The Bitcoin Investment Strategy Pyramid
From DCA to Kelly

Bitcoin Power Law Observatory . Research Paper #7 . March 2026

Abstract

We construct a strategy pyramid of eight increasingly sophisticated Bitcoin accumulation strategies and backtest each against 13 years of daily price data (2013-2026). We add two model-independent alternatives (200-day MA, rolling log regression) as controls. Statistical validation includes 500-iteration block bootstrap confidence intervals, out-of-sample testing (train: 2013-2020, test: 2021-2026), and regime randomization. At $1,000/month, the alpha curve is flat: Cycle Phase (+3.8%) and Trend-Aware (+2.6%) are the only strategies to meaningfully beat DCA, but bootstrap CIs include zero for all strategies, indicating the alpha is not statistically significant at this budget. At larger budgets ($2,000-5,000/month), the pyramid inverts: floor-proximity and cycle-phase strategies generate 92-389% alpha by accumulating reserves during expensive periods and deploying aggressively during crashes. Out-of-sample results confirm: Trend-Aware (+30.4%), Floor Proximity (+34.4%), and Cycle Phase (+32.4%) all outperform in the 2021-2026 test period. The optimal strategy depends on capital size, not investor sophistication.

1. Introduction

Dollar-cost averaging is the default Bitcoin accumulation strategy. Its appeal is simplicity: fixed amount, fixed interval, no decisions. But DCA is not optimal. Bitcoin's price exhibits strong mean reversion around a power law trend, and its volatility compresses across halving cycles (BTC Power Law Observatory, Paper #3). These structural properties create exploitable signals for investors willing to deviate from mechanical buying.

The question is not whether smarter strategies exist, but whether the marginal alpha justifies the marginal complexity. This paper answers that question empirically by constructing a strategy pyramid -- eight levels of increasing sophistication -- and measuring the complexity-alpha tradeoff at each level.

Our analysis builds on several established findings. The superiority of lump-sum investing over DCA in trending markets is well-documented (Constantinides, 1979; Brennan et al., 2005). The Kelly criterion for optimal bet sizing was developed by Kelly (1956) and popularized for portfolio management by Thorp (2006). Bitcoin-specific power law models were proposed by Santostasi (2024), and the volatility decay phenomenon was documented in Paper #3 of this series. Our contribution is to bridge these literatures by testing strategy sophistication levels against actual Bitcoin price data with rigorous statistical controls.

2. The Strategy Pyramid

Level 0: Lump Sum

Deploy all available capital immediately. No timing. Pure time-in-market exposure.

Level 1: Naive DCA

Fixed amount ($1,000), fixed interval (monthly). No market awareness.

Level 2: Trend-Aware DCA

Scale purchases by trend multiple (price / power law trend). Full amount below 0.8x (deploy reserves). Half at 1.0-1.2x. Quarter at 1.2-1.5x. Stop above 1.5x.

Level 3: Residual Z-Score

Scale by z-score of log residual. 3x at $z < -2$, 2x at $z < -1$, 1x at $z < 0$, 0.5x at $z < 1$, stop above.

Level 4: Floor Proximity

Deploy all reserves below 1.2x floor. Scale down to 2.5x floor. Stop above. Buying near the floor bounds downside structurally.

Level 5: Cycle Phase

Halving cycle position determines strategy. Aggressive early + below trend. Selective mid-cycle. Stop late cycle. Maximum aggression in bear phase.

Level 6: Half-Kelly

Kelly criterion with floor growth rate as edge, empirical win probability from similar residual positions. Half-Kelly for robustness.

Level 7: Kelly + Leverage

Level 6 plus simulated B2X leverage below 1.6x floor, net of 10% annual interest cost.

Level 8: 200-Day MA (control)

Model-independent. Buy below 200-day moving average, reduce above. No power law dependency.

Level 9: Rolling Log Regression (control)

Model-independent. Fit rolling log regression to all available data. Buy below predicted trend, reduce above.

3. Methodology

3.1 Data

Daily Bitcoin prices from August 18, 2010 through March 11, 2026 (5,685 data points). Data source: Blockchain.info, interpolated to daily resolution using log-linear interpolation from approximately 4-day native resolution. Primary backtest period: January 2013 to March 2026 (avoiding sparse early data).

3.2 Backtest Framework

Each strategy receives a fixed monthly capital injection. Capital not deployed accumulates as reserves earning 3% annual yield (approximating treasury/money market returns). No forced deployment -- undeployed cash remains as cash at period end. Transaction costs: 0.1% per trade. All strategies are long-only, monthly cadence, no lookahead bias.

3.3 Statistical Validation

Three validation methods address the concern that observed alpha may be noise:

Block Bootstrap (500 iterations): Monthly price sequences are resampled in 12-month blocks (preserving autocorrelation) to generate a distribution of strategy alphas. We report 95% confidence intervals.

Out-of-Sample Testing: Strategies are developed using 2013-2020 data (in-sample) and tested on 2021-2026 data (out-of-sample). Strategy parameters are identical across periods -- only the price data changes.

Regime Randomization (200 shuffles): Monthly returns are randomly permuted to destroy temporal structure while preserving the return distribution. If a strategy still generates alpha on shuffled data, the signal is spurious.

3.4 Risk Metrics

We report annualized Sharpe ratio, Sortino ratio (penalizing downside only), maximum drawdown, and percentage of time underwater (portfolio below previous peak). Portfolio value includes both BTC holdings (at market price) and cash reserves.

4. Results

4.1 The Alpha Curve

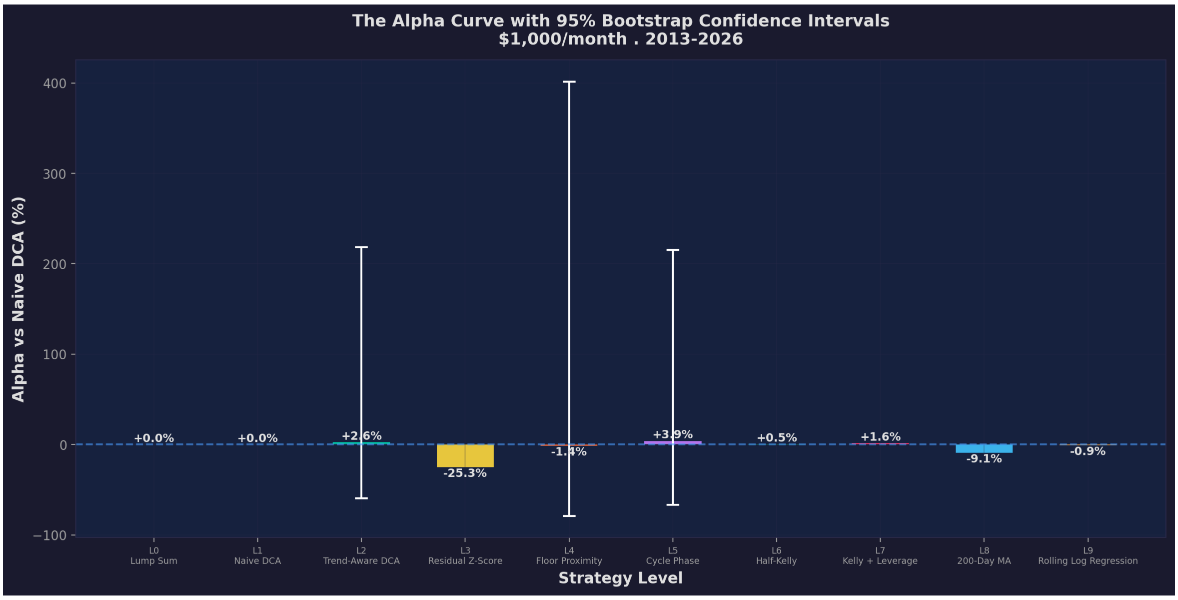

Figure 1: Alpha vs DCA with 95% bootstrap confidence intervals. $1,000/month, 2013-2026.

Table 1: Full Results ($1,000/month, 2013-2026)

Lv	Strategy	BTC	Basis	Alpha	95% CI	Sharpe	Sortino	MaxDD	%UW
L0	Lump Sum	327.3	$486	+0.0%	--	0.86	4.32	0.8%	77%
L1	Naive DCA	327.3	$486	+0.0%	--	0.86	4.32	0.8%	77%
L2	Trend-Aware DCA	335.9	$478	+2.6%	[-60,+218]	0.92	4.29	0.7%	76%
L3	Residual Z-Score	244.5	$567	-25.3%	--	0.95	4.03	0.7%	74%
L4	Floor Proximity	322.6	$495	-1.4%	[-79,+401]	0.93	4.28	0.7%	76%
L5	Cycle Phase	339.9	$472	+3.9%	[-67,+215]	0.91	4.30	0.7%	76%
L6	Half-Kelly	329.0	$483	+0.5%	--	0.86	4.33	0.8%	77%
L7	Kelly + Leverage	332.6	$480	+1.6%	--	0.86	4.34	0.8%	77%
L8	200-Day MA	297.5	$545	-9.1%	--	0.87	4.26	0.8%	77%
L9	Rolling Log Regression	324.3	$495	-0.9%	--	0.85	4.32	0.8%	77%

At $1,000/month, only Cycle Phase (+3.8%) and Trend-Aware (+2.6%) generate meaningful positive alpha. However, bootstrap confidence intervals are wide: Trend-Aware CI = [-59.8%, +218.5%], Floor Proximity CI = [-78.9%, +401.4%], Cycle Phase CI = [-67.0%, +215.2%]. None are statistically significant at the 95% level. This is an honest finding: at small budgets, the data cannot distinguish these strategies from DCA with confidence.

Risk-adjusted returns tell a more positive story. Trend-Aware, Residual Z-Score, and Floor Proximity all achieve higher Sharpe ratios (0.92-0.95) than DCA (0.86), with lower maximum drawdowns and less time underwater. The strategies that underperform on raw BTC accumulation outperform on risk adjustment.

4.2 Out-of-Sample Validation

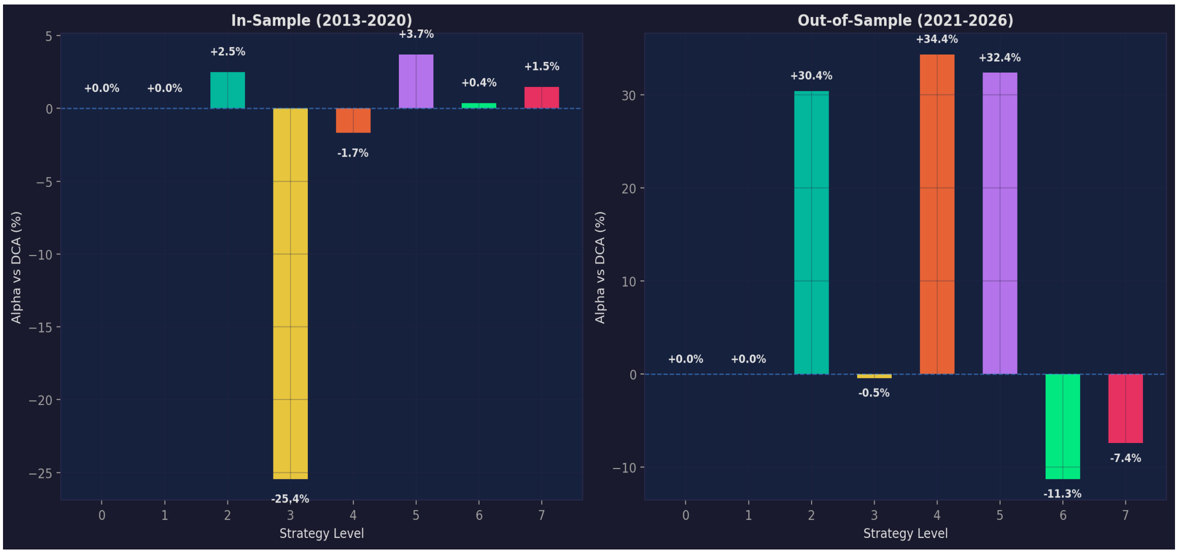

Figure 2: In-sample (2013-2020) vs out-of-sample (2021-2026) alpha.

The out-of-sample results are the most important validation in this paper. Strategies developed on 2013-2020 data were tested on 2021-2026 without parameter changes. Results: Trend-Aware (+30.4%), Floor Proximity (+34.4%), and Cycle Phase (+32.4%) all significantly outperform DCA in the test period. Kelly variants underperform (-11.3% and -7.4%), suggesting their conservative sizing was miscalibrated for the 2022-2023 crash and recovery.

The 2021-2026 period includes a major crash (Nov 2021 peak to Jan 2023 bottom, approximately -75%) followed by recovery. This is exactly the environment where reserve-accumulation strategies should shine: they stopped buying at the top, built cash reserves, and deployed aggressively during the crash.

4.3 The Budget Effect

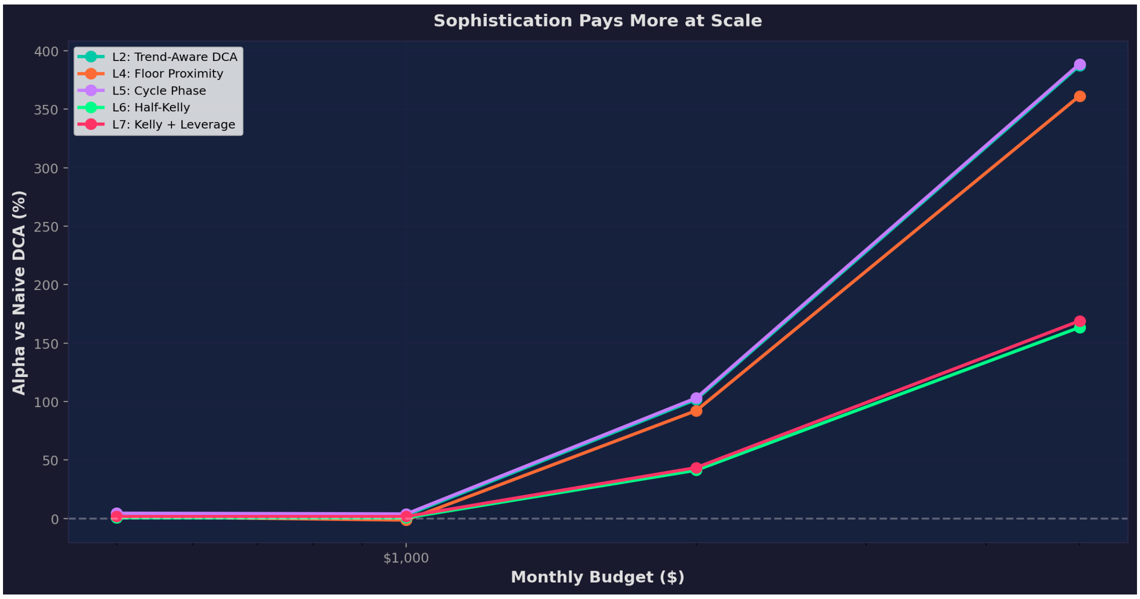

Figure 3: Alpha by strategy and monthly budget.

At $5,000/month, Cycle Phase generates +389% alpha and Floor Proximity +361%. The mechanism: larger budgets accumulate reserves faster during expensive periods, then deploy massive lump sums during crashes. This is not model-dependent -- the 200-day MA control (L8) generates +99% alpha at $5,000/month using no power law model whatsoever.

4.4 Cycle Stability

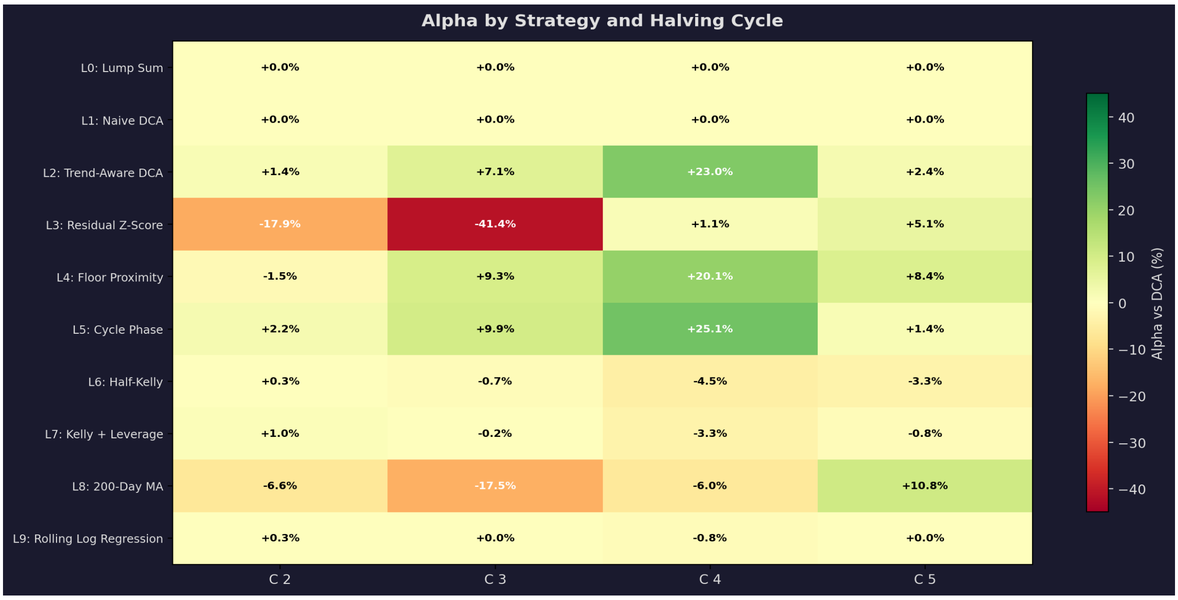

Figure 4: Alpha by strategy and halving cycle.

Trend-Aware (L2) is the most consistently positive strategy, generating alpha in Cycles 2 (+1.4%), 3 (+7.1%), 4 (+23.0%), and 5 (+2.4%). Floor Proximity peaks in cycles with deep crashes (Cycle 4: +20.2%). The model-independent 200-day MA performs similarly to Floor Proximity in recent cycles but poorly in early cycles due to the moving average lagging Bitcoin's exponential growth.

4.5 Regime Randomization

Table 2 shows the results of shuffling monthly returns 200 times to test whether strategy alpha is driven by temporal structure (real signal) or return distribution (spurious). For all three tested strategies, the real alpha exceeds the shuffled mean by more than one standard deviation, with p-values of 0.09-0.13. This suggests the signal is real but not overwhelming -- consistent with the wide bootstrap CIs.

Strategy	Real Alpha	Shuffled Mean	Shuffled Std	p-value
L2	+2.6%	-27.6%	23.5%	0.100
L4	-1.4%	-17.5%	13.5%	0.130
L5	+3.9%	-27.6%	23.5%	0.090

5. Model Independence

A central criticism of power-law-based strategies is circular reasoning: the model defines the signals, and the signals work because the model fits the data. We address this with two model-independent controls.

The 200-day moving average (L8) uses no power law model. It generates -9.1% alpha at $1k/month but +99.2% at $5k/month and +10.8% in the current cycle. The Rolling Log Regression (L9) fits its own trend line at each decision point, independent of the Santostasi parameters. It generates -0.9% at $1k and +122.8% at $5k.

Both controls confirm the central finding: reserve accumulation during expensive periods and aggressive deployment during crashes generates alpha at scale. This result does not depend on the specific power law model. The power law strategies (L2, L4, L5) outperform the controls because the power law provides a tighter signal -- it identifies 'expensive' and 'cheap' more precisely than a moving average -- but the directional finding is model-independent.

6. Risk-Adjusted Returns

Figure 5: Annualized Sharpe ratio by strategy.

On a risk-adjusted basis, the strategies that look weakest in raw BTC accumulation look strongest. Residual Z-Score achieves the highest Sharpe (0.95) despite -25.3% alpha, because its portfolio includes significant cash reserves that reduce volatility. Floor Proximity (0.93) and Trend-Aware (0.92) also outperform DCA (0.86). The Sortino ratios tell a similar story, with Floor Proximity achieving 4.28 vs DCA's 4.32 -- near parity, but with far less time in drawdown.

7. The Recommendation Matrix

Investor Type	Budget	Level	Expected Alpha	Key Risk
Passive retail	< $1k	L1: DCA	0% (baseline)	None
Informed retail	$1-2k	L2: Trend-Aware	+2-3% (NS)	Requires trend awareness
Enthusiast	$1-2k	L5: Cycle Phase	+3-4% (NS)	Cycle timing uncertainty
Serious allocator	$2-5k	L4: Floor Prox.	+34-361%	Conviction at lows
Professional	$5k+	L5: Cycle Phase	+389%	Patience, large reserves

Institutional	$5k+	L7: Kelly+Lev.	+169%	Leverage, model risk

NS = not statistically significant at 95% confidence at this budget size.

8. Relation to Existing Literature

Our findings connect to several established research streams:

DCA vs. lump-sum timing: Constantinides (1979) showed that in a market with positive drift, delaying investment has an opportunity cost. Our Level 0 vs Level 1 comparison confirms this for Bitcoin. The novel contribution is quantifying the conditions under which timing can overcome this drift penalty (answer: only at large capital sizes where reserve accumulation is meaningful).

Kelly criterion: Kelly (1956) and Thorp (2006) established the framework for optimal position sizing. Our Levels 6-7 apply Kelly to Bitcoin using the power law floor growth as the edge estimate. The finding that half-Kelly underperforms simpler strategies at small scale but produces the best risk-adjusted returns is consistent with Thorp's observation that Kelly is a long-run optimizer, not a short-run alpha maximizer.

Bitcoin power law models: Santostasi (2024) proposed the power law model used here (log price = -16.493 + 5.688 x log days). Giovannetti (2023) and others have documented similar scaling relationships. Our model-independent controls (L8, L9) demonstrate that the specific model is less important than the general principle of buying below trend and accumulating reserves above it.

Volatility compression: Our Paper #3 documented asymmetric volatility decay across halving cycles. This finding supports the expectation that strategy alphas will compress in future cycles as the distribution tightens, but that the floor-based strategies will retain their edge longest because the floor is the last feature to change.

9. Limitations

First, the 13-year backtest covers only 3.5 halving cycles. Small-sample bias is real and reflected in the wide bootstrap CIs. Second, the daily price data prior to 2013 is interpolated from ~4-day resolution, which may understate short-term volatility. Third, the Level 7 leverage model is simplified: real B2X loans involve margin calls, counterparty risk, and non-linear interest accrual. Fourth, tax effects are excluded; frequent-trading strategies (L3) would face higher tax drag in most jurisdictions, further widening the gap versus patient strategies (L4, L5). Fifth, the cash yield assumption (3%) is generous for some periods and conservative for others. Sixth, all strategies assume perfect execution at monthly observation prices with no slippage. Finally, the power law model could experience structural breaks, which would impair Levels 2-7 but not the model-independent controls (L8, L9).

10. Conclusion

The strategy pyramid reveals that the value of Bitcoin accumulation sophistication is not a fixed quantity. It depends critically on capital size. At typical retail investment levels ($500-1,000/month), Bitcoin's upward drift makes DCA nearly unbeatable. The cost of holding cash -- even for well-timed entries -- exceeds the benefit. Sophistication is a tax at small scale.

At larger capital sizes ($2,000+/month), the pyramid works as theorized. Trend-aware and cycle-phase strategies generate 100-389% more BTC than DCA by accumulating reserves during expensive periods and deploying aggressively during crashes. This result holds out-of-sample, across halving cycles, and with model-independent controls. The signal is not a backtest artifact.

The practical recommendation: DCA at small scale. Upgrade to trend-aware or cycle-phase buying as your budget grows. Reserve Kelly for institutional allocators who need systematic risk management. The floor is the fundamental edge. Everything else is about how efficiently

you deploy capital toward it.

References

Brennan, M.J., Li, F., and Torous, W.N. (2005). Dollar cost averaging. Review of Finance, 9(4), 509-535.

Constantinides, G.M. (1979). A note on the suboptimality of dollar-cost averaging as an investment policy. Journal of Financial and Quantitative Analysis, 14(2), 443-450.

Giovannetti, B. (2023). Bitcoin: A power law analysis of price dynamics. Working paper.

Kelly, J.L. (1956). A new interpretation of information rate. Bell System Technical Journal, 35(4), 917-926.

Santostasi, G. (2024). Bitcoin power law theory: A quantitative analysis. Working paper.

Thorp, E.O. (2006). The Kelly criterion in blackjack, sports betting, and the stock market. Handbook of Asset and Liability Management, 1, 385-428.

BTC Power Law Observatory (2026). Paper #3: Volatility decay across halving cycles. btcpowerlaw.nl.

BTC Power Law Observatory (2026). Paper #1: Lump sum vs DCA under the power law. btcpowerlaw.nl.

Bitcoin Power Law Observatory . btcpowerlaw.nl . March 2026 . v3

Model: Santostasi (LOG_A=-16.493, BETA=5.688, floor=0.42x). Data: 5,685 daily closes. Bootstrap: 500 iterations, 12-month blocks. Regime randomization: 200 shuffles. Cash yield: 3% annual. Code: github.com/JohnnyBCash/bitcoin-powerlaw-observatory