%	0.005066	0.003713

Note: Mean check-loss averaged across the four expanding-window cutpoints (January 1, 2022; January 1, 2023; January 1, 2024; January 1, 2025): per-cutpoint mean check-loss is computed on the post-cutpoint out-of-sample observations, and the four per-cutpoint means are then arithmetically averaged. Each model is fit on data strictly preceding the cutpoint and evaluated on data strictly following. Upper-tail losses (Q75%, Q95%, Q99%) are lower for the asymmetric specification at every cutpoint; lower-tail and central cutpoint-average losses appear higher for the asymmetric specification, but these averages are dominated by a single cutpoint (2022) where the asymmetric model produces a positive lower-tail curvature in training and incurs large losses during the 2022 bear market. The cutpoint average is not the appropriate summary at the lower tail under this evaluation design; Table 9b decomposes the result by cutpoint and shows that three of four cutpoints produce asymmetric-vs-linear loss ratios within 10% of unity at Q1%, Q10%, and Q25%.

Table 9b: Per-Cutpoint Out-of-Sample Loss (Transparency Decomposition)

Quantile	Cutpoint	Asymmetric Loss	Linear Loss	Ratio (Asym/Linear)	Direction
1%	2022	0.061712	0.008705	7.09	linear better
1%	2023	0.002508	0.002481	1.01	tied
1%	2024	0.003099	0.002899	1.07	tied
1%	2025	0.002849	0.002735	1.04	tied
10%	2022	0.068566	0.029761	2.30	linear better
10%	2023	0.019232	0.019563	0.98	tied
10%	2024	0.024057	0.022058	1.09	tied
10%	2025	0.021449	0.021372	1.00	tied
25%	2022	0.082836	0.057051	1.45	linear better
25%	2023	0.041978	0.042933	0.98	tied
25%	2024	0.047079	0.044004	1.07	tied
25%	2025	0.046042	0.045952	1.00	tied
50%	2022	0.146049	0.115361	1.27	linear better
50%	2023	0.068352	0.083017	0.82	asym better
50%	2024	0.086243	0.041653	2.07	linear better
50%	2025	0.055364	0.049746	1.11	linear better
75%	2022	0.085974	0.101778	0.84	asym better
75%	2023	0.060734	0.091524	0.66	asym better
75%	2024	0.027783	0.059508	0.47	asym better
75%	2025	0.035682	0.043413	0.82	asym better
95%	2022	0.022866	0.026097	0.88	asym better
95%	2023	0.019530	0.024739	0.79	asym better
95%	2024	0.013117	0.020544	0.64	asym better
95%	2025	0.010347	0.019743	0.52	asym better
99%	2022	0.004968	0.005764	0.86	asym better
99%	2023	0.004288	0.005429	0.79	asym better
99%	2024	0.003021	0.004662	0.65	asym better
99%	2025	0.002574	0.004410	0.58	asym better

Note: Per-cutpoint mean check-loss for each quantile and each of the four expanding-window cutpoints. Ratio = asymmetric loss / linear loss; a ratio of 1.00 indicates the two specifications produce essentially identical predictions on that cutpoint's out-of-sample window. Upper-tail wins (Q75, Q95, Q99) are present at every cutpoint (every cell of the 3 × 4 grid favors the asymmetric model; ratios from 0.47 to 0.88). Lower-tail losses (Q1, Q10, Q25) are concentrated at the 2022

cutpoint; the 2023, 2024, and 2025 cutpoints produce lower-tail predictions essentially indistinguishable from the linear baseline.

Table 10: Diebold-Mariano Test for Out-of-Sample Gains (Asymmetric vs. Linear Baseline)

Quantile	Asym. Loss	Linear Loss	Mean DM
1%	0.017542	0.004205	−37.15
10%	0.033326	0.023189	−4.15
25%	0.054484	0.047485	−2.34
50%	0.089002	0.072444	−5.33
75%	0.052543	0.074056	+32.84
95%	0.016465	0.022781	+65.18
99%	0.003713	0.005066	+74.62

Note: Diebold–Mariano test (Diebold and Mariano, 1995) computed separately within each of the four out-of-sample windows, with Newey–West HAC variance and per-window bandwidth $\ell = \lfloor 4 \cdot (T/100)^{2/9} \rfloor$. Loss differential $d_i = \text{check-loss}(\text{linear}) - \text{check-loss}(\text{asymmetric})$; positive = asymmetric has lower loss. Mean DM is the arithmetic average of the four per-window DM statistics. We do not report a combined p -value across cutpoints because the four expanding-window OOS samples are heavily nested; the substantive evidence is better read from Table 9b.

Table 9b is the central diagnostic: upper-tail check-loss is materially lower for the asymmetric specification at every cutpoint (ratios at or below 0.88 across the upper-tail $\times$ cutpoint grid; the gap widens at later cutpoints as training incorporates more cycle-compression evidence). Lower-tail check-loss is essentially indistinguishable from the linear baseline at the 2023, 2024, and 2025 cutpoints (ratios 0.98–1.09 at Q1%, Q10%, Q25%); the apparent lower-tail underperformance in the cutpoint averages is driven entirely by the 2022 cutpoint, where ratios reach 1.45 to 7.09. Mechanism: the asymmetric fit through 2021 produces $b^{\text{LO}} = +0.105$ (positive, reflecting the 2020–2021 acceleration), and the 2022 bear market then drives realized prices well below the projected floor. By the 2023 cutpoint, training includes 2022 and b^{LO} returns to -0.011 , with lower-tail OOS performance stabilizing. This is the limited-cycle identification caveat of Section 17.

The upper-tail finding is robust across training sets. The expanding-window estimates of b^{HI} are -0.17 (training through end-2021), -0.23 (through end-2022), -0.32 (through end-2023), and -0.35 (through end-2024), with the full-sample value -0.326 in the same range; all four are strictly negative and economically non-trivial. These full-refit estimates span a wider range than the Section 10.3 concentrated-bootstrap estimates ($[-0.35, -0.29]$) because the OOS evaluation re-estimates all parameters at each cutpoint while the concentrated bootstrap fixes intercepts and slopes at full-sample values. The two procedures answer related but distinct questions. The signed DM statistics in Table 10 are large and positive at every upper-tail quantile. The asymmetric specification captures a real and persistent feature of Bitcoin’s upper-tail conditional distribution

that the linear baseline misses; b^{LO} , however, is not stably identifiable from training windows ending near a cycle peak.

14. Comparison with Plan C’s Quantile Models

14.1 Context and Motivation

The present paper builds on prior quantile power-law work (Plan C, 2025a; Plan C, 2025b). The v1 Bitcoin Quantile Model (Plan C, 2025a) introduced probability-weighted price bands via linear quantile regression in log-log space; practitioner band-style indicators exist (e.g., Cowen, 2020b), but the formal quantile-regression treatment is due to Plan C.

We extend rather than replace that framework by asking whether the log-price / log-time relationship exhibits different curvature at the upper tail than at the lower. The asymmetry test in Section 8 supports this hypothesis, and the proposed specification captures the feature with a parsimonious three-parameter curvature structure that preserves the interpretability of the underlying power-law framework.

14.2 Formal Statistical Comparison

We compare three models: Plan C’s v1 linear quantile specification ($\ln(t)$ anchored to January 3, 2009), a stretched-exponential decay specification $Q\tau(t) = a\tau \cdot \ln(t) + b\tau \cdot \exp(-c\tau \cdot (t/T)^{d\tau})$ consistent with Plan C’s described v2 functional form (fit by check-loss minimization on the same dataset, with T the sample median of t), and our proposed rearranged asymmetric quadratic. Because the exact v2 parameters are unpublished (Plan C, 2025b), this is a comparison against the functional form rather than a specific fitted model.

In-sample, the proposed specification and the stretched-exponential are statistically indistinguishable (pseudo- R^2 within 0.01 at every quantile; Table 10b), both improving on or matching Plan C v1. The proposed model’s in-sample check-loss improvement over v1 ranges from essentially zero at Q25% to +7.26% at Q99%, concentrated at the upper tail. The out-of-sample comparison (Table 11) distinguishes the two flexible specifications.

Table 10b: In-Sample Pseudo- R^2 Comparison (Plan C v1 vs Stretched-Exponential vs Asymmetric)

Quantile	Plan C v1 R^2	Stretched-Exp R^2	Asymmetric R^2
1%	0.9024	0.9067	0.9049
10%	0.8902	0.8904	0.8905
25%	0.8612	0.8612	0.8612
50%	0.8069	0.8085	0.8083
75%	0.7347	0.7474	0.7468
95%	0.6352	0.6566	0.6537
99%	0.5896	0.6268	0.6194

Notes: Pseudo- $R^2 = 1 - (\text{model check-loss}) / (\text{constant-quantile check-loss})$, computed on the full in-sample dataset ($n = 5,788$). All three models fit in log-log space. Plan C v1 = linear quantile power law (Plan C, 2025a). Stretched-

exponential follows Plan C's described v2 functional form $Q\tau(t) = a\tau \cdot \ln(t) + b\tau \cdot \exp(-c\tau \cdot (t/T)^d)$, with T the sample median of t ; per-quantile Nelder-Mead with 15 random initializations. Asymmetric is the Section 4.3 rearranged asymmetric quadratic. Maximum pseudo- R^2 difference between Stretched-Exp and Asymmetric is 0.0074 (at Q99%).

14.3 Out-of-Sample Comparison

Table 11 reports OOS check-loss averaged across the four Section 13 expanding-window cutpoints. The structure mirrors Tables 9 and 9b: substantial upper-tail improvement at every cutpoint; lower-tail and median averages dominated by 2022 cutpoint behavior.

Table 11: Out-of-Sample Check-Loss Comparison (Expanding Window)

Q	Plan C v1 OOS Loss	Str. Exp. OOS Loss	Proposed OOS Loss	vs v1	vs Str. Exp.
1%	0.004111	0.011217	0.017542	-326.76%	-56.38%
10%	0.023024	0.055121	0.033326	-44.74%	+39.54%
25%	0.047421	0.099408	0.054484	-14.89%	+45.19%
50%	0.072246	0.108636	0.089002	-23.19%	+18.07%
75%	0.073971	0.059217	0.052543	+28.97%	+11.27%
95%	0.022763	0.035503	0.016465	+27.67%	+53.62%
99%	0.005060	0.007524	0.003713	+26.63%	+50.66%

Note: Averaged across cutpoints at January 1, 2022, 2023, 2024, 2025. Improvement = (competitor - proposed) / competitor $\times$ 100; negative = competitor lower. Str. Exp. = best-fit stretched-exponential (Plan C v2 functional form) fit per quantile by check-loss minimization on each training set; results consistent across 25 random seeds per quantile. Because the exact v2 implementation is unavailable, the stretched-exponential findings are evidence about the functional class rather than direct assessment of prior practitioner work.

The proposed specification exhibits consistently lower out-of-sample check-loss than Plan C v1 at the upper tail (+27% to +29% at Q75%, Q95%, Q99%, with the asymmetric model favored at every cutpoint; Mean DM statistics in Table 10). At the lower tail and median, cutpoint-averaged OOS performance is sensitive to the training-window composition: the 2022 cutpoint produces a positive lower-tail curvature estimate in the asymmetric fit, which is then penalized by the realized 2022 drawdown, while the 2023, 2024, and 2025 cutpoints produce lower-tail predictions essentially indistinguishable from v1. Against the more flexible stretched-exponential specification, the proposed model delivers lower OOS check-loss at six of seven quantiles (Q10%, Q25%, Q50%, Q75%, Q95%, Q99%), with the stretched-exponential lower at Q1% only. The stretched-exponential underperforms at Q1% under expanding-window evaluation because its four-parameter decay structure overfits cycle-specific features of the training window that do not persist out of sample: the characteristic risk of flexible decay specifications on training sets covering fewer than three complete cycles.

These results complement rather than compete with prior quantile power-law work. The contribution is concentrated in the single additional parameter b^{H} ; the lower-tail OOS sensitivity

to training-window composition is consistent with the full-sample finding that b^{LO} is not statistically distinguishable from zero (Section 7.2).

15. Theoretical Motivation for Asymmetric Curvature

15.1 Setup

This section presents an illustrative reduced-form mechanism that is consistent with the documented asymmetric curvature; it is offered as a sufficiency argument rather than an identification claim, and alternative mechanisms that produce observationally similar behavior are discussed at the end of the section. The framework combines a structural component (long-horizon monetary demand) producing near-zero log-time curvature with a speculative component whose amplitude decays as Bitcoin matures, in the spirit of the rational-bubble literature (Brunnermeier, 2009; Froot and Obstfeld, 1991; Biais et al., 2023). The speculative amplitude generates strictly negative upper-tail curvature in the conditional quantile. Notation: in this section, $\log$ denotes the natural logarithm unless written explicitly as $\log_{10}$.

Let t denote days since the genesis anchor. Bitcoin's circulating supply schedule is bounded (asymptoting at 21 million) and increases in discrete halving steps; over finite windows, however, the realized supply path is well approximated by a slowly varying function of t , and for analytical tractability we write $S(t) \propto t^\delta$ with $\delta > 0$ as a local approximation. Define the structural capital flow $F^s(t)$ as the cumulative demand from long-horizon investors who value Bitcoin for its monetary properties. We model this as a power law: $F^s(t) = A \cdot t^\alpha$, reflecting the gradual monetization of Bitcoin across investor cohorts. The structural price is then:

$$P^s(t) = F^s(t) / S(t) = (A/S_0) \cdot t^{\alpha-\delta}$$

Taking logarithms: $\log_{10}(P^s(t)) = c + ((\alpha - \delta)/\ln 10) \cdot \ln(t)$, which is linear in $\ln(t)$ with zero second derivative. The lower-tail quantile, which tracks structural support, thus predicts near-zero curvature in theory. This is consistent with the empirical estimate $b^{LO} = -0.024$, which is not statistically distinguishable from zero ($p = 0.258$).

15.2 Speculative Component and Diminishing Reflexivity

Upper-tail price observations correspond to periods of elevated speculative demand. In the spirit of Biais et al. (2023), we model the speculative premium as a multiplicative component with amplitude $A > 0$ and a shape function $g(t)$ that decays as Bitcoin's market capitalization grows. Define the reflexivity shape:

$$g(t) = 1 / (1 + (t / t^*)^k)$$

where t^* is a saturation scale and $k > 0$ governs the rate of reflexivity decay. This logistic-saturation form captures the empirical observation that early-period Bitcoin cycles saw percentage gains of tens of thousands of percent, while later cycles have seen progressively smaller relative amplitudes. When t is much less than t^* , $g(t) \approx 1$ (the speculative multiplier saturates near its

maximum value $1 + A$); when t is much greater than t^* , $g(t) \approx (t^*/t)^k$ (the multiplier decays toward 1, i.e., the structural price). The amplitude A is unidentified from price-only data but is needed for the multiplier to match observed early-period speculative excursions: with A in the range 10^2 – 10^4 , the speculative price can exceed the structural price by orders of magnitude before reflexivity attenuates, consistent with Bitcoin's early bull cycles. The speculative upper-tail price is:

$$P^H(t) = P^s(t) \cdot (1 + A \cdot g(t)) = P^s(t) \cdot (1 + A/(1 + (t/t^*)^k))$$

Taking logarithms and letting $x = \ln(t) - \mu$:

$$\log_{10}(P^H(t)) = c^H + a^H \cdot x + (1/\ln 10) \cdot \log(1 + A \cdot g(e^{x+\mu}))$$

15.3 Sign of Upper-Tail Curvature

The curvature of $\log_{10}(P^H)$ with respect to x is determined by the second derivative of $h(x) = \log(1 + A \cdot g(e^{x+\mu}))$. Define $u(x) = (t/t^*)^k = e^{k(x+\mu-\ln t^*)}$, so that $du/dx = k \cdot u$, $g(x) = 1/(1+u)$, and $1 + A \cdot g = (1+u+A)/(1+u)$. Applying the chain rule:

$$dh/dx = [1/(1+A \cdot g)] \cdot d(A \cdot g)/dx = [(1+u)/(1+u+A)] \cdot A \cdot [-1/(1+u)^2] \cdot (du/dx) = -A \cdot k \cdot u / [(1+u)(1+u+A)]$$

Applying the quotient rule to $dh/dx = -A \cdot k \cdot u / D$, where $D = (1+u)(1+u+A) = 1 + 2u + u^2 + A + Au$ and $dD/dx = k \cdot u \cdot (2+2u+A)$, and collecting terms over the common denominator $(1+u)^2(1+u+A)^2$:

$$d^2h/dx^2 = -A \cdot k^2 \cdot u \cdot (1 + A - u^2) / [(1+u)^2(1+u+A)^2]$$

The reflexivity term is not globally concave: the sign of d^2h/dx^2 switches at $u = \sqrt{1 + A}$ (concave for $u < \sqrt{1 + A}$; convex beyond). The inflection location, $t = t^* \cdot (1 + A)^{1/(2k)}$, depends on (k, t^*, A) , none of which is identified from price-only data. The sign reversal is specific to logistic-saturation multipliers with constant amplitude; other monotonically decaying reflexivity specifications need not exhibit an inflection. A larger A pushes the convex regime to later times.

For consistency with the Section 7–8 finding of negative upper-tail curvature, the empirically relevant range must satisfy $u < \sqrt{1 + A}$, equivalently $t^* \cdot (1 + A)^{1/(2k)} \geq T_{\text{end}}$ (where $T_{\text{end}} \approx 6,350$ days). This is a mild constraint amounting to Bitcoin not yet having fully traversed its reflexivity-compression regime. Within that range, the reflexivity component generates strictly negative upper-tail curvature; the lower tail (without the modeled reflexivity premium) has zero curvature. Alternative mechanisms producing diminishing upper-tail amplitude (herding (Bouri et al., 2019), saturating adoption) would generate similar qualitative behavior. Proposition 1 summarises:

Proposition 1. Under the structural-plus-reflexivity price model, evaluated over the range $u \leq \sqrt{1 + A}$ (equivalently $t^* \cdot (1 + A)^{1/(2k)} \geq T_{\text{end}}$), the log-price versus log-time relationship satisfies: (i) $b^{LO} = 0$ (lower-tail curvature is zero); (ii) $b^{HI} < 0$ (upper-tail curvature is strictly negative); and (iii) $b^{HI} < b^{LO}$ (asymmetry). The reflexivity term changes

sign at $u = \sqrt[3]{1 + A}$, so these properties are local to the pre-saturation regime rather than global features of the functional form.

15.4 Discussion

Proposition 1 establishes that the structural-plus-reflexivity framework generates the empirical pre-saturation pattern: zero lower-tail and strictly negative upper-tail curvature. This is a sufficiency result, not uniqueness; herding (Bouri et al., 2019), saturating adoption, and diminishing marginal liquidity could generate similar behavior. The model abstracts from macro factors, halving mechanics, and cross-asset correlations; tests distinguishing among competing mechanisms are left for future work.

16. Economic Interpretation

The results support a coherent economic narrative about Bitcoin's distributional evolution. Bitcoin's earliest price history was characterized by extreme reflexivity: small capital inflows could generate enormous percentage price changes when the asset base was tiny. As market capitalization has grown by orders of magnitude, this reflexivity has diminished. Progressively larger capital flows are required to produce comparable percentage moves, a pattern consistent with progressive compression of speculative peaks relative to structural support.

This mechanism operates asymmetrically. The lower tail (structural support) reflects Bitcoin's monetization trajectory, which evolves slowly relative to speculative-cycle frequency. The near-zero b^{LO} is consistent with structural support having risen approximately linearly in log-time, in line with the power-law characterization that Santostasi and others have documented. Liquidity dislocations can temporarily push prices below this trajectory (as documented in Section 12), but the trajectory itself is stable.

The upper tail (speculative peaks) exhibits a different dynamic. The significantly negative b^{HI} implies that each successive speculative peak, in log-price space, falls closer to the structural center than the prior peak. This is the statistical expression of what practitioners describe as Bitcoin's diminishing returns: percentage gains from cycle low to cycle high have been declining across cycles, from tens of thousands of percent in 2013, to thousands of percent in 2017, to progressively smaller relative amplitudes thereafter. The asymmetric curvature parameter b^{HI} places this pattern on a formal statistical footing: it quantifies the rate at which the upper-tail conditional quantile flattens in log-log space, and its block-bootstrap confidence interval ($[-0.455, -0.072]$) excludes zero, establishing that the diminishing-returns pattern is a statistically identified feature of the conditional distribution rather than a visual impression.

The implication for risk management is that models calibrated on early Bitcoin data, including both the OLS power law and the S2F model, implicitly embed assumptions about speculative reflexivity that no longer hold. As documented in Section 3, these models have been systematically too optimistic precisely because they project early-period upper-tail dynamics into a period of compressing speculative amplitude.

The asymmetric framework also offers a unified statistical perspective on diminishing returns. Practitioner discussions of Bitcoin's diminishing returns have historically focused on cycle-over-cycle ratios of cycle-low-to-cycle-high gains, drawdown depths, or rolling return windows. The b^{HI} estimate provides a complementary, fully distributional measure of the same phenomenon: rather than asking how much one cycle's peak return fell short of the prior cycle's, it asks how steeply the upper-tail conditional quantile bends inward as t grows. The two views are consistent (both diagnose a compressing upper tail) but the distributional view has the advantage that it pools information across all cycles into a single identified parameter with a confidence interval and a formal asymmetry test against the lower-tail null. Future work integrating cycle-ratio and quantile-curvature evidence may produce a sharper joint characterization of how Bitcoin's reflexivity attenuates as market capitalization grows.

17. Statistical Limitations

Several limitations of the present analysis should be noted.

First, the upper-tail curvature b^{HI} is not reliably identified on individual sub-periods (Section 10.1: profile-likelihood interval widths of ~ 27 units on the Mid sub-period and ~ 6 on Late, vs ~ 0.12 on the full sample). The full-sample identification rests on roughly four halving cycles; this is the deepest within-sample inference constraint. Section 10.3 mitigates it by showing $b^{\text{HI}} \in [-0.35, -0.29]$ across all 27 expanding windows with sufficient cycle coverage, but structural verification will require additional cycles. Structural-break tests (Bai and Perron, 1998) or time-varying specifications would be better suited to local curvature dynamics but require longer series.

A related sensitivity affects b^{LO} : across the Section 13 expanding windows, b^{LO} ranges from $+0.105$ (2022 cutpoint, training ends mid-acceleration) to -0.046 (2024 cutpoint, post-drawdown), more pronounced than b^{HI} variability. The full-sample $b^{\text{LO}} \approx 0$ is a long-run cross-cycle summary, not stable lower-tail linearity at any forecast horizon; prospective Q1% predictions are conditional on training-window cycle-phase composition.

Second, the tail-group partition is imposed rather than data-driven. Section 10.4 provides a data-consistency check; Section 10.4.1 reports a BIC-style comparison across all 877 partitions in which the imposed specification ranks fourth ($\Delta\text{BIC} = 5.26$ vs the minimizer, within Kass–Raftery's "positive but not strong" range). A data-driven clustering procedure with bootstrap-based selection uncertainty would provide a stronger characterization.

Third, the quadratic specification implies that as $t \rightarrow \infty$ the $b^{\text{HI}}x^2$ term dominates and predicted prices decline. This terminal behavior has no literal interpretation; the model is a finite-horizon distributional characterization, not a long-run structural model.

Fourth, block-bootstrap standard errors partially address autocorrelation but a complete treatment would require explicit residual modeling (e.g., GARCH or regime-switching innovations).

Fifth, the model is deliberately price-only. The expanding-window trajectory of b^{HI} (from -0.17 in 2021 to -0.35 by 2024; Section 13.2) may partly reflect the documented increase in Bitcoin’s correlation with global liquidity conditions (Alden, 2020–2025), a question left for future research.

Sixth, the analysis uses Bitcoin exclusively. Bitcoin is the only digital asset with enough history to span multiple cycles, and the structural component $F^s(t)$ in Section 15 specifically models monetization demand for fixed-supply assets; multi-asset validation across smart-contract platforms, utility tokens, or governance tokens would require asset-specific theoretical adaptations.

Seventh, the sample is unbalanced in macro and institutional coverage. The post-ETF window (US spot Bitcoin ETFs launched January 11, 2024) is roughly the final 15% of observations; most of the sample predates Bitcoin’s institutional integration and a prolonged recession with correlated liquidations. U.S. equity indices have declined materially in deep recessions (the S&P 500 fell $\sim 57\%$ in 2007–2009 and $\sim 49\%$ in 2000–2002); if Bitcoin exhibits similar macro sensitivity, dislocations exceeding the historical sample become plausible and the Q1% level should not be read as a floor under all macro scenarios.

Eighth, the first year of price history reflects an exchange-specific, thinly-traded market (Section 5); a more thorough exchange-by-exchange comparison is left for future research.

Ninth, partition-selection uncertainty is acknowledged but not propagated into the curvature standard errors. Section 10.4.1 reports the partition’s BIC-style rank (4 of 877; $\Delta\text{BIC} = 5.26$ vs the minimizer); a bootstrap-based selection procedure that propagates this uncertainty would provide a more complete characterization, which is left for future work.

18. Conclusion

This paper documents two main findings. First, three prominent prior Bitcoin price models, an OLS power law, the stock-to-flow (S2F) model, and its cross-asset extension (S2FX), have exhibited persistent optimistic out-of-sample bias since publication, with geometric mean price errors of $+32.1\%$, $+294.5\%$, and $+1,699\%$ respectively. Projecting constant-elasticity or scarcity-driven dynamics into a period of compressing speculative amplitude has been associated with upward forecast errors.

Second, we document a statistically significant asymmetry in Bitcoin’s conditional price distribution: $\Delta b = b^{\text{HI}} - b^{\text{LO}} = -0.302$ ($b^{\text{HI}} = -0.326$, $b^{\text{LO}} = -0.024$) is significant under both the full block bootstrap ($p = 0.012$) and the concentrated bootstrap ($p \leq 0.006$ across block lengths 14–90 days), with the 95% CI for Δb excluding zero. The lower-tail estimate is itself not distinguishable from zero ($p = 0.258$); this is consistent with, not acceptance of, lower-tail linearity. Out-of-sample, the asymmetric specification reduces upper-tail check-loss under expanding-window evaluation ($+26.7\%$ to $+29.1\%$ at Q75%, Q95%, Q99%, favored in every upper-tail $\times$ cutpoint cell), while lower-tail OOS performance is sensitive to training-window composition (Section 13). Quantile monotonicity is enforced through 2035 via rearrangement.

The framework extends rather than replaces the prior power-law research program: a linear lower-tail specification remains compatible with the historical lower conditional quantiles, and the contribution lies in identifying a parsimonious upper-tail correction that the linear specification misses. Section 15 develops a reduced-form structural-plus-reflexivity mechanism that is consistent with the observed asymmetry (a fundamental monetization trend augmented by a decaying speculative premium), offered as a sufficiency argument rather than a uniqueness claim, since alternative mechanisms (herding, saturating adoption, diminishing marginal liquidity) could generate similar qualitative behavior. The model is a distributional characterization, not a forecasting tool.

Four directions for future work: (i) integrating macro liquidity variables to explain the expanding-window trajectory of upper-tail curvature (Section 13.2); (ii) joint constrained quantile estimation as an alternative to post-estimation rearrangement; (iii) extending the methodology to other long-duration price processes where structural and speculative regimes plausibly differ (equity indices, commodities, emerging-market currencies during liberalization episodes); and (iv) cross-asset comparison to determine whether the asymmetry is generic or Bitcoin-specific.

References

- Alden, L. (2020–2025). Bitcoin and global liquidity cycles [Practitioner research, grey literature]. <https://www.lynalden.com/>
- Bai, J., & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1), 47–78.
- Baquero, C. (2026). Activity-warped power laws for Bitcoin price. Research Square preprint, posted February 10, 2026. <https://doi.org/10.21203/rs.3.rs-8845008/v1>
- Biais, B., Bisière, C., Bouvard, M., Casamatta, C., & Menkveld, A.J. (2023). Equilibrium Bitcoin pricing. *Journal of Finance*, 78(2), 967–1014. <https://doi.org/10.1111/jofi.13206>
- Bittel, J. (2021–2025). Research and commentary on Bitcoin liquidity dynamics [Practitioner research, grey literature]. <https://julienbittel.substack.com/>
- Bondell, H.D., Reich, B.J., & Wang, H. (2010). Noncrossing quantile regression curve estimation. *Biometrika*, 97(4), 825–838.
- Bouri, E., Gupta, R., & Roubaud, D. (2019). Herding behaviour in cryptocurrencies. *Finance Research Letters*, 29, 216–221. <https://doi.org/10.1016/j.frl.2018.07.008>
- Bouri, E., Saeed, T., Vo, X.V., & Roubaud, D. (2021). Quantile connectedness in the cryptocurrency market. *Journal of International Financial Markets, Institutions and Money*, 71, 101302. <https://doi.org/10.1016/j.intfin.2021.101302>
- Brunnermeier, M.K. (2009). Bubbles. In S.N. Durlauf & L.E. Blume (Eds.), *The New Palgrave Dictionary of Economics* (2nd ed.). Palgrave Macmillan. Provides a survey of rational

and near-rational bubble theory including conditions under which speculative premiums persist in equilibrium.

- Cannon, A.J. (2018). Non-crossing nonlinear regression quantiles by monotone composite quantile regression neural network, with application to rainfall extremes. *Stochastic Environmental Research and Risk Assessment*, 32(11), 3207–3225.
- Cheah, E.T., & Fry, J. (2015). Speculative bubbles in Bitcoin markets? An empirical investigation into the fundamental value of Bitcoin. *Economics Letters*, 130, 32–36.
- Chernozhukov, V., Fernández-Val, I., & Galichon, A. (2010). Quantile and probability curves without crossing. *Econometrica*, 78(3), 1093–1125. <https://doi.org/10.3982/ECTA7880>
- Cowen, B. (2020a). Bitcoin Letters 2020. Independent Research. <https://www.benjamincowen.com/reports/bitcoin-letters-2020>. Discusses the practitioner concept of fitting separate upper- and lower-envelope logarithmic regression bands to Bitcoin price history, with the lower band restricted to non-bubble observations; the two fitted bands exhibited visibly different curvature. Earliest video discussion of the upper/lower regression band framework: “Bitcoin: The path to \$1M using logarithmic regression,” posted April 15, 2020: <https://www.youtube.com/watch?v=J2FAHk6mWF8> (derives potential support and resistance zones using multiple logarithmic regression lines). Subsequent video, “Bitcoin logarithmic regression and under/over valuation of cryptocurrency marketcap,” posted June 3, 2020: https://www.youtube.com/watch?v=v7Ek__VFizE. Original visualization posted August 17, 2020: <https://x.com/benjamincowen/status/1295233606189633537>. This approach anticipates the intuition of lower-quantile regression and the asymmetric-curvature observation formalized in the present paper.
- Cowen, B. (2020b). Bitcoin logarithmic regression indicator [TradingView, grey literature]. <https://www.tradingview.com/v/wQWiGXRP/>
- Dette, H., & Volgushev, S. (2008). Non-crossing non-parametric estimates of quantile curves. *Journal of the Royal Statistical Society: Series B*, 70(3), 609–627.
- Diebold, F.X., & Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- Froot, K.A., & Obstfeld, M. (1991). Intrinsic bubbles: The case of stock prices. *American Economic Review*, 81(5), 1189–1214.
- Fulgur Ventures. (2024, May 26). Bitcoin Power Law Theory: Executive Summary [Practitioner research, grey literature]. Medium. <https://medium.com/@fulgur.ventures/bitcoin-power-law-theory-executive-summary-report-837e6f00347e>
- Kass, R.E., & Raftery, A.E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795. <https://doi.org/10.1080/01621459.1995.10476572>
- Koenker, R. (2005). *Quantile Regression*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511754098>

- Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1), 33–50.
<https://doi.org/10.2307/1913643>
- Kristoufek, L. (2015). What are the main drivers of the Bitcoin price? Evidence from wavelet coherence analysis. *PLOS ONE*, 10(4), e0123923.
<https://doi.org/10.1371/journal.pone.0123923>
- Liu, Y., & Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, 34(6), 2689–2727. <https://doi.org/10.1093/rfs/hhaa113>
- Liu, Y., Tsyvinski, A., & Wu, X. (2022). Common risk factors in cryptocurrency. *Journal of Finance*, 77(2), 1133–1177. <https://doi.org/10.1111/jofi.13119>
- Morillon, T.G., & Chacon, R.G. (2022). Dissecting the stock to flow model for Bitcoin. *Studies in Economics and Finance*, 39(3), 506–523. <https://doi.org/10.1108/SEF-10-2021-0409>
- Pal, R. (2020–2025). Global liquidity and digital assets [Practitioner research, grey literature]. Real Vision. <https://www.realvision.com/>
- Phillips, P.C.B., Shi, S., & Yu, J. (2015). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *International Economic Review*, 56(4), 1043–1078.
- Phillips, P.C.B., Wu, Y., & Yu, J. (2011). Explosive behavior in the 1990s Nasdaq: When did exuberance escalate asset values? *International Economic Review*, 52(1), 201–226.
- Plan C [pseudonym]. (2025a). Bitcoin Quantile Model v1 [Practitioner research, grey literature]. Posts at <https://x.com/TheRealPlanC>; updated v1 visualization (Jan 12, 2025): <https://x.com/TheRealPlanC/status/1878358716275392768>. Linear quantile regression in log-log space, anchored to January 3, 2009.
- Plan C [pseudonym]. (2025b). Bitcoin Quantile Model v2 [Practitioner research, grey literature]. Published October 15, 2025. <https://x.com/TheRealPlanC/status/1978449114179236198>. Piecewise stretched-exponential decay quantile regression with guaranteed non-crossing across the 1–99.9 percentile range; exact parameters not formally published. X content is not formally archived and may not persist.
- PlanB [pseudonym]. (2019). Modeling Bitcoin value with scarcity [Grey literature]. Medium. <https://medium.com/@100trillionUSD/modeling-bitcoins-value-with-scarcity-91fa0fc03e25>. Published parameters: $\log_{10}(P) = 3.4012 \cdot \log_{10}(SF) - 1.0456$.
- PlanB [pseudonym]. (2020). Bitcoin stock-to-flow cross asset model (S2FX) [Grey literature]. Medium. <https://medium.com/@100trillionUSD/bitcoin-stock-to-flow-cross-asset-model-50d260feed12>. Published parameters: $\ln(\text{mktcap}) = 12.7598 + 4.1167 \cdot \ln(SF)$.
- Santostasi, G. (2018–2024). Bitcoin Power Law Theory [Practitioner research, grey literature]. Original Reddit post (2018): https://www.reddit.com/r/Bitcoin/comments/9cqi0k/bitcoin_power_law_over_10_year_p

eriod_all_the_way/; Medium body of work (2019–2024):
<https://giovannisantostasi.medium.com/>

Shelton, A. (2024). Bitcoin return prediction: Is it possible via stock-to-flow, Metcalfe's Law, technical analysis, or market sentiment? *Journal of Risk and Financial Management*, 17(10), 443. <https://doi.org/10.3390/jrfm17100443>

Trolololo [pseudonym]. (2014). Logarithmic (non-linear) regression, Bitcoin estimated value [Grey literature]. BitcoinTalk forum, October 22, 2014.
<https://bitcointalk.org/index.php?topic=831547.0>. Published formula: $\log_{10}(P) = 2.9065 \cdot \ln(t) - 19.493$, where t is days since January 9, 2009.

Troster, V., Tiwari, A.K., Shahbaz, M., & Macedo, D.N. (2019). Bitcoin returns and risk: A general GARCH and GAS analysis. *Finance Research Letters*, 30, 187–193.
<https://doi.org/10.1016/j.frl.2018.09.014>