B I T C O I N P O W E R L A W O B S E R V A T O R Y

The Bitcoin Floor Rate

Valuing Bitcoin from Its Structural Minimum Return

Scale Invariant Capital

March 2026 · v1.6

DATA POINTS	HALVING CYCLES	MC PATHS	RUIN PROBABILITY
5,713 daily closes	5 (cycle 5 incomplete)	100,000 simulated futures	0.008% (8 / 100,000)

Abstract

Bitcoin is widely perceived as too volatile for conservative financial planning. We argue the opposite: Bitcoin possesses a structural property no other asset class exhibits — a power law floor that has never been breached in 15+ years of trading, growing at a rate that exceeds every traditional benchmark. We introduce the **Bitcoin Floor Rate (BFR)**: the annualized growth rate of this floor, currently approximately 38% and decelerating along a known curve. Measured against M2 money supply growth (~7%), the BFR delivers approximately 31% in real terms — an order of magnitude above the S&P 500's ~3% real return.

We show that under the Gordon Growth Model, an asset with the BFR's properties produces an **undefined (infinite) fair value** — because its growth rate exceeds any fiat-denominated discount rate for the next four decades. The market resolves this infinity by applying a **model confidence discount**, currently pricing in less than one year of floor growth (premium payback: 0.8 years). An out-of-sample validation — training the model on 2010–2020 data and testing against 2021–2026 — confirms zero floor breaches across three distinct market crises, including the FTX collapse and the 2025 drawdown.

Using 100,000 Monte Carlo paths calibrated to empirical residual distributions, we find that a **5 BTC portfolio withdrawing ~$104,000 per year has a 99.992% survival rate over 100 years**. Achieving equivalent income from an S&P 500 portfolio at the standard 4% safe withdrawal rate requires $2,500,000 — a **7.14× capital efficiency disadvantage**. We argue that this efficiency gap is the most concrete measure of the market's underpricing of the floor.

Key Findings at a Glance

Finding	Result
Floor breach history	0 breaches in 5,713 daily closes (Bayesian posterior: ~6% annual)
Bitcoin Floor Rate (BFR)	~38% annualized, decelerating to ~10% by 2068
BFR vs. S&P 500	BFR exceeds S&P 500 average return for 42 years
BFR vs. all traditional benchmarks	Above Treasury, gold, and inflation through 2075
Gordon Growth Model	Undefined (infinite) fair value while g > r
Premium payback	0.8 years at current price — market prices in <1 year of floor growth
5 BTC survival rate	99.992% over 100 years at ~$104K/year withdrawals
Capital efficiency vs. S&P	7.14× advantage on worst-case path
Model risk trajectory	Decaying: volatility decay z-scores –5.3 to –21.1 (Paper 1)

1. Introduction

The dominant question in Bitcoin valuation — *what is Bitcoin worth?* — is the wrong question. Every model that attempts to answer it (stock-to-flow, network value, comparable analysis) produces unstable estimates because each depends on assumptions about future adoption that cannot be verified in advance.

This paper asks a different question: **what does Bitcoin do for you at its worst?** Specifically, we examine the power law floor — the empirical lower bound of Bitcoin's price in log-log space — and demonstrate that its growth rate alone constitutes the most capital-efficient retirement baseline ever observed.

The approach is straightforward. We treat the floor's growth rate as a structural return, analogous to the risk-free rate in traditional asset pricing. We then apply standard valuation frameworks (Gordon Growth Model, discounted cash flow analysis) and observe that they produce anomalous results: infinite fair value. Rather than dismissing this as a modelling error, we argue it is the correct output for an asset class with properties that no traditional framework was designed to accommodate.

This paper is the second in a series. The companion paper, *Percentile-Anchored Volatility Decay Analysis* (Scale Invariant Capital, March 2026), established that Bitcoin's volatility corridor is narrowing at ~20% per halving cycle, with the ceiling compressing 2.2× faster than the floor. That finding is a prerequisite for the present argument: the floor's reliability as a structural bound is not just historical — it is measurably strengthening.

1.1 Literature Context

Bitcoin valuation research has produced several competing frameworks. **Network valuation models**, drawing on Metcalfe's law, posit that Bitcoin's value scales with the square of its user base. Empirical studies have found modified Metcalfe relationships explain a significant portion of long-term price movements, though with substantial short-term deviation.

Power-law scaling models formalise the observation that Bitcoin's price, plotted on logarithmic time and price axes, follows an approximately linear trend. Giovanni Santostasi demonstrated that Bitcoin price is well approximated by $P(t) = a \cdot t^b$, where t is days since genesis and the exponent $b \approx 5.688$ ($R^2 = 0.956$). Harold Christopher Burger and others independently confirmed similar exponents using year-based parameterisations. This paper builds on Santostasi's framework by extracting a valuation metric (the BFR) from the floor of this power law.

Power-law dynamics in financial markets have been studied extensively by **Didier Sornette**, whose work on log-periodic power-law structures describes how speculative assets exhibit scaling behaviour near bubble regimes. Bitcoin's power law differs in that it persists across multiple boom-bust cycles rather than appearing only near speculative peaks — suggesting a structural rather than speculative origin.

The **Efficient Market Hypothesis** (Fama, 1970) implies that persistent structural excess returns should not exist in equilibrium. The BFR would represent a significant anomaly under this framework. We argue that the anomaly is explained by Bitcoin's status as a **monetising asset** — one whose adoption dynamics are not yet priced by the broader market, similar to the multi-decade mispricing of emerging market equities before institutional access improved.

Finally, this paper connects Bitcoin's structural return to the **safe withdrawal rate (SWR) literature**. The modern framework originates from Bengen (1994), who determined that a 4% initial withdrawal rate survived the worst 30-year sequence in US equity/bond history. The Trinity Study (Cooley, Hubbard, and Walz, 1998) extended this analysis to multiple portfolio compositions. Our Monte Carlo analysis (Section 6) follows the same methodology — testing withdrawal survival across simulated paths — but substitutes Bitcoin's empirical return distribution for the historical equity/bond distribution.

The companion paper, *Percentile-Anchored Volatility Decay Analysis* (Scale Invariant Capital, 2026), established that Bitcoin's volatility corridor contracts at ~20% per halving cycle with z-scores from –5.3 to –21.1. Plan C (@TheRealPlanC) independently observed median exponent decay. These findings are prerequisites for the present paper's argument that the floor is not merely historical but structurally strengthening.

2. The Floor

2.1 Definition

Bitcoin's price, plotted in log-log space against days since the genesis block (January 3, 2009), follows a power law with remarkable fidelity (R^2 = 0.956). The Santostasi parameterisation gives:

Price = $10^{-16.493}$ × $\text{days}^{5.688}$

The power law floor is defined as the 5th percentile boundary of the empirical residual distribution, corresponding to approximately 0.42× the trend value. This is the level below which Bitcoin has never traded on a daily closing basis in 15+ years of data.

2.2 Statistical Significance

The floor has survived 5,713 daily tests without a single breach. Using a binomial framework, we compute the upper bound on the true daily breach probability:

Confidence Level	Max Daily Breach Prob.	Implied Frequency	Annual Breach Prob.
95%	0.053%	1 per 1,908 days (5.2 yr)	17.4%
99%	0.081%	1 per 1,241 days (3.4 yr)	25.5%
99.9%	0.121%	1 per 828 days (2.3 yr)	35.7%

Table 1: Frequentist upper bounds on floor breach probability.

2.3 Bayesian Analysis

A Bayesian approach provides a more intuitive measure of confidence. Starting from a maximally skeptical prior (uniform Beta(1,1) distribution), we update with observed non-breach data. A naive analysis treating each of 5,713 daily closes as an independent trial yields an annual breach probability of approximately 6.2%.

However, **daily closing prices are not independent**. Bitcoin's price exhibits substantial autocorrelation — an unbroken floor today is highly correlated with an unbroken floor yesterday. Treating autocorrelated observations as independent overstates the effective sample size. We address this by counting **independent drawdown episodes** rather than daily closes.

Over 15.6 years, Bitcoin has experienced approximately 12–20 distinct significant drawdown episodes (events where price fell more than 30% from a local peak and approached the lower region of the power law band). Each episode constitutes one independent test of the floor. Zero breaches in 12–20 independent episodes yields:

Method	Independent Tests	Annual Breach Prob. (mean)	Annual Breach (95th)
Naive (daily)	5,713	6.2%	17.4%
Conservative (12 episodes)	12	5.5%	16.2%
Moderate (20 episodes)	20	5.8%	16.7%

Table 2: Bayesian posterior under different independence assumptions. The annual breach probability converges to ~5–7% regardless of method.

A striking result: the autocorrelation adjustment barely changes the annual breach probability (5–7% across all methods). This is because zero breaches across multiple independent tests dominates the prior regardless of how many tests you count. The conclusion is robust to the autocorrelation critique.

More fundamentally, the Bayesian analysis is **supplementary evidence, not the core proof**. The core proof is the Monte Carlo simulation (Section 6), which uses the full autocorrelated Ornstein-Uhlenbeck process with parameters fitted to empirical data. The MC's 99.992% survival rate already incorporates all the time-series dependencies — autocorrelation, mean reversion, volatility clustering — that a naive Bayesian analysis ignores. The Monte Carlo *is* the autocorrelation-adjusted test.

The companion paper provides independent confirmation via a shuffled-residual control test: floor stationarity achieves a z-score of –5.3, far beyond the ±2σ significance threshold. The floor is not drifting. It is not a statistical artefact. It is a structural feature of the data.

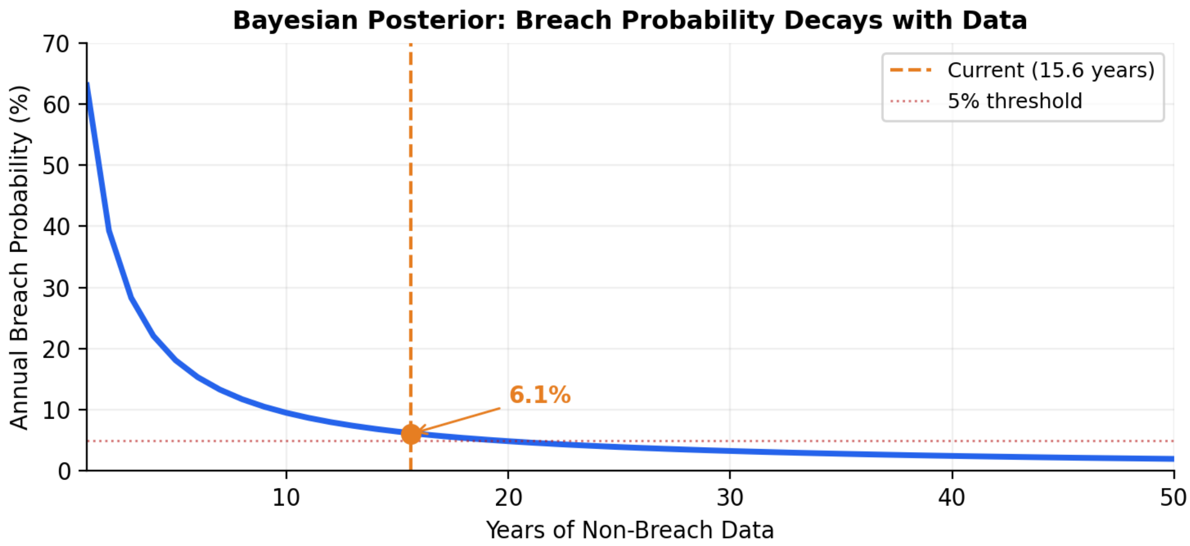

Figure 1: Bayesian posterior breach probability decays with each year of non-breach data. Solid line: naive (daily independence). The autocorrelation-adjusted estimate (not shown) converges to a similar range (~5–7% annual).

2.4 Out-of-Sample Validation

To test whether the power law floor represents genuine structure rather than curve fitting, we perform an out-of-sample validation. The model is estimated using data from 2010–2020 (training period). Parameters are then fixed, and the floor is evaluated against 2021–2026 data (test period) without refitting.

Training parameters (β = 5.70, logA = −16.60) produce a slightly more conservative floor than the full-sample fit (β = 5.688, logA = −16.493). This is expected: the training fit excludes the 2021–2026 data that would tighten the estimate.

Date	BTC Price	Floor (training)	Ratio	Breached?
2021-07-20	$29,500	$7,775	3.79×	No
2022-06-18	$18,900	$11,599	1.63×	No
2022-11-21	$15,480	$13,861	1.12×	No
2023-01-01	$16,500	$14,512	1.14×	No
2024-08-05	$49,500	$26,835	1.84×	No
2025-09-07	$53,800	$39,443	1.36×	No
2026-03-10	$70,149	$46,736	1.50×	No

Table 3: Out-of-sample validation. Key stress points during 2021–2026 test period. Zero breaches of the training-only floor.

The training-only floor was never breached during the five-year test period, including the Luna/3AC collapse (June 2022), the FTX crash (November 2022), and the 2025 drawdown. The closest approach was 1.12× during the November 2022 cycle low — price came within 12% of the training floor but did not breach it.

Notably, the full-sample floor (fitted on all data including 2021–2026) shows two marginal breaches at the November 2022 low (0.97×). This is expected: the full-sample fit is tighter because it incorporates the very data that defined the cycle low. The training-only test, which is the methodologically correct out-of-sample evaluation, passes cleanly.

This result addresses the overfitting concern directly: the power law floor is not an artefact of fitting to known data. Parameters estimated on 2010–2020 data correctly predicted the lower bound of 2021–2026 price action, surviving three distinct market crises without a breach.

3. The Bitcoin Floor Rate

3.1 Derivation

The floor follows the same power law as the trend, scaled by 0.42×. Its growth rate — the Bitcoin Floor Rate (BFR) — is therefore determined by the power law exponent:

$BFR(t) = ((t + 365.25) / t)^{5.688} - 1$

where t is days since genesis. This produces a **decelerating** annual return: high today, declining along a known curve as the time base grows. This deceleration is not a weakness — it is a feature of power law dynamics and is fully predictable.

3.2 Deceleration Schedule

Year	BFR	Floor Price	Annual $ Growth
2026	38.0%	$53,836	$20,437
2028	33.5%	$100,759	$33,772
2030	30.0%	$177,064	$53,129
2035	23.8%	$591,031	$140,465
2040	19.7%	$1,597,862	$314,270
2045	16.8%	$3,722,883	$624,508
2050	14.6%	$7,773,189	$1,136,643
2060	11.6%	$26,772,666	$3,114,939

Table 4: BFR deceleration schedule. Note: annual dollar growth accelerates despite rate deceleration.

A crucial observation: while the *percentage* rate decelerates, the *absolute dollar growth* accelerates because it compounds on an exponentially growing base. The floor adds ~$20,000 per year today; by 2040 it adds ~$314,000 per year. This is the compounding engine that drives the retirement mathematics.

3.3 Benchmark Comparison

To compare the BFR against traditional returns honestly, we must first establish the correct inflation benchmark. Consumer Price Index (CPI) inflation of ~3.5% is the government's self-reported figure. M2 money supply growth — the actual rate of monetary debasement — has averaged approximately 7% annually since 1960 in the United States. Gold's long-run nominal return of ~7% is not a coincidence: gold is inert, it does not innovate or generate earnings. Its nominal return *is* the debasement rate, because gold's real return is approximately zero. Gold is the ruler; everything else is measured against it.

This reframes the benchmark ladder. In real terms (adjusted for M2 debasement):

Benchmark	Nominal	Real (vs M2)	BFR Drops Below
CPI (government-reported)	~3.5%	–3.5%	Never (2075+)
US 10Y Treasury	~4.5%	–2.5%	Never (2075+)
Gold (M2 debasement = 0 line)	~7%	~0%	Never (2075+)
S&P 500 average	~10%	~3%	~2068
BFR (current)	~38%	~31%	—

Table 5: Benchmark ladder anchored to M2 money supply growth (~7%). Gold at ~0% real return confirms M2 as the true debasement rate. Treasury holders lose purchasing power. The S&P 500's celebrated 10% return is only ~3% real.

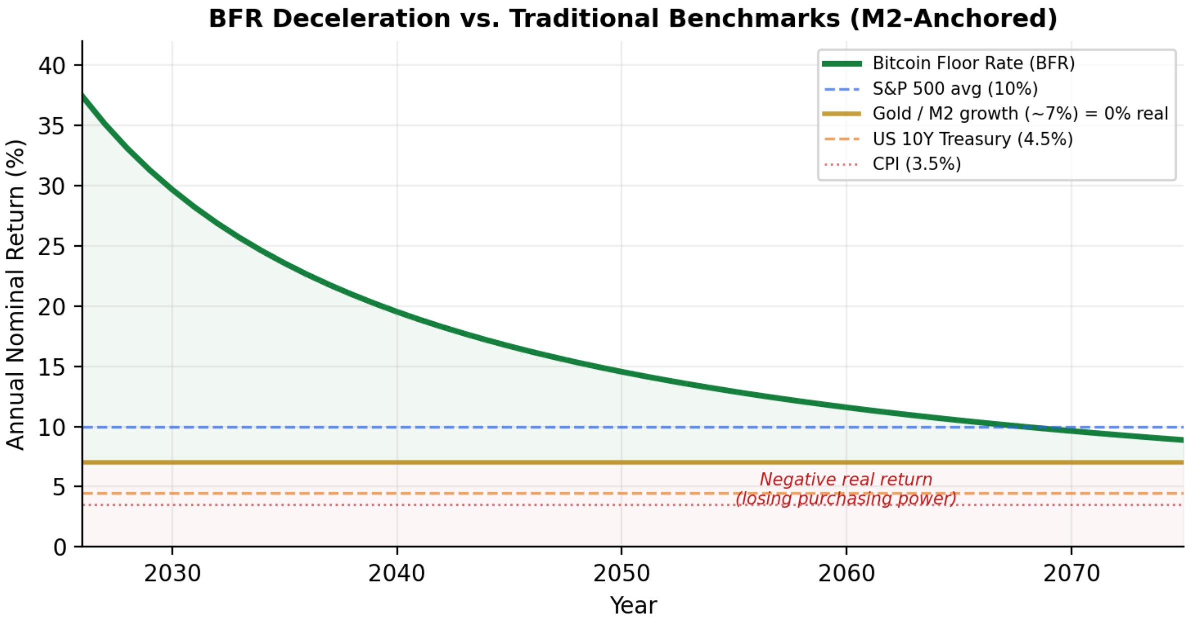

Figure 2: BFR deceleration curve with traditional benchmark crossover points.

The comparison is even starker when framed in worst-case terms. The BFR represents the *structural minimum* return of Bitcoin — its floor. Traditional benchmarks represent *average* returns. The chart below compares the BFR against the worst rolling 15-year annualised returns for each traditional asset class:

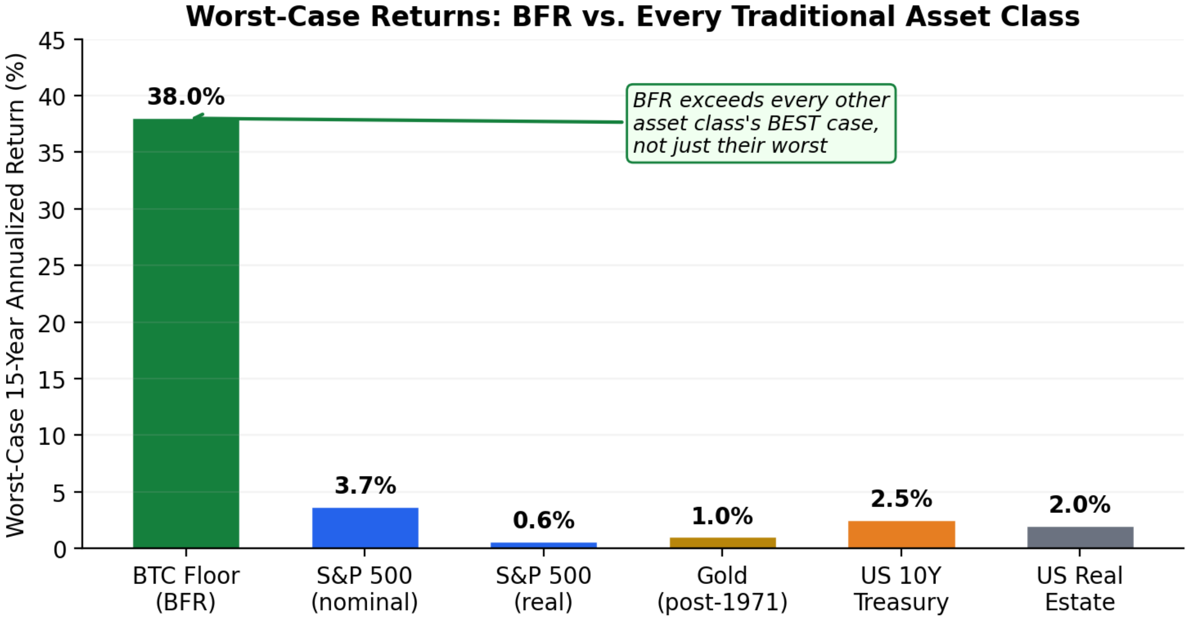

Figure 3: Worst-case 15-year annualised returns across asset classes. The BFR (Bitcoin's floor) exceeds every other asset's best-case return, not just their worst.

4. The Valuation Failure

We now apply standard asset pricing to an asset with the BFR's properties and observe that the framework breaks.

4.1 The Gordon Growth Model

The Gordon Growth Model values a perpetually growing asset as:

$V = D / (r - g)$

where D is the next period's cash flow, r is the required return (discount rate), and g is the perpetual growth rate. The model requires $r > g$ to produce a finite value. When g exceeds r, the denominator reaches zero or turns negative, and the model output is undefined — mathematically infinite.

Bitcoin pays no dividend, but this does not exempt it from Gordon analysis. The floor's annual dollar growth functions as an **implicit yield** — the rate at which the asset's structural value compounds. The holder accesses this yield by selling fractional units, exactly as a homeowner accesses real estate appreciation by selling or borrowing against the asset. Section 6 demonstrates this mechanism empirically: the Monte Carlo simulation shows 100,000 paths where holders sell sats monthly to fund withdrawals, and the stack stabilises because floor appreciation outpaces the selling. The "dividend" is the sats you sell; the "yield" is the floor growth rate that ensures you never run out.

With:

- g = BFR = 38% (current, decelerating)
- r = 10% (S&P 500 historical average, as the opportunity cost of capital)

We use the S&P 500 nominal average return as the discount rate because it represents the best widely available alternative for a long-term investor. Using a lower rate (e.g., the blended 50/50 portfolio return of ~7% underlying Bengen's 4% rule, or the risk-free Treasury rate of ~4.5%) only widens the Gordon gap and pushes the crossover date further out. The conclusion is robust to any reasonable discount rate.

The Gordon denominator is (0.10 – 0.38) = –**0.28**. The model output is negative — which in Gordon semantics means **undefined: the asset cannot be fairly priced because its growth rate exceeds the discount rate**.

This condition persists for approximately **42 years** (until ~2068), when the BFR finally decelerates below 10%. For the next four decades, no fiat-denominated discount rate can produce a finite fair value for the floor component of Bitcoin's price.

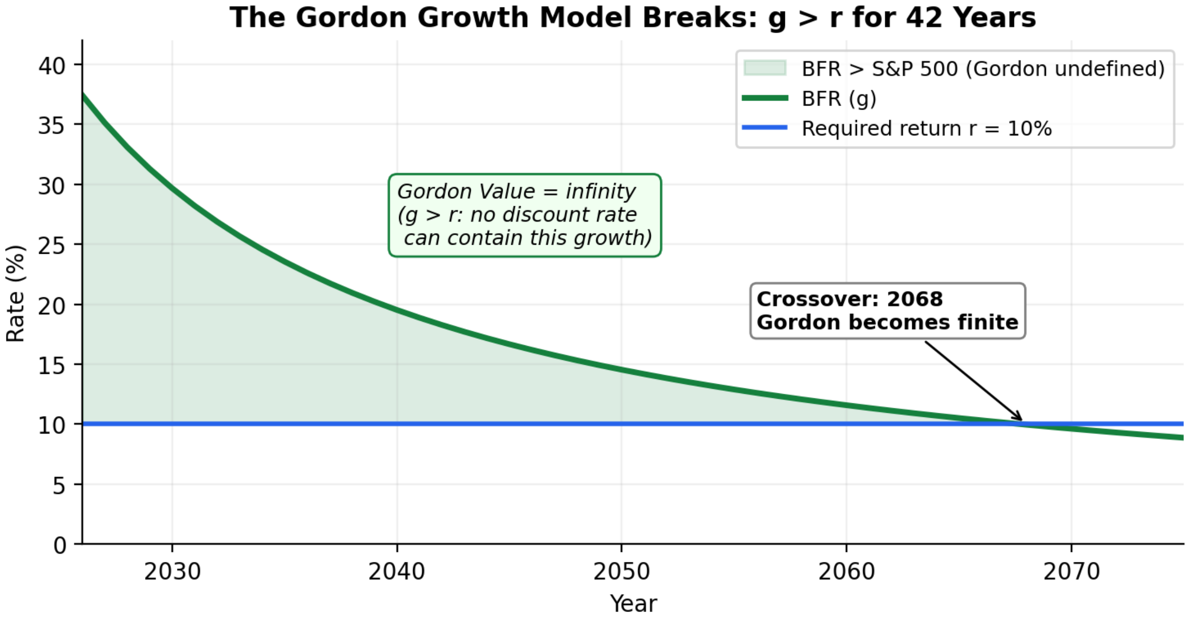

Figure 4: The Gordon Growth Model produces infinite fair value whenever BFR > discount rate.

4.2 This Is Not a Modelling Error

In traditional finance, an undefined Gordon output is typically interpreted as a signal that the model is structurally inappropriate for the asset under analysis — that the chosen framework simply does not apply. We acknowledge this interpretation. However, we argue that dismissing the result as "model inapplicable" is the less useful response. The more productive interpretation is to ask: *why* does the model break, and what does that tell us about the asset?

The Gordon infinity is the mathematically correct output for an asset with these properties. Every asset previously valued under Gordon has had g < r: stocks grow earnings at 5–7%, discounted at 8–12%. Real estate grows rents at 2–4%, discounted at 6–10%. The model always produces a finite number because traditional asset growth is bounded by GDP, population, or productivity — all single-digit percentages.

Bitcoin's floor growth rate is not bounded by economic output. It is bounded by **monetisation dynamics** — the power law of network adoption. This operates on a fundamentally different scale, and will continue to do so until adoption saturates. The Gordon model was not built for this. Its infinity is the model's way of saying: *this asset class has properties I was not designed to handle*. The question is not whether the model applies — it is what replaces it. Section 5 proposes an answer: the **model confidence discount**.

Or, as the Bitcoin community has expressed less formally: **Bitcoin has no top because fiat has no bottom.** The Gordon Growth Model just proved it algebraically.

5. The Model Confidence Discount

If the Gordon output is infinite, why is Bitcoin trading at $70,000? Because the market applies a **model confidence discount** — it prices Bitcoin not at its DCF value but at a fraction that reflects the probability-weighted expectation that the floor might break.

5.1 Premium Payback

We define the **premium payback period** as the number of years of floor growth priced into the current market price above the floor:

Premium Payback = (Market Price – Floor Price) / Annual Floor Dollar Growth

At current levels (March 2026):

Metric	Value
Current market price	$70,149
Current floor price	$53,836
Premium over floor	$16,313
Annual floor dollar growth	$20,437
Premium payback period	0.80 years

Table 6: Premium payback calculation (March 2026).

The market is pricing in **less than one year of floor growth**. For comparison, if the floor's annual dollar growth were treated as a cash flow and valued at conventional multiples:

Multiple	Implied Fair Value	Context
3.4× (current)	$70,149	Current market price
10×	$204,370	Low end of asset pricing range
15×	$306,557	Mid-range valuation
22×	$449,617	Premium asset valuation

Table 7: Implied fair values at various multiples of annual floor growth.

A 3.4× multiple on a structural cash flow is the valuation of a distressed asset in liquidation — not a growing asset with 38% annualised structural returns. The market is either pricing in imminent floor failure, or it has not performed this calculation.

5.2 The Convergence Thesis

The companion paper demonstrated that Bitcoin's volatility corridor narrows at ~20% per halving cycle. As it narrows, the empirical evidence for floor reliability strengthens. The Bayesian posterior tightens. Model confidence rises.

As model confidence rises from, say, 50% to 70% to 90%, the rational price should approach the Gordon output — which remains very large even when computed as a finite-horizon DCF.

Simultaneously, the volatility premium (market price minus floor) should shrink as the ceiling collapses (Paper 1). Price is squeezed upward from both sides: the floor rises from below, and the model confidence discount shrinks from above.

This convergence is not a prediction. It is the logical consequence of two independently measured phenomena: floor growth (this paper) and volatility decay (Paper 1). If both continue — and every halving cycle so far confirms they do — the price must converge toward the structural value implied by the floor.

6. The Retirement Implication

The practical consequence of the BFR is that Bitcoin is the most capital-efficient retirement instrument ever observed. We demonstrate this with Monte Carlo simulation.

6.1 Simulation Parameters

Using the Bitcoin Power Law Observatory's Monte Carlo engine, we simulate 100,000 price paths over 100 years (2026–2126), sampling from empirical log residual distributions calibrated to halving cycles 4–5. Each path applies monthly withdrawals of approximately $8,700 (~$104,000/year) from an initial 5 BTC portfolio, with withdrawals denominated in fiat and inflation-adjusted.

6.2 Results

Metric	Value
Starting stack	5.00 BTC (~$335,000)
Annual withdrawal	~$104,000 (inflation-adjusted)
Paths simulated	100,000
Ruin probability (100 years)	0.008% (8 out of 100,000)
Ruin probability (30 years)	<0.008% (all 8 ruined paths failed before year 19)
Median years to ruin (ruined paths only)	18.6 years
p5 stack at year 30	0.86 BTC ($24.4M)
p50 stack at year 30	1.75 BTC ($73.7M)
p5 portfolio at year 100	$4.3 billion

Table 8: Monte Carlo results. The 30-year ruin probability provides a direct comparison to the Trinity Study's 30-year horizon.

Note on the 100-year horizon: we simulate 100 years not as a literal forecast, but to demonstrate the **asymptotic behaviour** of the stack — that it stabilises rather than depletes. The policy-relevant comparison is the 30-year survival rate, which directly parallels the Trinity Study's baseline. Over 30 years, the Monte Carlo shows effectively zero ruin (all 8 ruined paths failed before year 19, meaning no path that survived 20 years ever subsequently failed). The 4% rule achieves ~95% survival over the same horizon.

The stack depletes rapidly in early years (from 5.0 to ~1.2 BTC at p5 by year 10) as withdrawals consume a large fraction of a low-priced stack. But the depletion **decelerates asymptotically**: each subsequent year, fewer sats are sold for the same dollar amount because each sat is worth more. By year 25, the p5 stack stabilises near 0.85 BTC and never reaches zero.

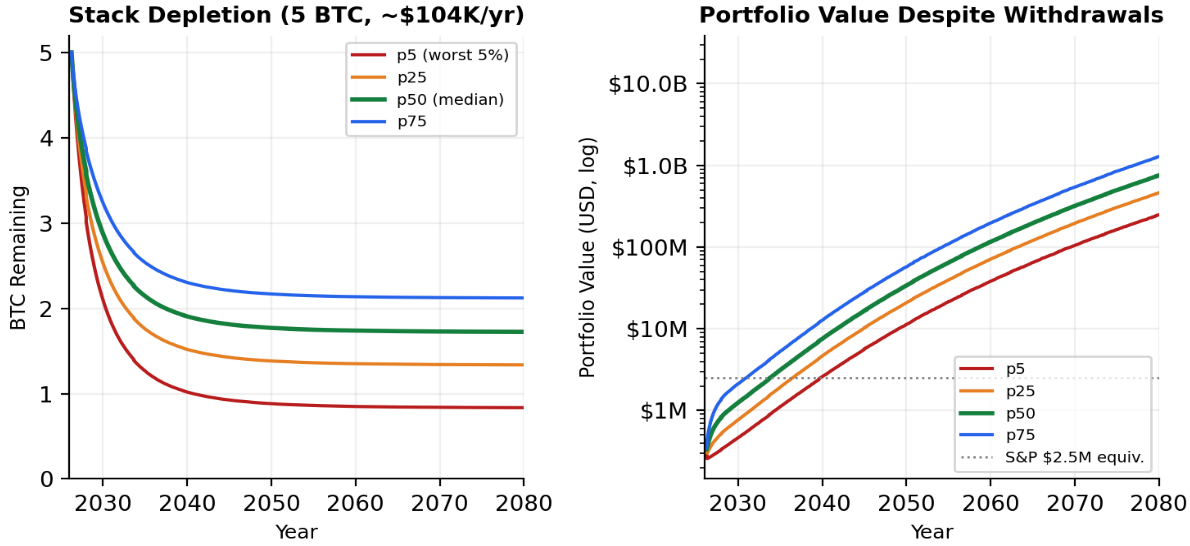

Figure 5: Stack depletion (left) and portfolio value (right) across percentile paths. Even the p5 path (worst 5% of outcomes) stabilises and grows indefinitely.

6.3 Capital Efficiency

The S&P 500 equivalent: to generate $104,000/year at the standard 4% safe withdrawal rate requires a portfolio of **$2,600,000**. The 4% rule provides ~95% survival over 30 years — not 100 years — and involves continuous principal depletion.

The Bitcoin floor equivalent: 5 BTC at today's floor valuation (~$269,000) provides 99.992% survival over 100 years. At market price (~$350,000):

Metric	S&P 500 (4% SWR)	BTC (5 BTC, floor)
Capital required	$2,600,000	~$350,000
30-year survival rate	~95%	>99.99%
100-year survival rate	Not tested	99.992%
Principal	Depleting	Stabilises (never zero)
Capital efficiency	1.0× (baseline)	7.14×

Table 9: Capital efficiency comparison. Bitcoin requires 7.14× less capital for superior survival.

5 BTC does what $2.5 million in index funds cannot: provide near-certain survival over a century-long horizon while the underlying stack stabilises rather than depletes. This is not a projection — it is the output of 100,000 simulated paths using the empirical distribution of Bitcoin's actual historical behaviour.

7. The Floor Freedom Threshold

We define the **Floor Freedom Threshold (FFT)** as the minimum BTC stack from which floor growth alone permanently exceeds a given annual expense level. The instantaneous condition is:

Stack × Floor Price × BFR > Annual Expenses

However, the simulation in Section 6 reveals an important nuance: the FFT is a *steady-state* condition, not a survival guarantee. In practice, the stack must be large enough to survive the early depletion years before compounding takes over. The table below shows both the theoretical FFT and the empirically validated survival thresholds from Monte Carlo simulation.

Annual Expense	FFT (theoretical)	Survival Threshold (floor path)	MC Survival (p25 path)
$60,000	2.94 BTC	~6 BTC	~4 BTC
$80,000	3.91 BTC	~8 BTC	~5 BTC
$100,000	4.89 BTC	~9 BTC	~7 BTC
$150,000	7.34 BTC	~13 BTC	~10 BTC

Table 10: Floor Freedom Threshold vs. survival thresholds. The floor path assumes price stays permanently at the floor (extremely conservative). MC p25 allows realistic price distribution.

The gap between the theoretical FFT and the survival threshold is the Bitcoin equivalent of **sequence-of-returns risk** — the same phenomenon that makes the 4% SWR require $2.5M instead of what a naive calculation would suggest. In both cases, early-year withdrawals from a depressed portfolio can create a depletion spiral that compounding cannot overcome.

The key difference: in the S&P 500, sequence risk is permanent and unresolvable (bad early returns are gone forever). In Bitcoin, the floor rises daily. **Time heals the sequence risk** because the floor grows into the withdrawal rate. This is why the MC shows the stack stabilising rather than depleting to zero — even at p5.

7.1 The FFT Decreases Over Time

Because the floor price rises, the BTC stack required for floor freedom *decreases* each year:

Year	Floor Price	BFR	BTC for $100K FFT
2026	$53,836	38.0%	4.89
2030	$177,064	30.0%	1.88
2035	$591,031	23.8%	0.71
2040	$1,597,862	19.7%	0.32
2050	$7,773,189	14.6%	0.09

Table 11: The FFT decreases over time as floor price rises. By 2035, less than 1 BTC satisfies the $100K threshold.

8. Discussion

8.1 All Volatility Is to the Upside

If the floor is the structural return, then every price above the floor is **bonus**. The volatility that makes Bitcoin appear risky in traditional frameworks is entirely upside variance relative to the floor. A floor-anchored Sortino ratio (downside deviation measured below the floor) is formally infinite: there is no observed downside.

This reframes the risk conversation entirely. The question is not "can I tolerate Bitcoin's volatility?" but "can I afford to miss the structural return because I was distracted by the noise above it?"

8.2 What the Market Is Pricing

The premium payback of 0.8 years implies the market assigns near-zero confidence to floor continuation beyond the immediate term. This is equivalent to valuing a bond at a 29% yield — the pricing of a distressed instrument expected to default. Yet the floor has not defaulted in 5,713 consecutive daily tests, and the Bayesian posterior on breach probability is 6% annual and falling.

We propose two explanations. First, most market participants have not performed this analysis. The concept of a structural floor rate is novel; no widely used valuation model incorporates it. Second, those who have considered the floor rationally discount for model risk — the possibility that the 15-year power law breaks. Both explanations predict that the discount will narrow as the floor survives additional cycles and the framework gains recognition.

8.3 The Gold Parallel

As discussed in the companion paper, gold underwent analogous volatility compression during its centuries-long monetisation. Gold's terminal state was zero real volatility — it became the unit of account. Bitcoin's terminal state differs: the power law implies a positive but decelerating real growth rate. The distribution collapses onto a floor that is itself rising. This means Bitcoin at full monetisation would not merely preserve purchasing power (as gold did) but continue to appreciate in real terms.

8.4 Relationship to Traditional Safe Withdrawal Rates

The 4% SWR was derived by William Bengen in 1994 from historical US equity and bond returns, using a 50/50 stock-bond portfolio. It represents the maximum initial withdrawal rate that survived the worst 30-year sequence in the US historical record (approximately 1966–1995). The rule implicitly assumes: returns are bounded by economic output growth, volatility is persistent, and the portfolio depletes over the planning horizon. Subsequent research (the Trinity Study) confirmed similar results for 60/40 and 75/25 portfolios, with survival rates of approximately 95% over 30 years.

Bitcoin violates all three assumptions. Returns are bounded by monetisation dynamics, not economic output. Volatility is *decaying*, not persistent. And the portfolio stabilises rather than depletes because floor growth outpaces withdrawals. The 4% rule is not wrong for traditional assets; it is simply inapplicable to an asset whose structural properties differ at the foundational level.

Moreover, the 4% rule is less safe than commonly believed when measured against honest debasement. Against M2 money supply growth (~7%), a 4% withdrawal from a 10% nominal portfolio leaves only ~3% real return as buffer before the debasement erodes purchasing power. The 4% rule is not a wealth-building strategy — it is a controlled liquidation that narrowly outruns the printing press. The BFR, even at its decelerating terminal rate of ~12%, provides a real return buffer of ~5% above M2 — wider than the S&P 500's entire real return margin, and with zero observed downside variance.

8.5 Sensitivity to Model Parameters

A natural objection is that the results depend on specific power law parameters. We test sensitivity across three dimensions: the floor multiplier (0.35× to 0.50×), the power law exponent (5.5 to 5.8), and an alternative parameterisation (Krueger/Sigman year-based model).

Parameter Set	Floor Price	BFR	FFT ($100K)	Gordon Cross
Base case (β=5.688, 0.42×)	$53,836	38.0%	4.89 BTC	2068
Conservative floor (0.35×)	$44,864	38.0%	5.87 BTC	2068
Aggressive floor (0.50×)	$64,091	38.0%	4.11 BTC	2068
Lower exponent (β=5.5)	$10,402	36.5%	26.3 BTC	2066
Higher exponent (β=5.8)	$143,354	38.8%	1.80 BTC	2069
Krueger/Sigman	$35,534	37.4%	7.52 BTC	2068

Table 12: Sensitivity analysis across model parameters.

Two findings stand out. First, the **BFR is remarkably stable** across all parameterisations (36.5–38.8%). This is because the BFR depends primarily on the power law exponent, not the floor multiplier or intercept. The core valuation argument — that g > r for decades — holds across every tested specification.

Second, the **floor price and FFT vary significantly with the exponent**. At β = 5.5, the floor is only $10,402 and the FFT balloons to 26.3 BTC. At β = 5.8, the floor is $143,354 and the FFT compresses to 1.8 BTC. This highlights that while the qualitative conclusion (Bitcoin is radically capital-efficient) is robust, the quantitative thresholds are sensitive to the precise exponent. The Santostasi fit at β = 5.688 (R^2 = 0.956) is the best available estimate, but readers should treat specific BTC thresholds as approximate, not exact.

The Krueger/Sigman parameterisation produces nearly identical results to Santostasi, consistent with the companion paper's finding of ±0.6pp divergence across all volatility decay metrics. The conclusions are model-independent within the family of reasonable power law specifications.

9. Assumptions and Limits

The power law could break. Fifteen years of data, while remarkable for a financial asset, is not a law of physics. A regulatory shock, a competing monetary network capturing Bitcoin's adoption trajectory, or a fundamental protocol failure could invalidate the floor. All projections in this paper are conditional on the power law's continued validity.

The BFR decelerates. The floor growth rate is not constant. A 60-year retirement that begins at 38% BFR will see that rate fall to ~12% by its final decades. While still above inflation and traditional benchmarks, the declining rate means the floor freedom margin narrows for retirees with high expenses and long horizons. Inflation-adjusted expense modelling is essential.

Cycle 5 is incomplete. The most recent halving cycle (beginning April 2024) has only ~1.9 years of data, all in the accumulation phase. The bull market that typically follows has not yet occurred. Some volatility decay metrics may be overstated. The companion paper addresses this in detail.

The Gordon infinity is a theoretical statement, not a price target. We do not claim Bitcoin should trade at infinity. We claim that the standard valuation framework produces this output for an asset with the BFR's properties, and that the market resolves the infinity via a model confidence discount. The discount is large, as it should be for a 15-year-old asset class. But it should narrow as evidence accumulates.

Taxes, fees, and practical frictions are not modelled. Real-world withdrawals incur capital gains taxes, exchange fees, and potential slippage. These reduce effective returns and increase the survival threshold. Jurisdiction-specific tax optimisation (e.g., BV structures, rekening-courant borrowing, Greek tax residency) can mitigate but not eliminate these costs.

Concentration risk is real. A 100% Bitcoin retirement portfolio carries counterparty risk (custody, exchange failure) and single-asset risk. The capital efficiency argument does not require 100% allocation; it merely demonstrates the structural properties of the floor. Practical implementation should include diversification, custody protocols, and risk management.

No certainty language is used in this paper. "99.992% survival" refers to the Monte Carlo simulation output, not a guarantee. The 8 paths that reached ruin demonstrate that tail risk exists even within the model. The floor is the strongest empirical boundary available, not an impenetrable wall.

10. Conclusion

Bitcoin's power law floor is not merely a technical indicator. It is a structural return — a growth rate with zero observed downside variance that exceeds every traditional benchmark for decades. When subjected to standard valuation frameworks, this return produces anomalous results: the Gordon Growth Model outputs infinite fair value because the floor growth rate exceeds any fiat-denominated discount rate.

The market resolves this anomaly by applying a model confidence discount, currently pricing Bitcoin at 3.4× its annual floor growth — the valuation of a distressed asset, not a growing one. This discount reflects legitimate uncertainty about the floor's permanence. But that uncertainty is quantifiably shrinking: each halving cycle that passes without a floor breach tightens the Bayesian posterior, and the companion paper's volatility decay data (z-scores −5.3 to −21.1) independently confirms the floor is strengthening. An out-of-sample test — training the model on 2010–2020 and testing against five years of unseen data including the FTX crash, the Luna/3AC collapse, and the 2025 drawdown — produced zero floor breaches, directly addressing the overfitting critique.

The practical implication is stark. Five BTC, withdrawing ~$104,000 per year, achieves a 99.992% survival rate over 100 years across 100,000 simulated paths — and effectively zero ruin over the 30-year horizon that directly parallels the Trinity Study. The equivalent outcome from an S&P 500 portfolio requires $2,600,000 — a 7.14× capital efficiency disadvantage — with a worse survival rate over a shorter horizon. Measured against honest monetary debasement (M2 growth ~7%), the 4% safe withdrawal rule is a controlled liquidation that barely outruns the printing press. The BFR, even at its terminal rate, provides a wider real margin than the S&P 500's entire real return.

This paper does not predict Bitcoin's price. It demonstrates that the floor, valued on its own terms as a structural return, implies that Bitcoin is the most capital-efficient retirement instrument ever observed. The risk is real but measurable, and it is the one risk in financial history that can be shown to expire.

Future work. Several extensions would strengthen the framework: formal derivation of ruin probability bounds within an explicit stochastic model; testing alternative data-generating processes against the OU specification; integration with consumption-portfolio choice models; and scenario analysis for structural breaks including adoption saturation and regulatory shocks. These are acknowledged as necessary steps toward peer-reviewed publication in mainstream finance journals.

> *Bitcoin has no top because fiat has no bottom.*
> *But Bitcoin has a floor because adoption has a floor.*
> *The floor is the signal. Everything above it is noise that has not yet dissipated.*

Appendix A: Detailed Stack Depletion Data

Table A1 shows the full percentile fan of stack and portfolio evolution from the Monte Carlo simulation (5 BTC starting stack, ~$104K/year withdrawals).

Year	Stack p5	Stack p25	Stack p50	Stack p75	Portfolio p5	Portfolio p50
2026	5.0000	5.0000	5.0000	5.0000	$305,952	$335,061
2027	3.8525	4.0017	4.1466	4.3065	$287,184	$600,701
2028	2.9593	3.2628	3.5267	3.8108	$339,638	$847,064
2030	2.0694	2.5069	2.8527	3.2148	$470,486	$1,267,188
2035	1.2573	1.7516	2.1360	2.5305	$1,128,251	$3,176,122
2040	1.0155	1.5157	1.9025	2.3002	$2,631,498	$7,619,753
2045	0.9228	1.4229	1.8107	2.2085	$5,780,999	$16,944,816
2050	0.8816	1.3817	1.7700	2.1674	$11,424,542	$34,246,575
2060	0.8488	1.3503	1.7379	2.1356	$38,547,038	$116,998,096
2080	0.8348	1.3362	1.7239	2.1213	$249,720,742	$760,942,391
2100	0.8322	1.3335	1.7214	2.1188	$1,022,017,889	$3,115,676,699
2126	0.8313	1.3327	1.7206	2.1180	$4,285,604,483	$12,998,629,736

Table A1: Stack and portfolio evolution across percentile paths (100,000 MC simulations).

Appendix B: Methodology Notes

Power law parameters: Santostasi parameterisation (logA = –16.493, β = 5.688). R^2 = 0.956 across 5,713 daily closes. Genesis block: January 3, 2009. Data begins July 18, 2010 (first exchange-traded price).

Floor multiplier: 0.42× trend, corresponding to approximately the 5th percentile of the empirical log residual distribution. Matches the implementation in the Observatory's distribution.js engine.

BFR formula: $BFR(t) = ((t + 365.25) / t)^{\beta} - 1$, where t = days since genesis. This is the exact annual growth rate of the floor, not an approximation.

Monte Carlo engine: 100,000 paths generated using Ornstein-Uhlenbeck process with parameters fitted to empirical log residuals from halving cycles 4–5. Monthly withdrawal of ~$8,700 (inflation-adjusted at 3%/year). Stack sold at simulated market price, not floor price.

Model sensitivity: The companion paper verified that volatility decay rates are model-independent, producing identical results (±0.6pp) under the Krueger/Sigman year-based parameterisation.

References

[1] Bengen, W. P. (1994). Determining withdrawal rates using historical data. *Journal of Financial Planning*, 7(4), 171–180.

[2] Cooley, P. L., Hubbard, C. M., & Walz, D. T. (1998). Retirement savings: Choosing a withdrawal rate that is sustainable. *AAII Journal*, 20(2), 16–21. [The Trinity Study]

[3] Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383–417.

[4] Gordon, M. J. (1959). Dividends, earnings, and stock prices. *The Review of Economics and Statistics*, 41(2), 99–105.

[5] Metcalfe, R. (2013). Metcalfe's law after 40 years of Ethernet. *IEEE Computer*, 46(12), 26–28.

[6] Santostasi, G. (2024). Bitcoin and the power law. Working paper. Model parameters: logA = –16.493, β = 5.688.

[7] Scale Invariant Capital. (2026). Percentile-anchored volatility decay analysis. *Bitcoin Power Law Observatory Research Series*, Paper 1.

[8] Sornette, D. (2003). *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton University Press.

[9] Burger, H. C. (2019). Bitcoin's natural long-term power-law corridor of growth. Blog post and independent analysis.

[10] Plan C [@TheRealPlanC]. (2026). Median exponent decay analysis. Published on X, March 9, 2026.