Kelly Sizing With Floor-Anchored Growth

A Walk-Forward Test of a Bitcoin, Gold and Cash Allocation Rule Against Buy and Hold

Scale Invariant Capital Research | 2026-11-06 | Paper 19 | v1.3

Claims used: C1.1, C1.3, C1.4, C6.1 to C6.8, C7.2, C7.3 (Observatory Claims Registry v2.7). Drafted as Paper 17 and published as Paper 19, because the site's Paper 17 is Playbook Optimization. *Revised after two external reviews, a blind replication and an independent pre-publication review; see the Revision notes.*

Abstract

The Kelly criterion sizes a position from its expected return and variance. Applied to Bitcoin with the historical average return, it recommends roughly 100% Bitcoin almost all the time. We replace the historical return with **floor-anchored growth**: the growth from today's price to a power-law floor four years ahead. This input is high near the floor, negative near and above trend, and falls automatically as Bitcoin ages. In the model's own four-year forecast it sits near the 14th percentile, so it acts as a cautious, low-quantile estimate of growth. This is not textbook Kelly, which needs the expected return. It is a conservative, robust-Kelly-style sizing rule: Kelly's mean-variance objective, solved with a low-quantile growth estimate in place of the mean.

We test the rule on a Bitcoin, gold and cash portfolio from 2012 to October 2026. Decisions are monthly and use only data available at each date. The power law and the floor are re-estimated each month (walk-forward). Weights are long only, with no leverage, 0.25% trading costs and no taxes. The question that matters to Bitcoin holders is whether the rule ends with **more Bitcoin** than buying once and holding.

We separate the evidence into three levels.

1. **Strong: the floor timing adds value beyond rebalancing.** A constant mix with the same average weights ended ahead of buy and hold in 9.7% of 5,671 start and end month pairs since 2014. Walk-forward floor-Kelly ended ahead in 98%. Its log-growth advantage over the matched mix is +7.5% per year ($p < 0.01$). This advantage stays significant after a Holm correction across all 12 full-Kelly floors we explored (adjusted $p = 0.002$), for seven of the eight power-law floors and none of the four other shapes, and 53 of the 56 comparable configurations in the full search beat their own matched mix.
2. **Suggestive: full floor-Kelly accumulates more Bitcoin than buy and hold.**
   - **Historically:** it was ahead from every January start from 2012 to 2025, in all three non-overlapping four-year windows, and with model parameters frozen at four past dates and never recalibrated. From 2014 it ended at 2.16× the Bitcoin of a holder (90% bootstrap interval 1.05× to 4.54×).
   - **But not proven:** the log-growth advantage over holding, +6.0% per year ($p = 0.04$), does not survive a Holm correction across the 12 floors (adjusted $p = 0.46$). Nor does it survive a

deflated Sharpe ratio over all 57 distinct configurations explored (0.09, or 0.48 if they count as only five independent trials).

3. **Unproven, and weak in simulation: the edge continues.** With the floor held at its October 2026 fit, the simulated rule keeps its edge only while Bitcoin's cycles persist. It loses Bitcoin if cycles disappear or if Bitcoin re-rates permanently above its trend, and it is a coin flip under regime switching. Refit inside each simulated path (yearly, on monthly closes; the backtested and tracked rule refits monthly), it ends behind holding at the median in 9 of 14 worlds, and at only 1.25× a holder's Bitcoin with cycles as large as history.

The edge is specific to the power-law floor (four walk-forward floors with other shapes fail), and half Kelly is a risk-control rule (drawdowns of -37% to -43%), not an accumulation rule.

A blind re-implementation from the paper and data alone reproduced every number it tested, once one unstated input was clarified. A monthly prospective record of the frozen rule starts with its November 2026 allocation.

Keywords: Bitcoin, gold, Kelly criterion, power law, walk-forward backtest, portfolio allocation, mean reversion, rebalancing, backtest overfitting.

Key findings

1. **Floor-anchored growth turns Kelly into a counter-cyclical rule.** It goes all in near the floor and cuts Bitcoin to 39% to 56% at cycle peaks (Figure 4).
2. **Walk-forward, full floor-Kelly ended with more Bitcoin than buy and hold** in 98% of start and end pairs since 2014 (median 1.34×) and from every January start from 2012 to 2025, at 1.11× to 2.16× (Figures 5 and 6).
3. **The accumulation edge is not statistically established:** +6.0% per year over holding, Holm-adjusted p = 0.46 across 12 floors, deflated Sharpe ratio 0.09 to 0.77 (Tables 7 and 9).
4. **Freezing the parameters did not remove the result.** It beat holding with parameters frozen at December 2016, 2018, 2020 and 2022, at 1.11× to 1.65× (Figure 7).
5. **The timing is the edge, not rebalancing.** The same average weights, rebalanced mechanically, ended ahead in 9.7% of pairs (Figure 6).
6. **The edge depends on the power law.** Over starts from 2016 to 2022, four walk-forward floors with other shapes ended ahead in 0% to 41% of pairs, against 99% for the power-law floor (Figure 12).
7. **Two input choices matter.** A two-year horizon or a 60-month covariance window weakens the result (41% and 74% of pairs ahead). Other changes to the Kelly inputs barely move it (Figure 11).
8. **In simulation, the edge needs cycles and a stable floor.** With the floor frozen at October 2026, it holds until cycle amplitude has shrunk by about 90% over 20 years, fails under a permanent twofold repricing above trend, and is a coin flip under regime switching (Figures 14 to 17). With the floor refit inside each path (yearly, on monthly closes), it is behind holding at the median in 9 of 14 worlds (Table 14).
9. **The floor input is a cautious quantile.** It sits near the 14th percentile of the model's four-year forecast. Kelly fed the model's 25th or 10th percentile directly also beat holding, in 97% and 77% of pairs (Figures 2 and 13).

10. **The whole search is disclosed.** All 57 distinct configurations explored are shown. The headline ranks 11th of 56 comparable configurations on its edge over holding, and the probability of backtest overfitting is 0.19 (Figure 8).

1. Introduction

1.1 The problem with Kelly on Bitcoin

The Kelly criterion (Kelly, 1956) chooses the portfolio that maximises the expected logarithm of wealth, and with it long-run compound growth (Rotando and Thorp, 1992; Thorp, 2006). For several assets with arithmetic excess returns m and covariance Σ, the unconstrained Kelly weights are $w^* = \Sigma^{-1}m$. Practitioners scale this down to a fraction of Kelly because the rule is very sensitive to errors in m (MacLean, Thorp and Ziemba, 2010).

Feed Kelly Bitcoin's historical average return and it recommends close to 100% Bitcoin throughout its history. In our backtest, full historical-average Kelly holds 100% Bitcoin and is buy and hold; at half Kelly it holds 90% on average and ends from 2012 with 0.37× a holder's Bitcoin (Table 2). A rule that always says "all in" is not an allocation rule.

1.2 The idea: growth anchored to a floor

The power law gives a structural alternative to the historical average. Near the floor the expected return is far higher than trend growth; near a cycle peak it is far lower. The input we test is:

> Assume the price will sit on the floor four years from now. The planning growth rate per year is the growth from today's price to that future floor.

The full floor growth is earned only from the floor. From trend, the floor needs years to rise above today's price, so the floor-anchored growth over short horizons is negative. That waiting time grows as Bitcoin ages. The input keeps Kelly in check by construction.

This is conservative only within the power-law framework. If the trend itself fails, a floor-anchored rule can be badly wrong in either direction. Sections 3.8 and 3.9 test both directions.

1.3 Two questions

Institutional allocators ask whether the rule improves risk-adjusted returns and drawdowns. Bitcoin holders ask whether **using the cycles increases the number of Bitcoin they own**. Buying early and holding is hard to beat on that measure, because any time spent partly in gold or cash costs coins during a bull market. We answer both questions, and we measure the second directly.

1.4 What this paper contributes

1. A floor-anchored growth input for Kelly sizing, with its closed form and horizon behaviour (Section 2.2).
2. A walk-forward backtest from 2012 to 2026 with matched-exposure controls, every start and end month pair, frozen-parameter tests and three independent windows (Sections 3.1 to 3.5).

3. A statistical treatment that separates what is established from what is suggestive: a bootstrap of the Bitcoin stack ratio, and a Holm correction across every floor variant explored (Section 3.5).
4. Robustness to the Kelly inputs and a model-free alternative signal (Section 3.7).
5. Forward stress tests built to break the rule: shrinking cycles, a random walk, a broken trend, a permanent repricing, regime switching, fat tails, a structural break and a long low-volatility spell (Section 3.8).
6. An interpretation of the floor input as a low quantile of a conditional forecast, which links the rule to robust Kelly (Section 2.3), tested directly with quantile inputs (Section 3.7).
7. The full specification search, with a deflated Sharpe ratio and the probability of backtest overfitting (Section 3.5).
8. Four alternative valuation floors, including an on-chain realised-price floor (Section 3.7).
9. A blind replication (Section 3.10) and a frozen, append-only prospective record (Section 7).

1.5 Related work

- **Growth-optimal investing:** Kelly (1956); Rotando and Thorp (1992); Thorp (2006); MacLean, Thorp and Ziemba (2010).
- **Drawdown and risk constraints:** Grossman and Zhou (1993) derive growth-optimal policies under a maximum-drawdown constraint. Busseti, Ryu and Boyd (2016) impose a drawdown-probability constraint on Kelly betting and find it beats fractional Kelly at equal risk.
- **Mean reversion:** under log utility the optimal policy is myopic even when returns mean-revert (Merton, 1971). More risk-averse investors add an intertemporal hedging demand (Kim and Omberg, 1996; Wachter, 2002), which we do not model.
- **The rebalancing premium:** the excess growth from rebalancing a constant mix of volatile assets is documented by Fernholz and Shay (1982) and Bouchey, Nemtchinov, Paulsen and Stein (2012). Our matched-mix control isolates it.
- **Bitcoin as an asset:** Liu and Tsyvinski (2021) show that cryptocurrency returns are not explained by standard asset-pricing factors. Platanakis and Urquhart (2020) find gains from adding Bitcoin to stock and bond portfolios under several allocation rules.
- **The power law:** the empirical power-law corridor was popularised by Burger (2019) and Santostasi (*The Physics of Bitcoin*, n.d.). The Observatory fit (β = 5.688) and its out-of-sample performance are documented in Paper 12.
- **Robust optimisation and robust Kelly:** mean-variance portfolios that plan against the worst expected return in an uncertainty set (Goldfarb and Iyengar, 2003; Tütüncü and Koenig, 2004), and growth-optimal portfolios that guard against a set of plausible return distributions (Rujeerapaiboon, Kuhn and Wiesemann, 2016; Sun and Boyd, 2018). Section 2.3 reads the floor input in this way.
- **Backtest overfitting:** we use the deflated Sharpe ratio (Bailey and López de Prado, 2014) and the probability of backtest overfitting (Bailey, Borwein, López de Prado and Zhu, 2017), following the concerns of Harvey, Liu and Zhu (2016).
- **Regime switching:** our regime world follows the Markov-switching idea of Hamilton (1989).
- **Covariance estimation:** our shrinkage robustness check follows Ledoit and Wolf (2004).
- **Floor stability out of sample:** Paper 20 freezes floors at 172 past month-ends and carries them forward. The 1% quantile line, this paper's floor, was below price on a median 14% of later days and held for four years at 5 of 94 freeze dates (C7.2); the trend times its lowest-1% multiple was

the most stable floor tested (C7.3). Paper 18 refits the trend itself every month and finds its exponent settled in 2016.

2. Data and Method

2.1 Data

- **Bitcoin:** `bitcoin-powerlaw-site/shared/datasets/btc_historical.json`, 5,951 daily closes from 2010-07-18 to 2026-11-01, sampled at month-end.
- **Gold:** month-end close of the SPDR Gold Shares ETF (GLD), 203 months from 2009-12 to 2026-10 with no gaps. GLD holds bullion valued at the LBMA Gold Price PM, so its returns are those of an investable gold holding, net of the 0.40% fee. It tracks COMEX gold futures to a median level gap of -0.03%, with a monthly return correlation of 0.991.
- **Cash:** 3-month US Treasury bill rate (FRED TB3MS).
- **On-chain valuation (v1.2 alternative floor only):** daily MVRV ratio (market value to realised value) from the Coin Metrics community API, series `CapMVRVCur`, 5,955 days from 2010-07-18 to 2026-11-05. Realised price is price divided by MVRV.
- **Published power law (reference model):** $\log_{10}(P) = -16.493 + 5.688 \cdot \log_{10}(d)$, where d is the number of days since 2009-01-03 (C1.1).
- **Official floor (reference):** 0.42× trend (C1.3). The Observatory defines a floor as a level price spends little time below and climbs back above quickly, not a boundary (`research/floor-definition.md`). Bitcoin has spent 3.7% of days since July 2010 below this one, in 23 visits, none longer than 86 days; since 2016, 0.3% of days, with a longest visit of 7 days (C1.4).

2.2 Floor-anchored growth

At decision date t, with Bitcoin price P_t and horizon H years, the floor-anchored log growth is

$$g(t) = \ln(\text{Floor}(t + H) / P_t) / H.$$

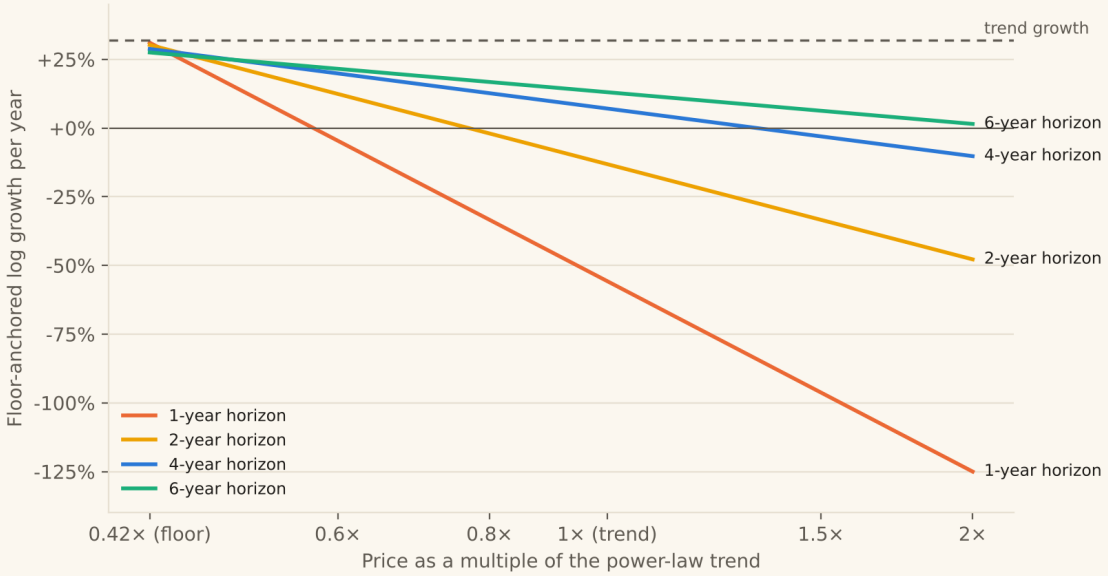
For a floor that is a fixed multiple f of a power-law trend, this separates into a valuation term and a trend term:

$$g(t) = [\ln(f / m_t) + \beta \cdot \ln(1 + H / a_t)] / H,$$

where m_t is the price as a multiple of trend and a_t is Bitcoin's age in years. Figure 1 plots g against the trend multiple for several horizons.

Figure 1. Floor-anchored growth keeps Kelly in check

Growth per year if the price sits on the floor (0.42× trend) H years from now, as of October 2026.
Only from the floor does it equal the floor's full growth. Near or above trend it is negative for years.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 1. Floor-anchored growth by trend multiple and horizon, October 2026, for the 0.42× floor. Note that only at the floor does *g* equal the floor's full growth. At trend it is negative for horizons shorter than about three years.

Table 1. Floor-anchored growth, October 2026, 0.42× floor (Bitcoin age 17.8 years).

Price today	Years until the floor passes today's price	*g* at H = 1	*g* at H = 2	*g* at H = 4
0.42× (on the floor)	0	+31%	+30%	+29%
0.59×	1.1	-3%	+13%	+20%
1.0× (at trend)	2.9	-56%	-13%	+7%
1.5×	4.5	-96%	-33%	-3%

At trend, the time until the floor passes today's price was 0.8 years in January 2014 and is 2.9 years in October 2026. It reaches 5.1 years in January 2040 (C6.1).

2.3 What the floor input is: a low conditional quantile

The floor input is not an expected return in the usual sense. It is the growth to one point of the price distribution four years ahead. To locate that point, we fit two models each month, on data up to that month only: the power law by OLS on daily closes, and an AR(1) model of the monthly $\log_{10}$ residual around it. This gives a normal conditional distribution for $\log_{10}$ price $n = 12H$ months ahead:

$$\text{mean} = \text{trend}(t + H) + \mu + \varphi^{n} (r_t - \mu), \quad \text{variance} = s^2 (1 - \varphi^{2n}) / (1 - \varphi^2),$$

where r_t is today's residual, μ its long-run mean, φ the monthly persistence and s the innovation standard deviation. The quantile of the headline floor in this distribution is

$$q(t) = \Phi\left((\log_{10} \text{Floor}(t + H) - \text{mean}) / \sqrt{\text{variance}} \right).$$

Figure 2. The floor input is a low quantile of the model's own forecast

Where the headline floor four years ahead sits in the walk-forward conditional distribution of log price (power law plus AR(1) residual, both fitted on data up to each month). Monthly decisions, January 2014 to September 2026.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_v12_figures.py).

Figure 2. The quantile of the headline floor four years ahead within the walk-forward conditional forecast, each month from 2014. Note that it stays between 11% and 20% in most months: the floor and the forecast largely rise and fall together.

The median of q is 14%. Across months its 10th to 90th percentile is 11% to 20%, and its full range is 3% (mid-2014) to 26% (early 2017). The floor input is therefore close to a one-in-seven downside outcome of the model's own four-year forecast.

Because the forecast is normal in log price, the same fact can be written as a penalty. The floor input equals the forecast's central growth rate minus a multiple κ of the forecast's standard deviation:

$$g_floor(t) = g_central(t) - \kappa(t) \cdot \sigma_forecast(t),$$

where both growth rates and $\sigma_forecast$ are per year over the horizon H.

- **The penalty is about one standard deviation.** From 2014 to 2026 the median κ is 1.06 (10th to 90th percentile: 0.85 to 1.25).
- **In growth terms it is large.** $\sigma_forecast$ is about 20% per year, so the penalty is about 20 percentage points of growth per year (range 14 to 42).
- **Central forecast against floor.** The central forecast implies a median of 49% growth per year; the floor input implies 29%.

This gives the rule a standard interpretation. Planning against an expected return about one standard deviation below the central estimate is analogous to robust mean-variance optimisation with an ellipsoidal uncertainty set for the mean (Goldfarb and Iyengar, 2003). The analogy is not exact: here the penalty uses the spread of the four-year outcome rather than the estimation error of the mean, and it applies to Bitcoin only. Applied to Kelly, it is a form of robust growth-optimal investing, which maximises growth against a pessimistic member of a set of plausible distributions (Rujeerapaiboon, Kuhn and Wiesemann, 2016; Sun and Boyd, 2018). Two consequences follow.

1. **The power law does not need to be the mean.** The rule needs only a good estimate of a low quantile of the four-year outcome. A floor fitted to the lowest closes aims at the low tail of the distribution, though in practice it lands near the 14th percentile of this forecast, not the 1st.

Paper 20 finds that the same line, frozen at a date, was above price on a median 14% of later days, an out-of-sample match for this reading.

2. **The quantile can be chosen directly.** Section 3.7 feeds Kelly the 50th, 25th, 10th, 5th and 1st percentiles of the same conditional distribution and compares them with the floor.

The link to robust optimisation can be made exact for one uncertainty set. Because Bitcoin's weight is never negative, planning against the worst Bitcoin mean in an interval is the same as plugging in the interval's lower end. Floor-Kelly is therefore the robust solution for an interval uncertainty set whose lower end is the floor growth (Tütüncü and Koenig, 2004). What is not standard is how that end is set: by the floor, not by estimation error.

The interpretation has limits. The AR(1) model is fitted on the same history, and the normal assumption understates fat tails. The quantile describes the rule; it does not guarantee its behaviour.

2.4 Floors: the headline and the reference

	Headline (walk-forward)	Reference (published model)
Floor	Quantile regression at $\tau = 0.01$ of $\log_{10}$ price on $\log_{10}$ days, **refit each month on data up to that month**	0.42× the published trend
Parameters known at decision time?	Yes	No: the 0.42 multiple and the trend were chosen with the full history
Floor today	0.44× trend, slope 5.954, rising toward trend	0.42× trend, constant multiple

The headline is the version an investor could have run in real time. The reference is shown for continuity with the Observatory's published model.

2.5 The Kelly allocation

Kelly needs arithmetic returns. Each month we convert log growth to an arithmetic excess return, $m = g + \sigma^2/2 - \ln(1 + r)$, where σ^2 is the asset's variance and r the annual T-bill rate. This is the standard lognormal relation between median and mean growth. It holds exactly in continuous time and closely at a monthly frequency when returns are near-lognormal. Gold's expected log growth is the **expanding** mean of its monthly log returns, from August 2010 (the first month with both a Bitcoin and a gold return) to the decision month, times 12. r is the TB3MS rate of the decision month; cash earns $r/12$ in the following month.

Σ is the sample covariance of monthly log returns over the trailing 36 months, annualised. Until mid-2013 (decisions before July 2013), when fewer than 36 months of joint returns exist, it uses all available months. The gold-mean window is the one input choice that a blind replication found was not stated in v1.1, and it matters: the expanding mean gives the lowest Bitcoin stack of the four estimated windows we test, from both 2012 and 2014 (Section 3.7). The weights solve

maximise $w^{\mathsf{T}}m - w^{\mathsf{T}}\Sigma w / (2c)$ subject to $w_\text{BTC} \geq 0$, $w_\text{gold} \geq 0$, $w_\text{BTC} + w_\text{gold} \leq 1$.

This is the second-order (mean-variance) approximation to the fractional-Kelly objective E[log wealth]. The solution with $c = 1$ is the growth-optimal portfolio under lognormal returns (Merton, 1971; Thorp, 2006). We solve on a 1% grid; Section 3.7 confirms that an exact solver gives the same results. Any remainder sits in T-bills. $c = 1$ is full Kelly and $c = 0.5$ is half Kelly. The headline uses $H = 4$ years, one halving cycle.

2.6 Strategies and controls

Strategy	Rule
Buy and hold Bitcoin	Buy once and hold
50/50, rebalanced yearly	Reset to 50/50 Bitcoin and gold every January
Kelly, historical average	Kelly with Bitcoin's trailing average return (the textbook plug-in)
Kelly, trend growth	Kelly with the trend's growth rate, β / age
Floor-Kelly	Kelly with floor-anchored growth (Section 2.2)
Matched mix	A constant Bitcoin and gold mix, rebalanced monthly, at the same average weights as the floor-Kelly run it is compared with

The matched-mix weights are the floor-Kelly run's average weights from its start to October 2026, not over each pair's window. Any rule that sells some Bitcoin on rallies and buys on dips earns a rebalancing premium. The matched mix earns that premium at the same average exposure but with no timing. Matching average weights does not match the time path of exposure. That is the point: the timing is what we test.

2.7 Execution rules and the Bitcoin stack

- Decisions use month-end prices known at that date and apply to the following month. A run "from January 2012" buys at the December 2011 close, so January is the first month held.
- The Bitcoin month-end price is the last daily close in the month. The floor four years ahead is evaluated at the month-end age plus 4 × 365.25 days.
- Long only, no leverage, monthly rebalancing.
- 0.25% cost on the notional traded in Bitcoin and gold, including the initial purchase, deducted before the new holdings are set; the cash leg is free. 1% as a stress test.
- Drawdowns are measured on month-end marks. "Average weight" is the mean of the target weights at decision dates.
- No trade at the final mark.
- No taxes. All results are before tax, by design; Section 5 states what that leaves out.
- An engine test suite runs five tests of four properties: the buy-and-hold stack is constant, zero-cost arithmetic is exact, no decision can see later prices (for the reference floor, the walk-forward floor and the moving-average floor), and the exact Kelly solver matches the grid.

Bitcoin stack vs hold is the portfolio's value in Bitcoin at each month's price, divided by the Bitcoin a buy-and-hold investor with the same starting capital owns. Above 1.00× means more Bitcoin than holding.

2.8 Robustness design

1. **Start dates:** every January from 2012 to 2025.
2. **Every start and end month pair:** first months held from 2014-01 to 2022-10, end months at least 48 months after the first month held, 5,671 pairs. These overlap heavily, so Section 3.5 adds independent tests.
3. **Frozen parameters:** the power law and floor fitted only on data up to December 2016, 2018, 2020 and 2022, then held fixed, with no recalibration. Each frozen run starts at the freeze date's close, so the first month held is the following January. Only the floor is frozen; the covariance and gold mean keep updating, as they would for a real investor.
4. **Independent windows:** three non-overlapping four-year windows.
5. **Inference:** a paired stationary block bootstrap (Politis and Romano, 1994; mean block 12 months, 5,000 draws) and a Holm correction across all 12 full-Kelly floors explored: eight power-law variants and four other shapes.
6. **Kelly inputs:** covariance windows, EWMA, shrinkage, the solver, horizons and costs.
7. **Model-free signal:** a 200-week moving-average floor.
8. **Forward Monte Carlo:** 20 years and 1,000 paths in each of 14 worlds (Table 14). Monthly log residuals follow the fitted AR(1) around the trend (φ = 0.932, half-life 9.8 months). Bitcoin and gold innovations are drawn together in 12-month blocks from 2012 to 2026. The start point is October 2026 at 0.49× trend. In the simulation the strategy uses the rising floor fitted on data to October 2026 and holds it fixed. Section 3.8 also reports the rule with the floor refit inside each path (yearly, on monthly closes), the closest the simulation comes to the monthly refit of the backtest.
9. **Added in v1.2:** the full specification search with a deflated Sharpe ratio and the probability of backtest overfitting; the gold-window test; the floor input as a conditional quantile and as a penalised central forecast; Kelly on conditional quantiles; three more walk-forward floors; four more simulated worlds; a blind replication; and a frozen prospective record.

3. Results

3.1 Headline: $1M from January 2012

Table 2. $1M from January 2012 to October 2026.

Strategy	End value	CAGR	Max drawdown	Worst year	Sharpe	Avg BTC weight	BTC stack vs hold
Buy and hold Bitcoin	$16.3B	92.3%	-81%	-73%	0.78	100%	1.00×
50/50, rebalanced yearly	$4.5B	76.3%	-60%	-38%	0.69	54%	0.28×
Kelly, historical average (half)	$6.0B	79.8%	-75%	-73%	0.78	90%	0.37×
Kelly, trend growth (full)	$17.1B	92.9%	-78%	-73%	0.78	98%	1.05×
Floor-Kelly, walk-forward (full)	**$20.6B**	**95.4%**	**-66%**	**-59%**	**0.87**	**84%**	**1.26×**
Floor-Kelly, walk-forward (half)	$821.4M	57.2%	-43%	-34%	0.91	46%	0.05×
Floor-Kelly, published model (full)	$19.6B	94.7%	-60%	-54%	0.90	76%	1.20×

The Sharpe ratio uses monthly returns in excess of the T-bill rate.

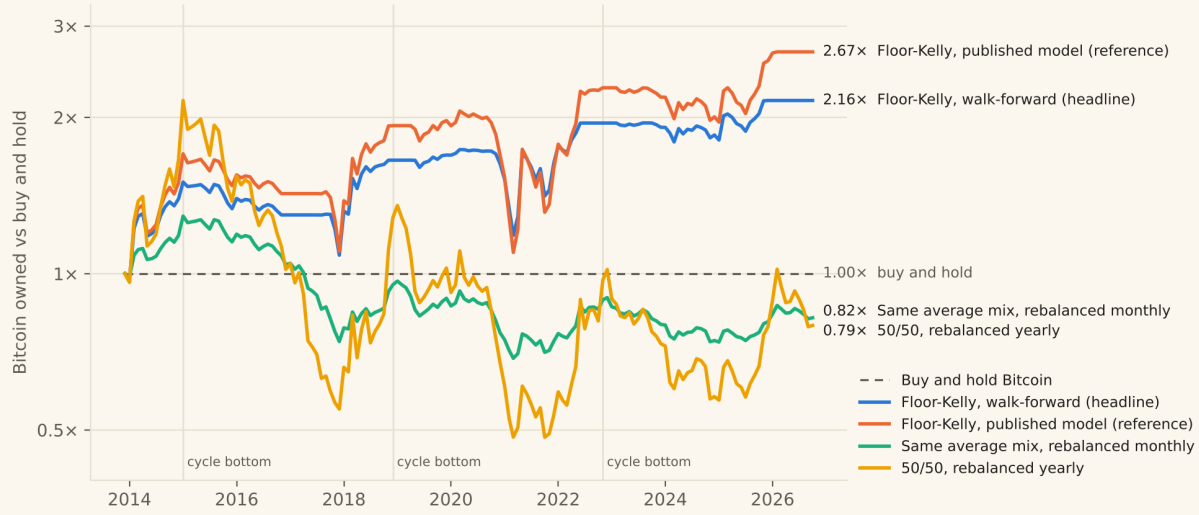
Walk-forward full floor-Kelly ended with 26% more Bitcoin than holding, and the deepest drawdown was 15 points shallower. The 2012 start is the most extreme case: Bitcoin rose about 16,000-fold, which rewards full exposure above all else.

3.2 How the rule gains Bitcoin

Figure 3 shows the stack over time from 2014, and Figure 4 the weight the rule chose at each valuation.

Figure 3. Walk-forward floor-Kelly grew the Bitcoin stack; plain rebalancing did not

$1M from January 2014, monthly decisions, 0.25% costs. Portfolio value in BTC divided by the BTC a buy-and-hold investor owns.

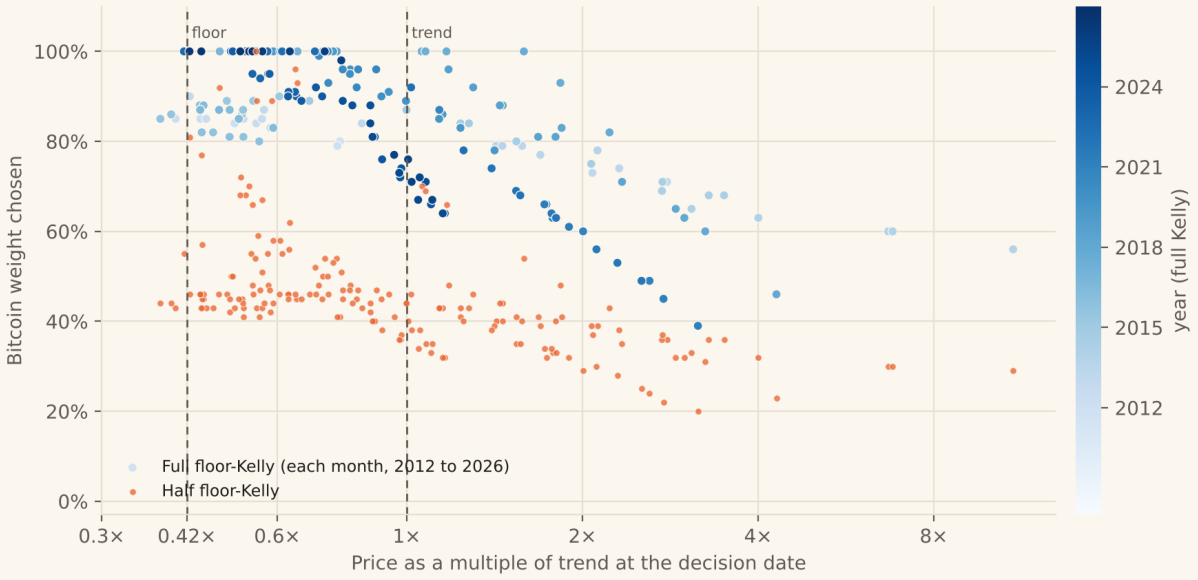


Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 3. Bitcoin stack relative to buy and hold, $1M from January 2014. Note that floor-Kelly loses ground in each bull run, recovers it in each bear market, and ends each cycle further ahead. The matched mix and the yearly 50/50 spend most of the period below 1×.

Figure 4. The rule is counter-cyclical: all in near the floor, light near peaks

Bitcoin weight chosen each month from January 2012 against the trend multiple at that date. The rest sits in gold or T-bills.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 4. Bitcoin weight chosen each month from 2012 by walk-forward floor-Kelly, against the trend multiple. Note the near-100% weights at and just above the floor, and the decline as the price rises above trend.

Table 3. Walk-forward full floor-Kelly from January 2012 at cycle turning points.

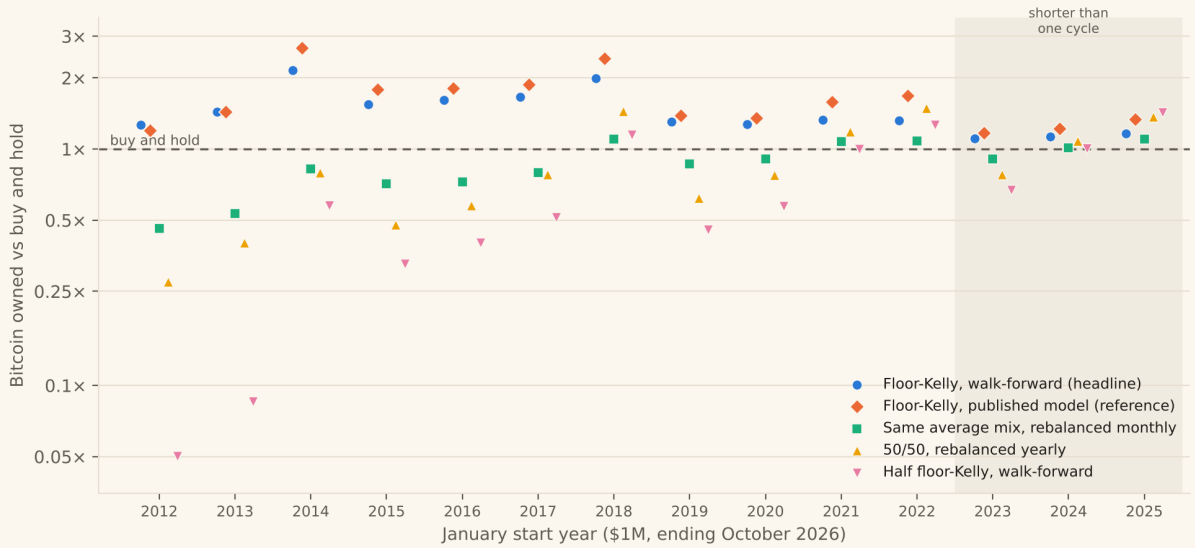
Month	Market	BTC weight	Stack vs hold
2013-11	Peak	56%	0.49×
2015-01	Bottom	90%	0.88×
2017-12	Peak	46%	0.64×
2018-12	Bottom	100%	0.97×
2021-03	Peak	39%	0.70×
2022-11	Bottom	100%	1.14×
2026-10	End	100%	1.26×

The rule sells Bitcoin on the way up and buys it back near the floor with gold and cash. Measured from one cycle bottom to the next, the stack rose from 0.88× to 0.97× to 1.14×.

3.3 Every start year and every start and end month

Figure 5. Walk-forward floor-Kelly ended ahead of buy and hold from every start year

Bitcoin stack at October 2026 relative to buy and hold, for each January start from 2012 to 2025.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 5. Bitcoin stack at October 2026 relative to buy and hold, for each January start from 2012 to 2025. Note that walk-forward floor-Kelly (blue circles) sits above 1× in every year.

Table 4. End value of $1M by January start year, ending October 2026.

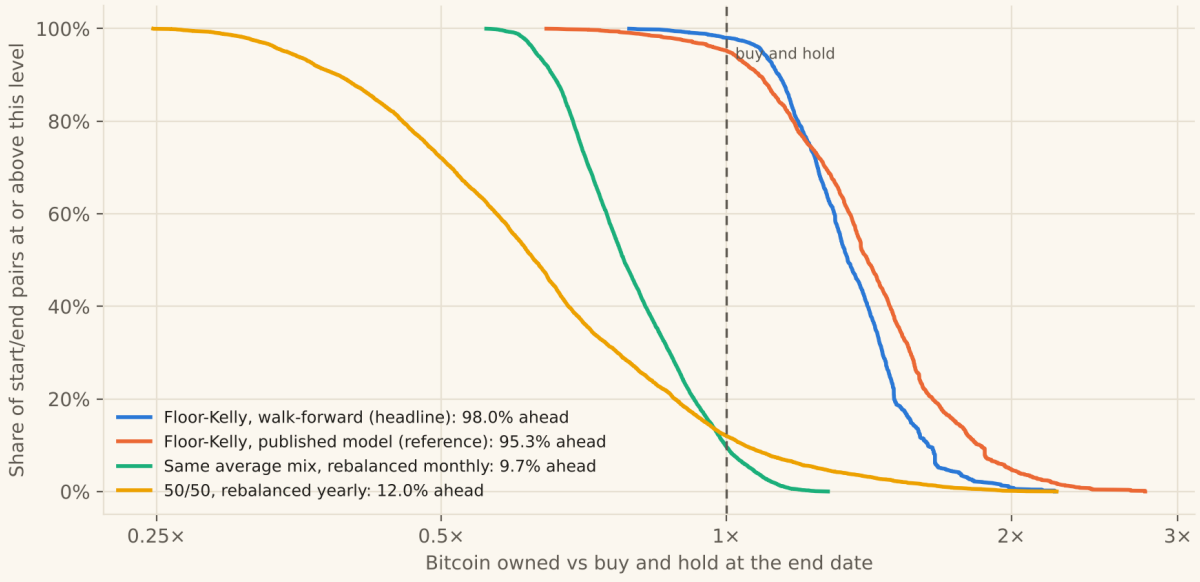
Start	Buy and hold	50/50 yearly	Floor-Kelly full, walk-forward	Floor-Kelly half, walk-forward	Floor-Kelly full, published model
2012	$16.3B	$4.5B	$20.6B	$821.4M	$19.6B
2014	$101.7M	$80.8M	$219.2M	$58.8M	$272.0M
2016	$179.1M	$103.3M	$289.5M	$71.9M	$324.4M
2018	$5.6M	$8.1M	$11.0M	$6.4M	$13.5M
2020	$10.6M	$8.2M	$13.5M	$6.1M	$14.4M
2022	$1.6M	$2.4M	$2.2M	$2.1M	$2.7M
2024	$1.8M	$2.0M	$2.1M	$1.8M	$2.2M
2025	$0.8M	$1.1M	$1.0M	$1.2M	$1.1M

Starts from 2023 on are shorter than one four-year cycle. All 14 start years are in `analysis/f007_start_years.md`.

The start year changes the size of the result enormously: holding compounded at 92.3% per year from 2012 and 10.7% from 2022. It does not change the ranking on Bitcoin stack. Walk-forward full floor-Kelly ended ahead from every start year, at 1.11× to 2.16×. For the full-cycle starts from 2012 to 2022, its maximum drawdown was -59% to -66%, against -65% to -81% for holding.

Figure 6. Across 5,671 start and end months, floor-Kelly was ahead in most

Every start month from January 2014 to October 2022 and every end month at least four years later.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 6. Distribution of the Bitcoin stack relative to buy and hold across 5,671 start and end month pairs. Note that both floor-Kelly curves sit mostly right of 1×, and both rebalancing benchmarks mostly left of it.

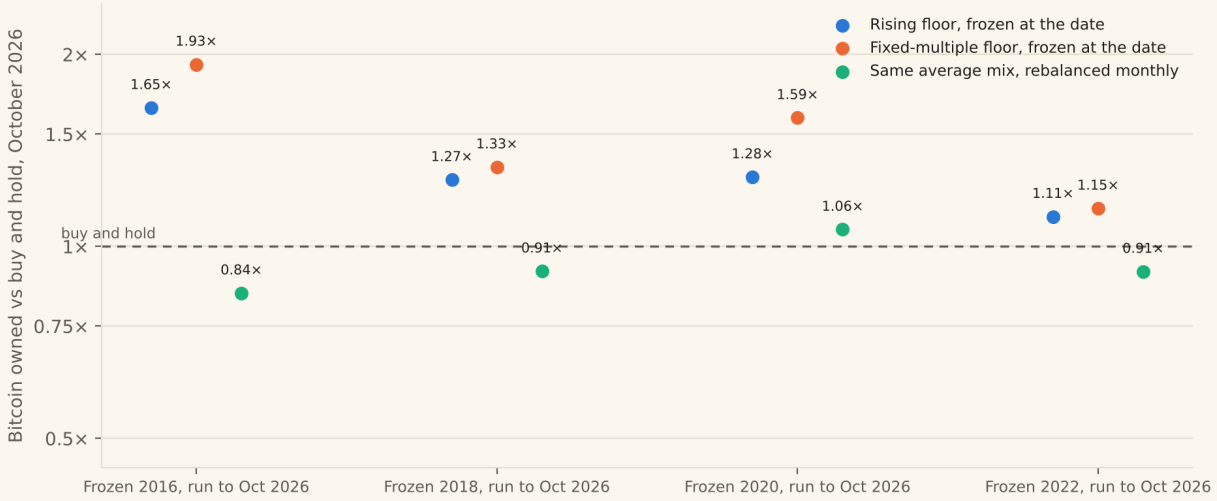
Table 5. Bitcoin stack vs hold over all start and end pairs (end at least four years after start).

Strategy	Pairs ahead of holding	Median	10th to 90th percentile
Floor-Kelly full, walk-forward (headline)	**98%**	1.34×	1.14× to 1.63×
Floor-Kelly full, published model (reference)	95%	1.41×	1.08× to 1.85×
Matched mix (headline weights)	9.7%	0.78×	0.66× to 1.00×
50/50, rebalanced yearly	12%	0.63×	0.39× to 1.05×
Floor-Kelly half, walk-forward	6%	0.47×	0.27× to 0.84×

3.4 Could an investor have done this at the time?

Figure 7. Parameters frozen in the past, never recalibrated, still beat buy and hold

Power law and floor fitted only on data up to December of each year, then held fixed. Strategy starts the next month.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 7. Bitcoin stack at October 2026 relative to buy and hold, with the power law and floor fitted only on data up to December of each year and never recalibrated. Note that both frozen floors beat holding from all four freeze dates, while the matched mix does so only from December 2020.

Table 6. Frozen-parameter test. Parameters are estimated only on data up to the freeze date, then held fixed to October 2026.

Frozen at	Power law fitted then (logA, β)	Rising floor slope	Rising floor: stack vs hold	Fixed-multiple floor: stack vs hold	Matched mix	Max drawdown, rising floor (hold)
Dec 2016	-16.303, 5.618	6.108	1.65×	1.93×	0.84×	-62% (-75%)
Dec 2018	-17.092, 5.873	6.121	1.27×	1.33×	0.91×	-62% (-73%)
Dec 2020	-16.682, 5.742	6.139	1.28×	1.59×	1.06×	-62% (-73%)
Dec 2022	-16.923, 5.819	5.993	1.11×	1.15×	0.91×	-42% (-48%)

This is the most direct answer to the hindsight concern. An investor who fitted the model in 2016, 2018, 2020 or 2022 and never touched it again would have ended with more Bitcoin than a holder by October 2026. The power-law corridor was published in 2019 (Burger, 2019), so the December 2020 and 2022 freezes are also out of sample for the model class.

3.5 What the statistics support

The 5,671 start and end pairs overlap heavily: they share the same three or four cycles. We therefore add three tests that do not rely on counting pairs.

Independent windows. In three non-overlapping four-year windows, walk-forward floor-Kelly ended ahead of holding in each one:

- **2014 to 2017:** 1.09×, with a drawdown of -60% against -74%.
- **2018 to 2021:** 1.50×, with -62% against -75%.
- **2022 to October 2026:** 1.32×, with -59% against -65%.

Three windows out of three is consistent with an edge, but three is a small number.

Bootstrap of the Bitcoin stack ratio. Resampling the monthly log-return differences in 12-month blocks, the 2014 to 2026 stack ratio of 2.16× has a 90% interval of 1.05× to 4.54×. 4.2% of draws end at or below 1×.

Table 7. Paired block bootstrap from January 2014, with a Holm correction across all 12 full-Kelly floors explored.

Floor variant	vs buy and hold: log growth per year	p	Holm-adjusted p	vs matched mix: log growth per year	p	Holm-adjusted p
Rising floor, walk-forward (headline)	+6.0%	0.042	0.46	+7.5%	< 0.001	0.002
0.42 floor, published model	+7.7%	0.069	0.69	+10.3%	< 0.001	0.002
0.42-type floor, refit	+5.4%	0.134	0.92	+7.9%	0.002	0.011
Expanding 1st-percentile floor, refit	+1.8%	0.394	1.00	+5.8%	0.032	0.16
0.42 floor, cash only	+6.4%	0.069	0.69	+9.9%	< 0.001	0.002
Rising floor, full sample	+7.3%	0.069	0.69	+9.8%	< 0.001	0.002
Rolling 8-year multiple	+6.6%	0.131	0.92	+9.9%	< 0.001	0.002
Rolling 8-year regression	+8.8%	0.023	0.28	+11.2%	< 0.001	0.002
200-week moving average	-5.4%	0.890	1.00	-3.3%	0.791	1.00
Quadratic log-time regression	-0.1%	0.506	1.00	+2.9%	0.204	0.82
Rolling 8-year exponential regression	-0.1%	0.858	1.00	-0.1%	0.767	1.00
On-chain realised price	-24.1%	0.999	1.00	-17.9%	0.978	1.00

A reported p of < 0.001 means fewer than 5 of 5,000 bootstrap draws went the other way. The first eight rows are power-law floors; the last four are the other shapes of Section 3.7. The moving-average floor holds 100% Bitcoin until April 2015, when it first has an anchor.

This table is the paper's statistical bottom line.

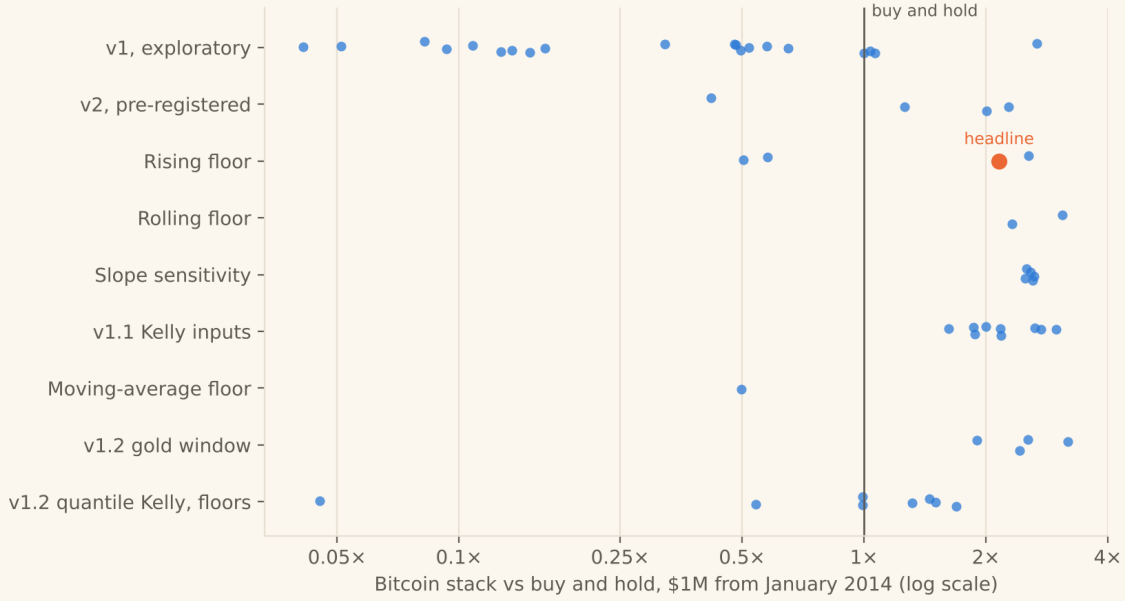
- **Against the matched mix,** seven of the eight power-law floors are significant after the correction, and none of the four other shapes. The floor timing adds value beyond rebalancing when the floor is a power law.

- **Against buy and hold,** no variant is significant after the correction. The accumulation result is consistent across every test we ran, but with three or four cycles it is not statistically established. The corrections for the full search below point the same way.

The full specification search. Table 7 corrects only across floor variants. The project also explored Kelly fractions, horizons, input types, slopes, covariance estimators, a moving-average floor and, in v1.2, the gold window, quantile inputs and three more floors. Figure 8 and Table 8 show all of them: 58 runs, of which two turned out to be identical (the full-sample rising floor and the frozen 5.954 slope), so 57 distinct configurations. For each one we computed monthly returns from January 2014 to October 2026 (154 months) and applied two corrections for the search.

Figure 8. Every configuration explored, not only the one reported

All 57 distinct configurations run for F007 and Paper 19, by phase. 32 of 57 end with more Bitcoin than holding. The headline ranks 11th of 57 on its edge over holding. Points are jittered vertically within each phase.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_v12_figures.py).

Figure 8. Bitcoin stack relative to buy and hold from January 2014 for every configuration explored, by research phase. Note that the headline sits inside the cluster of floor-Kelly variants, not at its top.

Table 8. The specification search: all 57 distinct configurations, $1M from January 2014.

Phase	Configurations	Ended with more Bitcoin than holding	Stack vs hold, range
v1, exploratory: input types, Kelly fractions and horizons on the 0.42 floor	20	3	0.04× to 2.67×
v2, pre-registered: refit floors and a cash-only variant	4	3	0.42× to 2.27×
Rising floor, full and half Kelly (includes the headline)	4	2	0.50× to 2.55×
Rolling eight-year floors	2	2	2.32× to 3.09×
Slope sensitivity (5.954 counted under the rising floor)	5	5	2.50× to 2.63×
v1.1 Kelly inputs: covariance, solver, horizon, costs	9	9	1.62× to 2.99×
v1.1 moving-average floor	1	0	0.50×
v1.2 gold window	4	4	1.90× to 3.19×
v1.2 quantile inputs and alternative floors	8	4	0.05× to 1.69×
All	**57**	**32**	**0.04× to 3.19×**

The headline's stack of 2.16× ranks 19th of the 32 configurations that ended ahead. On its monthly information ratio against holding (0.34 annualised) it ranks 11th of 56 comparable configurations, and on its Sharpe ratio 32nd of 57.

Deflated Sharpe ratio. The deflated Sharpe ratio (Bailey and López de Prado, 2014) is the probability that a strategy's true Sharpe ratio exceeds the best one that *N* independent trials would produce by luck alone. It corrects for the skewness and fat tails of the strategy's returns. We take the spread of Sharpe ratios across all distinct configurations and vary only *N*, because the configurations are far from independent: their active returns have a mean pairwise correlation of 0.76. We use the conventional threshold of 0.95.

Table 9. Deflated Sharpe ratio of the headline, by the assumed number of independent trials.

Independent trials *N*	2	5	10	20	All distinct
Excess return over T-bills (57)	0.9998	0.9994	0.9991	0.9986	0.9977
Edge over buy and hold, log-return difference (56)	0.77	0.48	0.31	0.19	0.09
Edge over the matched mix, log-return difference (56)	0.97	0.88	0.76	0.62	0.43

The count in brackets is the number of configurations in each row. Historical-average full Kelly holds 100% Bitcoin throughout, so it equals both buy and hold and its own matched mix, and it drops out of the last two rows.

- **Against T-bills,** the headline survives any correction. That says little: any portfolio that holds mostly Bitcoin does.
- **Against buy and hold,** the deflated Sharpe ratio is below 0.95 for every *N*: 0.77 even if the search counts as only two independent trials. An eigenvalue count of the correlation matrix gives 1.6 effective trials, but we do not rely on it: correlated return series can still reflect different choices. The edge over holding is not established.
- **Against the matched mix,** each configuration is compared with a constant mix at its own average weights. The headline's edge passes 0.95 only at *N* = 2 (0.97) and falls to 0.43 over the whole search, so the headline alone does not establish it. The evidence for the timing edge is that it is shared across the search: 53 of the 56 configurations beat their own matched mix.

Probability of backtest overfitting. Combinatorially symmetric cross-validation (Bailey, Borwein, López de Prado and Zhu, 2017) splits the 154 months into 16 blocks and forms all 12,870 ways of choosing half of them as the in-sample period. In each split it selects the configuration with the best in-sample information ratio against holding and records its rank out of sample.

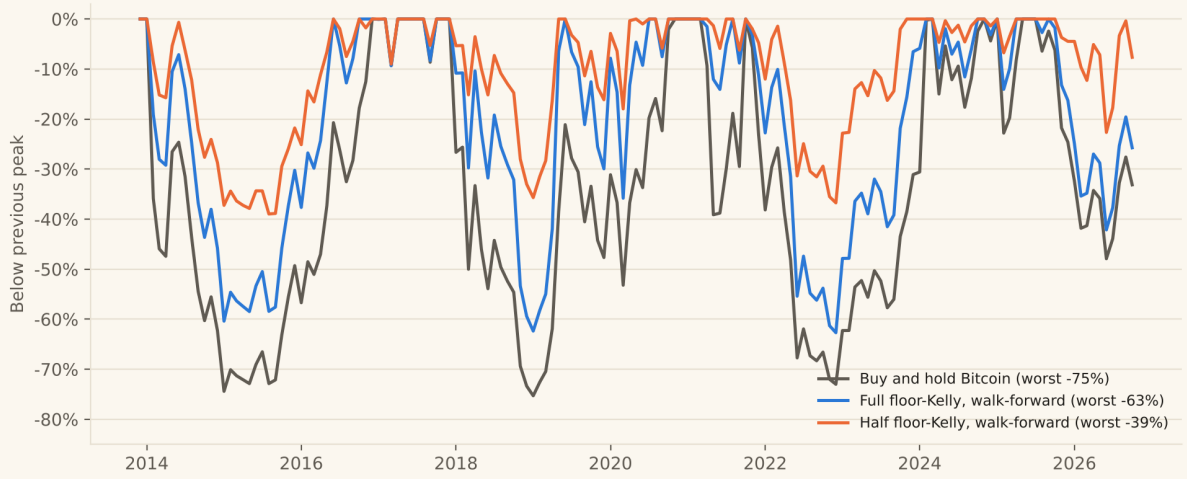
- **Against holding:** the selected configuration ranked below the out-of-sample median in 19% of splits (PBO = 0.19), and it lost to holding out of sample in 11%.
- **Against the matched mix:** PBO is 0.32, and the selected configuration lost to its matched mix out of sample in 6% of splits.

Selecting a configuration on part of this history and applying it to the rest mostly works: PBO is well below the 0.5 of a random pick. But the halves share the same three or four cycles, so this measures stability within the sample, not success on a new cycle. It does not change the deflated Sharpe result above.

3.6 Drawdowns and the risk-return trade-off

Figure 9. Drawdowns: full floor-Kelly 12 points shallower than holding, half Kelly half as deep

Decline from the previous portfolio peak, $1M from January 2014.

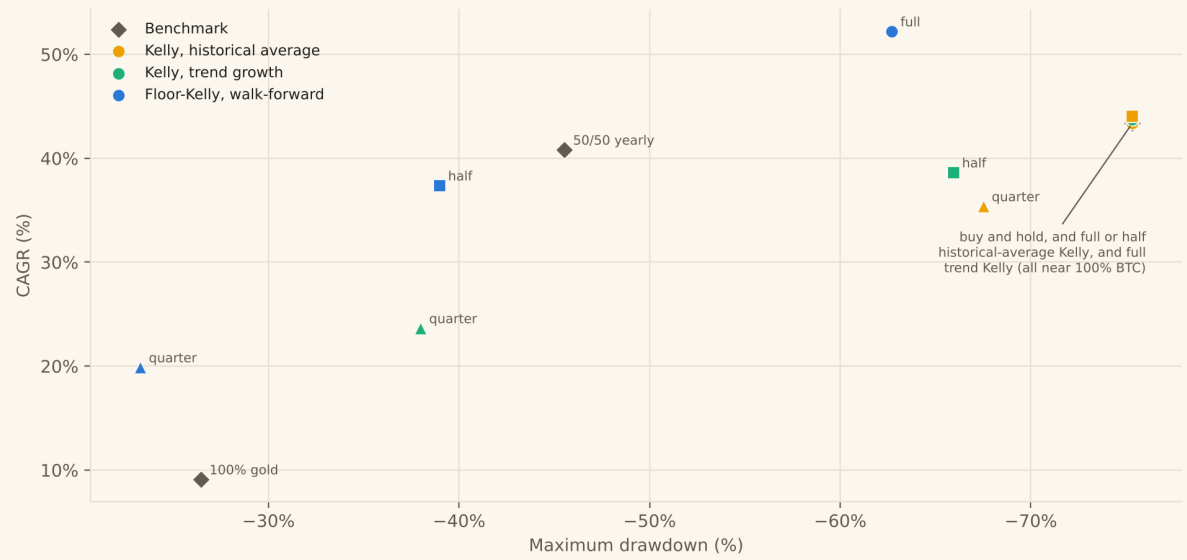


Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 9. Decline from the previous peak, $1M from January 2014. Note that walk-forward full floor-Kelly's worst drawdown (-63%) is 12 points shallower than holding's (-75%), and half floor-Kelly's (-39%) is about half as deep.

Figure 10. Walk-forward floor-Kelly sits on the upper edge of growth for the drawdown taken

CAGR against maximum drawdown, January 2014 to October 2026. Markers: circle full, square half, triangle quarter Kelly.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

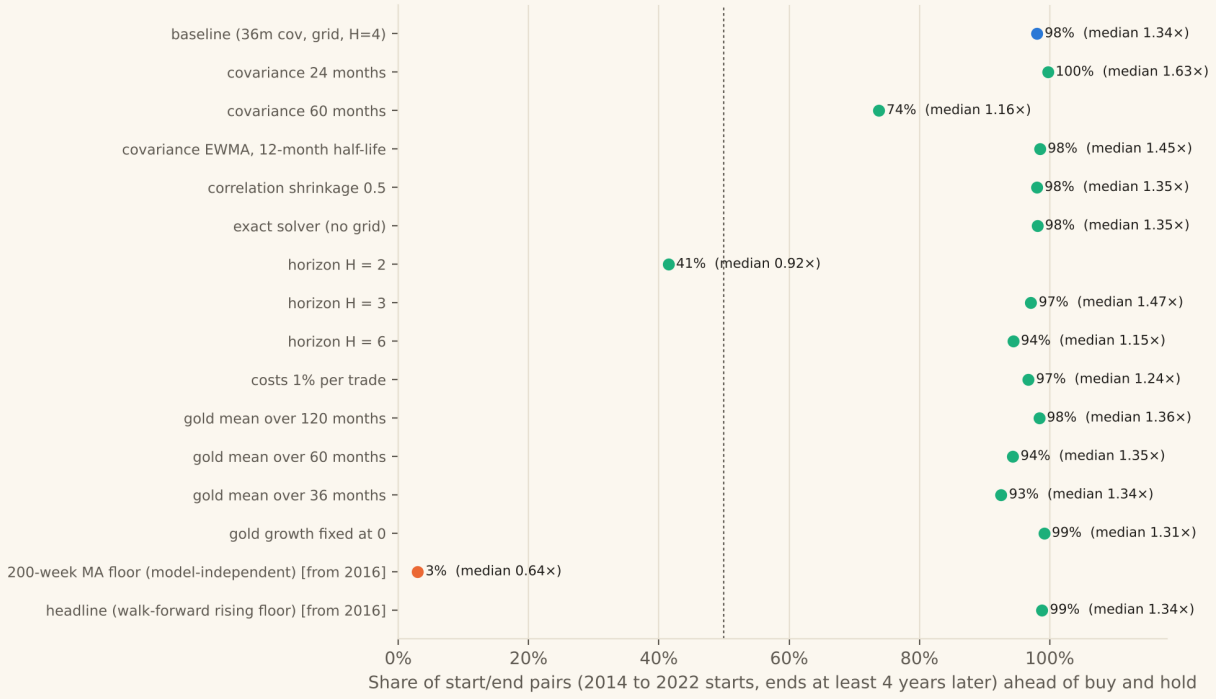
Figure 10. CAGR against maximum drawdown from January 2014. Note that historical-average Kelly at full and half, and trend Kelly at full, sit on top of buy and hold.

Figure 9 shows the drawdowns and Figure 10 the trade-off. From 2014, walk-forward floor-Kelly at full Kelly had a CAGR of 52.2% with a -63% drawdown, against 43.4% and -75% for holding. At half Kelly it had 37.4% with -39%. Historical-average Kelly at full and half Kelly matched holding's -75% drawdown at the same CAGR.

3.7 Robustness: Kelly inputs and a model-free floor

Figure 11. Robust to most Kelly inputs; not to a 2-year horizon or a model-free floor

Walk-forward rising-floor Kelly with one input changed at a time (blue: baseline; the baseline gold mean is expanding).
Orange: the 200-week moving-average floor, compared with the headline on the same 2016 to 2022 starts.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 11. Share of start and end pairs ahead of buy and hold when one input of the walk-forward strategy is changed. Note the two sensitive inputs, a two-year horizon and a 60-month covariance window, and the failure of the moving-average floor.

Table 10. Sensitivity of walk-forward full floor-Kelly (starts 2014 to 2022, all pairs).

Change	Pairs ahead	Median	Stack vs hold from 2014	Max drawdown from 2014
None (36-month covariance, expanding gold mean, grid, H = 4)	98%	1.34×	2.16×	-63%
Covariance over 24 months	100%	1.63×	2.99×	-61%
Covariance over 60 months	**74%**	1.16×	1.88×	-60%
EWMA covariance, 12-month half-life	98%	1.45×	2.64×	-62%
Correlation shrinkage, intensity 0.5	98%	1.35×	2.17×	-63%
Exact solver instead of the 1% grid	98%	1.35×	2.18×	-63%
Horizon H = 2 years	**41%**	0.92×	2.00×	-55%
Horizon H = 3 years	97%	1.47×	2.73×	-59%
Horizon H = 6 years	94%	1.15×	1.62×	-71%
Costs of 1% per trade	97%	1.24×	1.88×	-63%
Gold mean over the last 120 months	98%	1.36×	2.42×	-65%
Gold mean over 60 months	94%	1.35×	2.54×	-62%
Gold mean over 36 months	93%	1.34×	3.19×	-62%
Gold expected growth fixed at 0	99%	1.31×	1.90×	-65%

Two inputs matter.

- **A two-year horizon** makes the floor input too pessimistic near trend. The rule holds 62% Bitcoin on average, against 84% at four years, and falls behind in most pairs. The four-year horizon matches the halving cycle. It was the pilot default in the design note (`research/F007-kelly-btc-gold-design.md`) before any backtest, but this sensitivity has to be reported.
- **A 60-month covariance** window reacts slowly to falling volatility and so holds less Bitcoin (80% on average).

Every other change leaves the share of pairs ahead between 93% and 100%. The size of the edge moves more: 1% costs cut the 2014 stack from 2.16× to 1.88× (that row charges the holder 1% too; against a holder paying 0.25% it is 1.87×), and a 24-month covariance raises it to 2.99×.

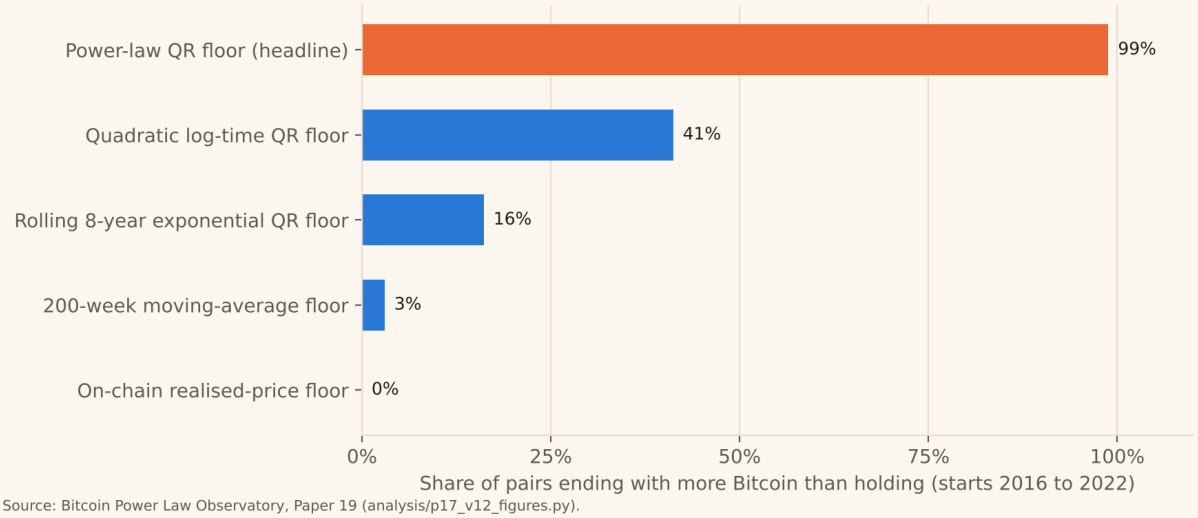
The gold input. The blind replication (Section 3.10) found that v1.1 did not state the window for gold's expected growth. The window barely changes the distribution over pairs: 93% to 99% ahead, with medians of 1.31× to 1.36×. It does move single end points. From 2012 the stack ranges from 1.26× with the expanding mean (published) to 1.95× with a 36-month mean, which holds more gold after gold rallies (21% average gold weight, against 13%). The published expanding mean gives the lowest result of the four estimated windows from both 2012 and 2014. Only fixing gold's growth at zero is lower from 2014 (1.90×).

Other valuation models. v1.1 tested one floor without the power law, the 200-week moving average. v1.2 adds three more. All are refit each month on data up to the decision date and fed through the same Kelly machinery.

- **Quadratic log-time floor:** quantile regression at $\tau = 0.01$ of $\log_{10}$ price on $\log_{10}$ days and its square. It contains the power law as a special case and lets the growth path bend.
- **Rolling eight-year exponential floor:** $\tau = 0.01$ regression of $\log_{10}$ price on time, not log time, over the last 96 months. This is the shape that a constant growth rate implies.
- **On-chain realised-price floor:** the floor multiple is the 1st percentile of MVRV since 2012, up to the decision date. Realised price is projected forward at its trailing 12-month growth.
- **200-week moving-average floor (v1.1):** the 1st percentile of price relative to the 200-week average since 2012, projected at the average's trailing 12-month growth.

Figure 12. Only the power-law floor produces the edge

Five walk-forward floors, each refit monthly on data up to the decision date, with the same Kelly machinery. 3,403 pairs each. The exponential floor holds about 100% Bitcoin, so it ends at holding minus costs.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_v12_figures.py).

Figure 12. Share of start and end pairs, starts 2016 to 2022, ending with more Bitcoin than buy and hold for five walk-forward floors. Note that only the power-law floor is ahead in most pairs.

Table 11. Walk-forward floors compared, full Kelly, starts 2016 to 2022 (3,403 pairs).

Floor	Pairs ahead of holding	Median stack	Stack vs hold from 2016	Average BTC weight from 2016
Power-law QR floor (headline)	**99%**	**1.34×**	**1.62×**	**85%**
Quadratic log-time QR	41%	0.95×	0.96×	76%
Rolling eight-year exponential QR	16%	1.00×	0.99×	100%
200-week moving average	3%	0.64×	0.62×	79%
On-chain realised price	0%	0.15×	0.06×	60%

Only the power-law floor produces the edge, and the failures are informative.

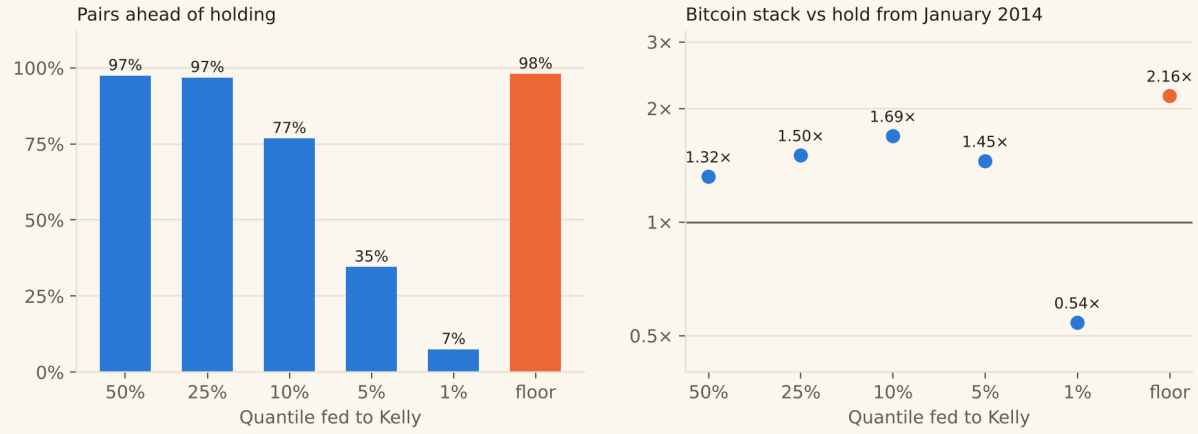
- **The exponential floor** extrapolates eight years of fast growth. The floor four years ahead sits far above the price, so the rule holds about 100% Bitcoin and ends at buy and hold minus costs.
- **The moving-average and realised-price floors** project trailing growth. That growth is highest after a bull market and lowest after a crash, so the input is most optimistic near peaks and most pessimistic near bottoms. The timing is reversed, and the realised-price rule ends with a small fraction of a holder's Bitcoin.
- **The quadratic floor** contains the power law as a special case but adds a curvature term. Refit monthly, it ends behind holding in most pairs. That makes the result specific to the power-law shape, and also fragile to it: a small, data-chosen bend in the growth path removes the edge, though it leaves the rule close to holding (median 0.95×) rather than far behind. Paper 20 finds that bending models forecast worse than the straight line at most freeze dates (C7.5), so the data do not favour the bend either.

The edge is therefore neither generic mean reversion nor generic valuation. It depends on a floor whose four-year growth comes from a stable long-run shape, and so far only the power law supplies one. This is the paper's most important model dependence. If the power law's shape is wrong, the evidence here gives the rule no edge.

Kelly on a chosen quantile. Section 2.3 placed the floor near the 14th percentile of the walk-forward conditional forecast. Feeding Kelly that forecast's quantiles directly gives a family of rules with an explicit level of caution (Figure 13, Table 12). These rules were specified after the floor results were known, so they test the interpretation rather than offer an independent edge.

Figure 13. Kelly fed with a low quantile works down to about the 10th percentile

Floor-anchored input = growth to the q-quantile of the walk-forward conditional price distribution four years ahead.
Full Kelly, 5,671 pairs (starts 2014 to 2022). The floor sits near the 14th percentile on average.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_v12_figures.py).

Figure 13. Kelly fed with growth to a chosen quantile of the walk-forward conditional price forecast four years ahead, against the floor. Note that the Bitcoin stack from 2014 peaks at the 10th percentile, while the share of pairs ahead falls steeply below the 25th.

Table 12. Conditional-quantile Kelly, full Kelly, starts 2014 to 2022 (5,671 pairs).

Growth to	Pairs ahead of holding	Median stack	Stack vs hold from 2014	Max drawdown from 2014	Average BTC weight from 2014
Median (50th percentile)	97%	1.10×	1.32×	-75%	98%
25th percentile	97%	1.15×	1.50×	-69%	93%
10th percentile	77%	1.11×	1.69×	-63%	79%
5th percentile	35%	0.90×	1.45×	-60%	67%
1st percentile	7%	0.43×	0.54×	-56%	44%
Floor (headline)	**98%**	**1.34×**	**2.16×**	**-63%**	**84%**

- **The median input** behaves like holding with a small edge: it holds 98% Bitcoin on average.
- **The 25th and 10th percentiles** add caution and, from 2014, more Bitcoin.
- **Below the 10th percentile** the input becomes too pessimistic and the rule falls behind.
- **The floor beats every fixed quantile.** Its quantile is not fixed: it moves between 3% and 26% (Figure 2), lowest after the 2013 bubble and highest in 2017. Part of the floor's edge may come from that variation, which this simple AR(1) forecast does not capture.

Every Kelly input under identical assumptions. Table 13 runs each input through the same engine from January 2014: 36-month covariance, the $\sigma^2/2$ term, the expanding gold mean, the 1% grid, long only and 0.25% costs. The robust family is represented by the quantile inputs; distributionally robust and drawdown-constrained Kelly were not run.

Table 13. Kelly inputs and fractions compared, $1M from January 2014 to October 2026.

Input	Kelly fraction c	Stack vs hold	CAGR	Max drawdown	Sharpe
Buy and hold		1.00×	43.4%	-75%	0.81
Historical average (textbook plug-in)	1	1.00×	43.4%	-75%	0.81
Historical average	0.5	1.06×	44.0%	-75%	0.82
Historical average	0.25	0.48×	35.4%	-67%	0.77
Trend growth (β = 5.688, fitted on the full history)	1	1.03×	43.7%	-75%	0.81
Trend growth	0.5	0.65×	38.6%	-66%	0.81
Trend growth	0.25	0.15×	23.6%	-38%	0.79
Conditional 25th percentile	1	1.50×	48.0%	-69%	0.87
Conditional 10th percentile	1	1.69×	49.3%	-63%	0.92
Floor-anchored (headline)	**1**	**2.16×**	**52.2%**	**-63%**	**0.93**
Floor-anchored	0.5	0.58×	37.4%	-39%	0.95
Floor-anchored	0.25	0.10×	19.9%	-23%	0.88

- **Textbook Kelly is buy and hold.** The historical average says 100% Bitcoin at full Kelly. At half Kelly it ends 1.06× ahead with the same −75% drawdown; at quarter Kelly, 0.48×.
- **Lower fractions cut drawdowns** and, except half historical-average Kelly, end with fewer Bitcoin. No fractional rule ends more than 1.06× ahead.
- **Among full-Kelly rules, the floor input gives the largest stack (2.16×) and the best Sharpe ratio (0.93).** Half floor-Kelly has the best Sharpe ratio overall (0.95) and a −39% drawdown, but ends with 0.58× a holder's Bitcoin.

3.8 Forward stress tests

Figure 14. Forward test: the edge needs cycles; a rising floor helps if they shrink

1,000 simulated 20-year paths per world from October 2026, full Kelly. Dot: median. Bar: 10th to 90th percentile (cut at 0.02×).

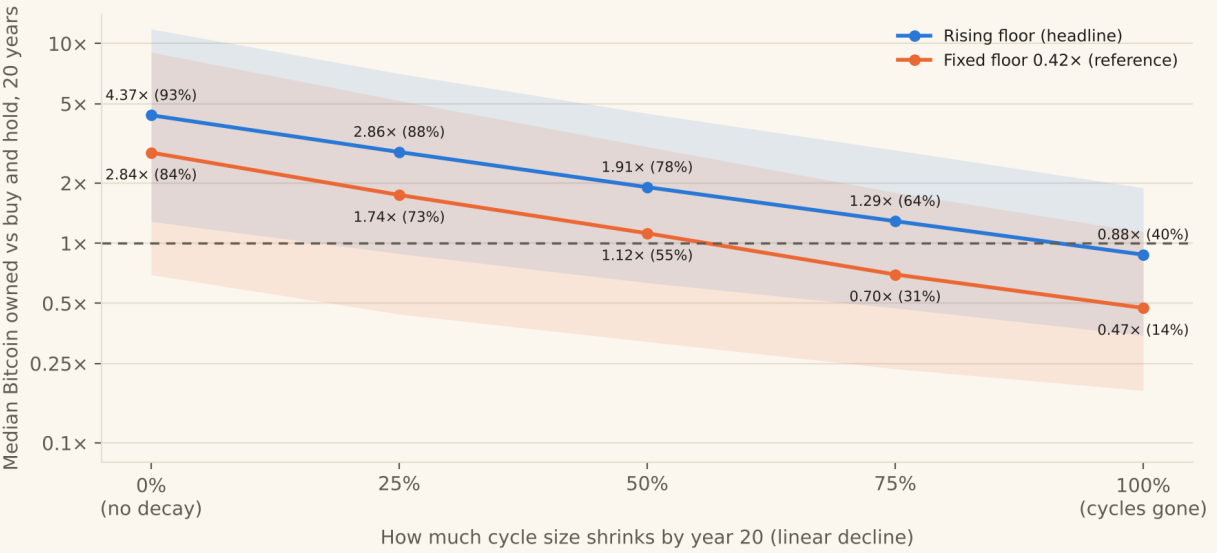


Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 14. Bitcoin stack after 20 years relative to buy and hold, full Kelly with the floor frozen at its October 2026 fit, 1,000 paths per world from October 2026. The dot is the median and the bar spans the 10th to 90th percentile.

Figure 15. Where the edge runs out: gradually shrinking cycles

Cycle amplitude declines linearly over 20 years. Labels: median (share of 1,000 paths ahead of buy and hold). Band: 10th to 90th percentile.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 15. Median Bitcoin stack after 20 years when cycle amplitude declines linearly. Note that the frozen rising floor stays ahead until cycles have shrunk by roughly 90%. The fixed floor breaks even at about 55%.

Figure 16. The strongest case against the rule: a lasting monetary repricing

Residual mean lifted twofold above its usual level (to about 1.84× trend) over six months. Same innovations and seed as Figure 14. Full Kelly, 1,000 paths.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 16. Bitcoin stack after 20 years when the price re-rates twofold above its usual level around trend (to about 1.84× trend), temporarily (10 years) or permanently.

Table 14. Forward Monte Carlo, full Kelly. Each cell gives the median stack vs hold and the share of paths ahead.

World	What happens	Rising floor, frozen at October 2026	Rising floor, refit in path (v1.3)	Fixed 0.42 floor (reference)
BASE	Cycles as large as 2012 to 2026	4.37×, 93%	1.25×, 55%	2.84×, 84%
Gradual decay to 75%	Cycle size falls linearly to 75% by year 20	2.86×, 88%	0.88×, 46%	1.74×, 73%
Gradual decay to 50%	… to 50%	1.91×, 78%	0.66×, 39%	1.12×, 55%
Gradual decay to 25%	… to 25%	1.29×, 64%	0.54×, 29%	0.70×, 31%
Gradual decay to 0%	Cycles disappear by year 20	0.88×, 40%	0.44×, 19%	0.47×, 14%
DAMPED	Cycle size halves immediately	0.92×, 45%	0.82×, 40%	0.45×, 16%
BROKEN	Price grows at half the trend rate, cycles continue	8.1×, 100%	5.54×, 93%	9.0×, 100%
NO_REVERSION	Random walk around the trend	1.24×, 57%	0.19×, 33%	1.22×, 56%
REPRICE_10Y	Price lifted twofold (about 1.84× trend) for 10 years, then reverts	1.93×, 75%	1.32×, 60%	1.15×, 55%
REPRICE_PERMANENT	Price lifted twofold (about 1.84× trend) for good	**0.63×, 36%**	0.58×, 32%	0.33×, 17%
REGIME_SWITCH (v1.2)	Markov switching between high-growth, mature and reserve-asset states	**1.07×, 53%**	0.71×, 37%	0.53×, 28%
FAT_TAILS_T3 (v1.2)	Student t(3) innovations, same variance	5.64×, 97%	2.31×, 73%	3.83×, 90%
STRUCTURAL_BREAK (v1.2)	At a random date, trend growth and cycle size both halve	2.52×, 96%	1.67×, 70%	2.33×, 91%
LOW_VOL_SPELL (v1.2)	Cycles shrink to 30% for five to ten years, then return	1.76×, 72%	0.57×, 35%	0.93×, 47%

The BASE world draws innovations from 2012 to 2026 and likely overstates future volatility. Its median buy-and-hold drawdown of -92% is deeper than any drawdown in the last decade.

The frozen column is not the backtested rule. The backtested rule, and the FULL prospective track, refit the floor every month. Refit inside each simulated path (every 12 months, on monthly closes), the rule does much worse: 1.25× with cycles as large as history, with 55% of paths ahead, and behind holding at the median in 9 of the 14 worlds. Section 3.9 explains why an estimated floor hurts going forward. The rest of this section describes the frozen floor.

Three worlds define where the frozen-floor rule stops working.

- **Cycles fade:** the frozen rising floor stays ahead until cycle amplitude has shrunk by about 90%.
- **Cycles become a random walk:** the outcome is a coin flip with a fat left tail.
- **Bitcoin re-rates permanently above its trend:** the rule sells Bitcoin into a repricing that never reverses, and ends with about two-thirds of a holder's Bitcoin.

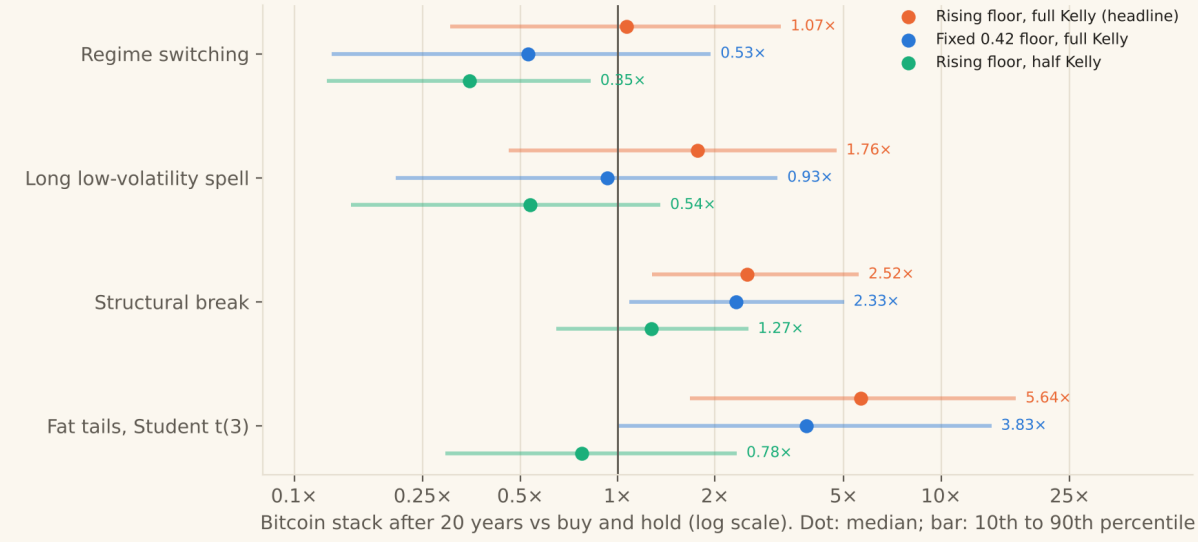
This last world is the strongest argument for buy and hold, and Section 4.4 addresses it.

Four more worlds (v1.2). The second review asked for conditional simulations with regime changes. We built four more worlds, each with 1,000 paths, from the same start point, bootstrap and seed.

- **REGIME_SWITCH:** a three-state Markov chain (Hamilton, 1989). A high-growth state lifts the residual by 0.10 in $\log_{10}$ (a 1.26× lift, to about 1.16× trend) with full-size cycles. A mature state has half-size cycles, and a reserve-asset state quarter-size cycles. Paths start in the high-growth state; the monthly probability of staying in a state is 98% to 99%. Paths spend 36%, 39% and 24% of the time in the three states.
- **FAT_TAILS_T3:** innovations from a Student t distribution with three degrees of freedom, scaled to the historical variance.
- **STRUCTURAL_BREAK:** at a random month between years 2 and 15, trend growth halves and cycle size halves.
- **LOW_VOL_SPELL:** cycle size falls to 30% for five to ten years, starting between years 1 and 10.

Figure 17. Four more ways the future could differ from the past

1,000 paths per world from October 2026, same seed and bootstrap as Figure 14. Regime switching alternates high-growth, mature and reserve-asset states; the break halves trend growth and cycle size at a random date.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_v12_figures.py).

Figure 17. Bitcoin stack after 20 years relative to buy and hold in the four v1.2 worlds, for the headline, the fixed 0.42 floor and half Kelly. Note the regime-switching world, where the headline's median sits near 1×.

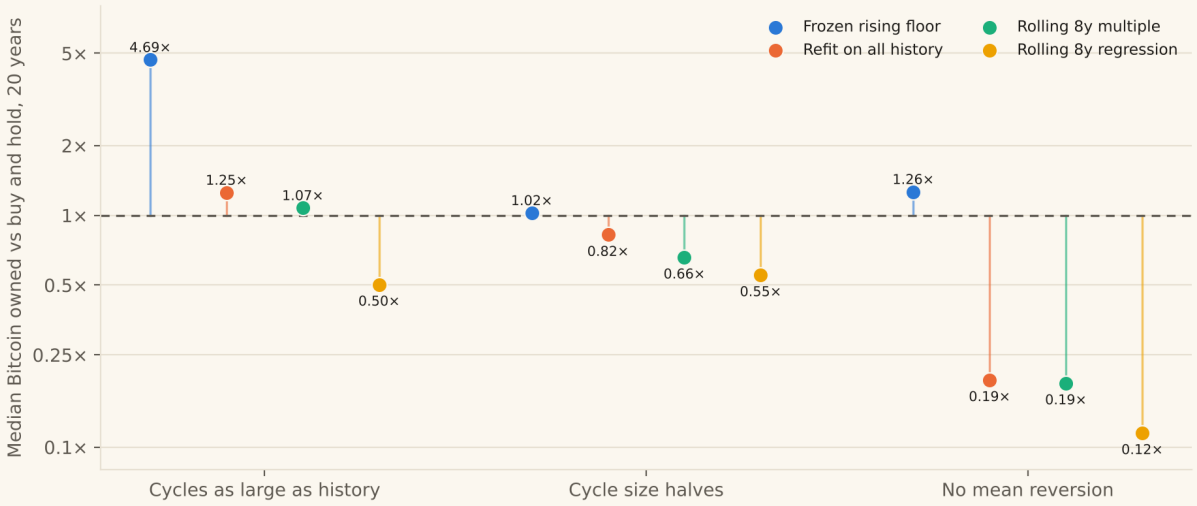
- **Regime switching is a coin flip.** The rule sells Bitcoin into the high-growth spells, when the price sits above trend, and the buy-back level does not always return. The median is 1.07×, with 53% of paths ahead and a 10th to 90th percentile of 0.30× to 3.20×.
- **Fat tails help the rule.** Deeper crashes leave more to buy back: the median is 5.64×.
- **A structural break leaves the rule ahead** (2.52×). Before the break it trades full-size cycles. After it, slower trend growth brings the price toward the rising floor, *g* rises, and the rule holds more Bitcoin (74% on average, against 55% in BASE).
- **A long calm spell costs part of the edge** (1.76×, 72% of paths ahead). The fixed floor falls behind (0.93×).
- **Half Kelly ends behind holding in three of the four worlds** (0.35× to 0.78×, and 1.27× after a break).

3.9 What the simulations say about specifying the floor

The headline's rising floor is re-estimated each month in the backtest. In the simulation, re-estimating the floor inside each path makes results worse in every world (Figure 18; for the headline rule, behind holding at the median in 9 of 14 worlds, Table 14). The deepest lows dominate a 1st-percentile fit, so an estimated floor falls after crashes and cuts Bitcoin when buying matters most. Figure 18 compares the frozen floor with an expanding regression, a rolling eight-year multiple and a rolling eight-year regression.

Figure 18. Re-estimating the floor makes it worse in every world

Each simulated path refits its floor from its own prices. All floors use monthly closes here, so the frozen floor reads 4.69× rather than Figure 14's 4.37× (daily fit). Full Kelly, 1,000 paths. Stalks run from 1× to each median.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 18. Median Bitcoin stack after 20 years when each simulated path re-estimates its own floor. Note that every re-estimated floor does worse than the frozen floor.

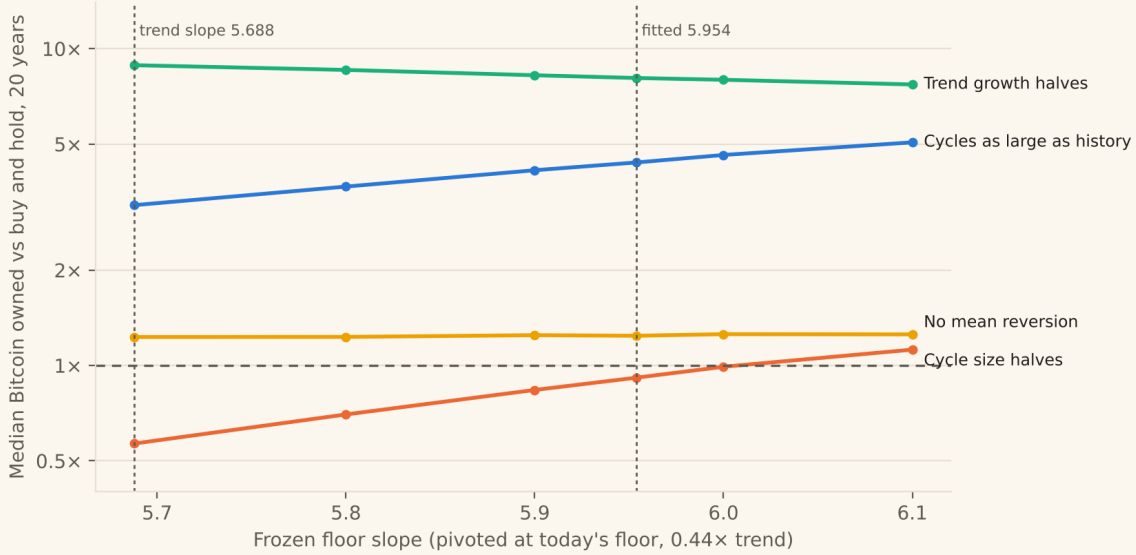
The rolling eight-year regression shows the risk of selection on history:

- **Historically it looked best:** 98% of pairs ahead, and +8.8% per year over holding (unadjusted p = 0.023).

- **Its slope is unstable:** it swings between 4.8 (April 2025) and 6.6 (February 2020) over the decision months from December 2013.
- **Forward, it is the worst floor in every simulated world.**

Figure 19. History cannot pick the slope; the future depends on it

Historically every slope ends ahead in 95% to 96% of start and end pairs. Forward, a steeper floor holds more Bitcoin.



Source: Bitcoin Power Law Observatory, Paper 19 (analysis/p17_figures.py).

Figure 19. Median Bitcoin stack after 20 years for a floor frozen at today's level (0.44× trend), by slope. Note the flat historical result against the steep forward dependence.

History cannot choose the slope: every slope from 5.688 to 6.1 ends ahead in 95% to 96% of historical pairs. Forward, the slope sets how much Bitcoin the rule commits to. With cycles halved, a steeper slope raises the median from 0.57× to 1.12×, with break-even near 6.0. Figure 19 shows that dependence. The case for a floor that converges toward trend rests on the Observatory's quantile-regression evidence (F002, scan S010), not on this backtest.

Paper 20 tests frozen floors directly. The 1% line is not calibrated out of sample, and the trend times its lowest-1% multiple is the most stable floor tested (C7.2, C7.3). The rule still beat holding from almost every freeze date with both floors, because it uses the floor's shape and swings, not its level.

3.10 Blind replication

An independent agent re-implemented the method from the v1.1 text and the four data files, without access to our code or results. It reproduced every quantity it tested to rounding:

- Table 1;
- Table 2, with an end value of $20.61B against $20.6B and a stack of 1.263× against 1.26×;
- all seven turning points of Table 3 and the start years of Table 4;
- all four frozen tests of Table 6, including the fitted slopes;
- the 5,671-pair distribution of Table 5 (98.0% ahead, median 1.344×).

It found one material gap. v1.1 said gold's expected growth was its "trailing sample mean" without giving the window. Read as 36 months, the 2012 headline becomes 1.95× instead of 1.26×, and only the expanding mean reproduces the paper. Section 2.5 now states it, and Section 3.7 reports every window.

The replication also listed conventions it had to infer: the start date, the cost base, early covariance, cash timing, the pair rule and the start of the frozen tests. Sections 2.5, 2.7 and 2.8 now state each one. It did not test the bootstrap, the Holm correction or the Monte Carlo.

The replicator was a separate AI agent working under a blindness protocol, not a human team. Its work is an independent re-implementation, not an independent review of the method's assumptions.

---

4. Discussion

4.1 Why floor-anchored Kelly works

Three properties combine.

1. **The input is valuation-aware.** Near the floor g is large and Kelly leans in. Near and above trend g is small or negative and Kelly steps out.
2. **The input is pessimistic within the model.** It assumes the price lands on the floor rather than returning to trend: about the 14th percentile of the model's own four-year forecast (Section 2.3). Using a low quantile instead of the mean is a built-in guard against overestimating returns, the main weakness of Kelly in practice.
3. **Kelly converts the signal into size.** Variance enters the denominator, so the rule also scales back when volatility rises.

The matched mix shows that the combination, not rebalancing alone, creates the edge. The four alternative floors show that the power law's long-horizon shape is essential.

4.2 Levels of confidence

Claim	Evidence	Level
The floor timing adds value beyond the rebalancing premium	Matched-mix comparison significant for seven of the eight power-law floors after a Holm correction across all 12 floors; 98% vs 9.7% of pairs ahead; 53 of 56 configurations beat their own matched mix; the in-sample choice loses to its matched mix out of sample in 6% of splits. The headline's own deflated Sharpe is 0.43 to 0.97	**Strong** (historical, as a property of the method)
The edge requires the power-law floor	Starts 2016 to 2022: four walk-forward floors with other shapes ahead in 0% to 41% of pairs, against 99%	**Strong** (historical)
Full floor-Kelly accumulates more Bitcoin than buy and hold	Ahead from every start year, in all three independent windows and all four frozen-parameter tests; bootstrap interval 1.05× to 4.54×; Holm-adjusted p = 0.46; deflated Sharpe 0.09 to 0.77 (below 0.95 at every trial count); PBO 0.19	**Suggestive**
The edge continues	Requires cycles to persist; fails under a permanent repricing and when cycles vanish; a coin flip under regime switching; with the floor refit in each path, behind holding at the median in 9 of 14 simulated worlds	**Unproven hypothesis**, now tracked prospectively (Section 7)

4.3 Two audiences, two Kelly fractions

	Full floor-Kelly (walk-forward)	Half floor-Kelly (walk-forward)
Goal	Grow the Bitcoin stack	Control risk
Pairs ahead of buy and hold since 2014	98%	6%
Max drawdown, 2012 to 2020 starts	-63% to -66%	-37% to -43%
Sharpe ratio, 2012 to 2020 starts	0.69 to 1.08	0.74 to 1.11
Average Bitcoin weight	84% to 85%	45% to 48%
Median 20-year stack vs hold in the four v1.2 simulated worlds (floor frozen at October 2026)	1.07× to 5.64×	0.35× to 1.27×

4.4 The case for buy and hold

The strongest counterargument is that Bitcoin is not a mean-reverting asset around a stable trend but an asset undergoing a long monetary repricing. In that world, selling at high valuations sells an appreciating asset too early, and the buy-back level may never return. The REPRICE_PERMANENT world quantifies the cost: the rule ends with 0.63× the Bitcoin of a holder with the floor frozen (36% of paths ahead), and 0.58× with it refit in each path.

Buy and hold also carries no model risk and no execution risk. The DAMPED and NO_REVERSION worlds show the rule's other failure modes.

Our evidence does not settle this argument. It shows that, for the history we have, floor-anchored timing would have added Bitcoin, and that the same rule loses Bitcoin if the history stops repeating. A holder who believes in a permanent repricing should hold. A holder who believes the power law's cycles persist has historical support for the rule, but not proof.

4.5 Investor suitability: utility, drawdowns and liquidity

Full Kelly maximises expected log wealth. That is the right objective only for an investor with a relative risk aversion of one, a long horizon and no drawdown limit. Three adjustments matter for real investors.

1. **Risk aversion.** With constant relative risk aversion γ and lognormal returns, the optimal portfolio is the growth-optimal portfolio scaled by $1/\gamma$, which is our Kelly fraction c (Merton, 1971; MacLean, Thorp and Ziemba, 2010). Half Kelly corresponds to $\gamma = 2$, a common institutional calibration. On this reading the accumulation result belongs to log-utility investors. Bitcoin holders who measure wealth in Bitcoin and already accept -75% drawdowns are close to that profile; most institutions are not.
2. **Drawdowns.** Full floor-Kelly's worst drawdown was -59% to -66% across full-cycle starts: shallower than holding's, but beyond most mandates. Drawdown-constrained growth policies (Grossman and Zhou, 1993; Busseti, Ryu and Boyd, 2016) cap drawdowns more directly than a fixed fraction. We have not tested them here.
3. **Liquidity and execution.** The rule trades at most once a month, in spot Bitcoin and a gold ETF, two of the most liquid markets available. From 2012 the median month traded 8% of the portfolio, the 90th percentile month 31%, and 16% of months needed no trade. Quadrupling the cost assumption to 1% per trade leaves the result in place, with 97% of pairs ahead (Table 10). Market impact for very large portfolios is not modelled.

The low-quantile input also bears on the usual objection to full Kelly: Kelly with estimated means over-bets because estimation errors inflate m. The floor input already sits well below the model's central forecast. It is therefore less exposed to that error than a mean estimate, though not immune to an error in the model itself.

5. Limitations

1. **Thin history.** Bitcoin has three to four complete cycles. Start and end pairs overlap heavily. Only three independent four-year windows exist.

2. **Multiple testing.** 57 distinct configurations were explored (Table 8). After a Holm correction across all 12 full-Kelly floors, none beats buy and hold at the 5% level (Table 7). The deflated Sharpe ratio of the edge over holding is 0.09 to 0.77, below 0.95 however many independent trials the search represents (Table 9). The timing edge over the matched mix rests on its breadth across the search (53 of 56 configurations), not on the headline alone.
3. **Hindsight in the model class.** Walk-forward fitting and frozen-parameter tests remove hindsight in the parameters. They do not remove it in the choice of a power-law model, the quantile-regression floor, or the four-year horizon. Four floors with other shapes fail (Table 11), so the result rests on the power-law shape itself, which was chosen on the same history.
4. **The model's predictive power is limited.** The in-sample R^2 of the power law is 0.956, but the out-of-sample R^2 is 0.546 (S097). The model explains about 55% of variance on data it has not seen.
5. **The floor is a level, not a boundary.** Bitcoin has spent 3.7% of days since July 2010 below the published floor (0.42× trend), in 23 visits, none longer than 86 days; since 2016, 0.3% of days, with a longest visit of 7 days (C1.4). That line was fitted on the whole history, so this record describes the past. The headline floor does worse out of sample: the 1% quantile line, fitted at a date and carried forward, was below price on a median 14% of later days and held for four years at 5 of 94 freeze dates (Paper 20, C7.2). Visits below the published floor are expected. Since 2016 it was visited 6 times in 10.8 years, none for longer than 7 days (`analysis/p18_visits.json`). On a Poisson model at that rate, at least one visit within 30 years is close to certain, and a visit of a week or more has a probability of about 94% (one such visit since 2016). The Observatory's standard caveat R026 puts the 30-year breach probability at 85% under a Poisson model; it predates the floor definition, and its source analysis is not in this repository.
6. **Cycle decay is uncertain.** The Observatory's per-cycle decay estimate of about 20% has confidence intervals that span zero for each individual transition (S017). Which forward world applies is unknown.
7. **Stylised simulation worlds.** The Monte Carlo uses an AR(1) residual around the official trend, with block-bootstrapped innovations estimated from the same history. Its worlds are conditional stress tests, not forecasts. The forward results also depend on how the floor is maintained: frozen at its October 2026 fit, the rule does well while cycles persist; refit inside each path, it falls behind holding in most worlds (Table 14).
8. **Replication by an AI agent, not a human team.** A blind re-implementation reproduced every number it tested (Section 3.10). No human third party has yet reproduced or reviewed the results. All code and data are in the repository.
9. **Monthly granularity.** Decisions and returns are monthly, so intra-month moves and liquidity are not modelled.
10. **Gold proxy.** GLD's New York close differs in timing from Bitcoin's UTC close and from the London fix.
11. **Before tax.** All results are before tax, by design. At each cycle peak the rule cuts its Bitcoin weight by 44 to 61 percentage points (Table 3) and buys back near the floor. Where realised gains are taxed, this turnover could shrink or reverse the stack result. The effect depends on jurisdiction and account type, and we do not estimate it. Custody costs and market impact at size are not modelled.
12. **The quantile interpretation is model-based.** It rests on a normal AR(1) forecast fitted to the same history (Section 2.3).

13. **Prospective evidence is only starting.** The frozen record begins in November 2026 (Section 7) and will need at least one full cycle before it can say anything about accumulation.

6. Implications

- **For Bitcoin holders.** Floor-anchored Kelly is a disciplined way to use the cycles, with historical support for adding Bitcoin. It requires full Kelly, a four-year horizon and conviction that the cycles persist. Expect to trail a holder through every bull run.
- **For institutional allocators.** Half floor-Kelly is a risk-control rule: drawdowns of -37% to -43%, Sharpe ratios of 0.74 to 1.11, and an auditable valuation logic.
- **For Observatory tools.** The institutional page's `kellyAllocation()` (`shared/js/powerlaw.js`) treats a log growth rate as an arithmetic return, omitting $\sigma^2/2$. It also fixes σ at 0.75 while realised 36-month volatility is about 0.46. It should be corrected, or replaced by the walk-forward floor-anchored rule.
- **For research.** The decisive open question is forward: does the floor converge, and do cycles persist? The prospective record in Section 7 is the cleanest out-of-sample test available. Two further studies would sharpen the result: a drawdown-constrained Kelly rule, and a joint test of the power-law shape against other long-run growth models on more assets.

7. Prospective record

The cleanest test of the rule is one that history cannot contaminate. On 2026-11-06 we froze the rule in `research/p17_frozen_rule.json`: the walk-forward τ = 0.01 power-law floor, H = 4, 36-month covariance, expanding gold mean, 1% grid and 0.25% costs, at full and half Kelly.

At each month-end close, `analysis/p17_prospective.py` appends a row to `research/p17_prospective_record.csv`. Each row holds:

- the decision for each track;
- a $1M paper portfolio for each track and for buy and hold, all started at the October 2026 close;
- the inputs behind the decision;
- a SHA-256 hash of the frozen spec.

Rows are append-only, and the script refuses to add a row if the spec has changed. The spec was revised twice on the day it was frozen, both before any prospective month: once to state where the gold window starts, and once to add the TREND track below. The decision code did not change. The first row was recorded on 6 November 2026 from the 31 October close. Its `--check` mode, run before each monthly update, verifies that its decisions match the backtest engine month by month.

Table 15. First entry of the prospective record: decision at the October 2026 close, held through November 2026.

BTC close	Walk-forward floor today	Price / floor	Floor in four years	Floor growth g	Full Kelly	Half Kelly
$77,223	$70,158	1.10×	$234,230	+27.7% per year	100% Bitcoin	75% Bitcoin, 25% gold

A third track, TREND, was added the same day. Paper 20 found the trend line times the lowest-1% multiple of price over trend to be the most stable floor out of sample (C7.3), so TREND runs the same rule on that floor, fitted on data to October 2026 and frozen: a floor of $44,504, with Bitcoin 1.74× above it. Its first allocation is 82% Bitcoin and 18% gold (Table 15 shows the FULL and HALF tracks).

Future versions of this paper will report the record against criteria registered before any outcome was known (`research/p19_prospective_criteria.md`). A fair test needs at least one full cycle. Until the next cycle peak and trough, the record can show whether the rule behaves as specified, but not whether it adds Bitcoin.

8. Conclusion

Kelly sizing with floor-anchored growth is a counter-cyclical allocation rule that the power law makes possible.

- **Walk-forward:** full floor-Kelly ended with more Bitcoin than buy and hold in 98% of start and end month pairs since 2014, from every January start from 2012 to 2025, and with parameters frozen at four past dates.
- **The edge comes from the floor timing,** not from rebalancing. That claim survives a Holm correction across all 12 floors for seven of the eight power-law variants and holds in 53 of 56 configurations in the full search, though the headline's own deflated Sharpe ratio is only 0.43 to 0.97.
- **Ending with more Bitcoin than a holder is suggestive, not established.** It does not survive the multiple-testing corrections. It also depends on Bitcoin's cycles persisting: the rule loses Bitcoin if cycles vanish or if Bitcoin re-rates permanently above its trend. In simulation, with the floor refit inside each path (yearly, on monthly closes), it falls behind holding in most worlds.
- **The edge is a power-law result.** Four floors with other shapes fail, and the input behaves like a low quantile of a power-law forecast, though fixed quantiles do less well than the floor (Table 12).
- **Half Kelly** turns the same signal into a risk-control rule for institutions.
- **The test now moves forward.** A frozen, append-only record tracks the rule each month from November 2026.

Revision notes

v1.2 to v1.3

These changes respond to an independent pre-publication review (score 7.5 before revision) and to the companion Paper 20.

Review point or finding	Change in v1.3
Paper number collision	Published as Paper 19 (drafted as Paper 17); the companion floor paper is Paper 20
Companion paper not cited; floor wording predates the Observatory definition	Paper 20 cited in Sections 1.5, 2.3, 3.7, 3.9, 5 and 7; the floor described as a level, not a boundary (Section 2.1, Limitation 5)
Forward simulations used a frozen floor, not the backtested refit rule	New column in Table 14: refit inside each path, behind holding at the median in 9 of 14 worlds; Abstract, Key findings, Sections 3.8, 4.2, 5 and 8 updated
Holm family covered 8 of 12 floors explored	Table 7 now covers all 12; headline vs matched mix adjusted p = 0.002, vs holding 0.46; seven of eight power-law floors significant against the matched mix
No single Kelly comparison under identical assumptions	Table 13: historical-average, trend, quantile and floor inputs at full, half and quarter Kelly
“Expected growth” wording contradicted Section 2.3	“Floor-anchored growth” throughout; the rule described as a robust-Kelly-style sizing rule, with the exact interval-set link (Section 2.3)
Statistical verdict missing from Key findings	Key finding 3 added
Tax limitation terse	Limitation 11 states what before-tax results leave out
Prospective record had no pass or fail criterion	Criteria registered in `research/p19_prospective_criteria.md` before any outcome
Blind replication not archived	Report and code in `research/p19_blind_replication/`
Text mismatches and disclosures	Slope range 4.8 to 6.6; sensitivity range 93% to 100%; cost-row note; rank out of 56; January values in Section 2.2; matched-mix weighting; H = 4 pilot default; registry C6.4 updated

v1.1 to v1.2

These changes respond to a second external review of v1.1 and to a blind replication.

Review point or finding	Change in v1.2
Multiple-testing accounting incomplete	Full specification search of 57 distinct configurations, a deflated Sharpe ratio over a range of trial counts, and the probability of backtest overfitting (Section 3.5, Figure 8, Tables 8 and 9)
No theory for the floor-derived expected return	The floor input as a low quantile of a walk-forward conditional forecast, linked to robust Kelly (Section 2.3, Figure 2)
Experiment 1: probability-weighted Kelly	Kelly fed with the 50th to 1st percentiles of the conditional forecast (Section 3.7, Figure 13, Table 12)
More independent valuation models	Quadratic log-time, rolling exponential and on-chain realised-price floors; all fail (Section 3.7, Figure 12, Table 11)
Experiment 3: regime switching and conditional simulations	Regime-switching, fat-tail, structural-break and low-volatility worlds (Section 3.8, Figure 17)
Full Kelly too aggressive for institutions	Utility, drawdown and liquidity discussion (Section 4.5)
Experiment 2: prospective validation	Frozen rule and an append-only monthly record from November 2026 (Section 7)
Blind replication: gold window unstated	Stated (Section 2.5); every window reported (Section 3.7)
Blind replication: other inferred conventions	Start date, cost base, early covariance, cash timing, pair rule and frozen-test start stated (Sections 2.5, 2.7, 2.8)

v1.0 to v1.1

The following changes respond point by point to an external academic review of v1.0.

Review point	Change in v1.1
Parameter hindsight in the headline	The headline is now the walk-forward rising floor; the published-model floor is the reference (Sections 2.4, 3.1)
No implementable-in-real-time test	Frozen-parameter tests at four dates (Section 3.4, Figure 7)
Overlapping observations; p = 0.07 overstated	Independent windows, a bootstrap of the stack ratio, a Holm correction, and three explicit confidence levels (Sections 3.5, 4.2)
Kelly implementation needs justification and sensitivity	Derivation note; covariance windows, EWMA, shrinkage, exact solver, horizons and costs (Sections 2.5, 3.7)
Model dependence: test a model-independent signal	200-week moving-average floor; it fails (Section 3.7)
Repricing and gradual decay scenarios	Two repricing worlds and a gradual-decay sweep (Section 3.8, Figures 15 and 16)
“Conservative” only within the framework	Stated explicitly (Sections 1.2, 4.1)
Literature review too thin	Section 1.5 expanded, with every citation verified
Taxes	Out of scope by design (Section 2.7)

References

- Bailey, D. H., Borwein, J., López de Prado, M. and Zhu, Q. J. (2017). The probability of backtest overfitting. *Journal of Computational Finance*. doi:10.21314/JCF.2016.322.
- Bailey, D. H. and López de Prado, M. (2014). The deflated Sharpe ratio: correcting for selection bias, backtest overfitting and non-normality. *Journal of Portfolio Management*, 40(5), 94 to 107.
- Bouchey, P., Nemtchinov, V., Paulsen, A. and Stein, D. M. (2012). Volatility harvesting: why does diversifying and rebalancing create portfolio growth? *Journal of Wealth Management*, 15(2), 26.
- Burger, H. C. (2019). Bitcoin’s natural long-term power-law corridor of growth. Medium, Quantodian Publications.
- Busseti, E., Ryu, E. K. and Boyd, S. (2016). Risk-constrained Kelly gambling. *Journal of Investing*, 25(3), 118 to 134.
- Fernholz, R. and Shay, B. (1982). Stochastic portfolio theory and stock market equilibrium. *Journal of Finance*, 37(2), 615 to 624.
- Goldfarb, D. and Iyengar, G. (2003). Robust portfolio selection problems. *Mathematics of Operations Research*, 28(1), 1 to 38.
- Grossman, S. J. and Zhou, Z. (1993). Optimal investment strategies for controlling drawdowns. *Mathematical Finance*, 3(3), 241 to 276.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357 to 384.
- Harvey, C. R., Liu, Y. and Zhu, H. (2016). … and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5.
- Kelly, J. L. (1956). A new interpretation of information rate. *Bell System Technical Journal*, 35(4), 917 to 926.

- Kim, T. S. and Omberg, E. (1996). Dynamic nonmyopic portfolio behavior. *Review of Financial Studies*, 9(1), 141 to 161.
- Ledoit, O. and Wolf, M. (2004). Honey, I shrunk the sample covariance matrix. *Journal of Portfolio Management*, 30(4), 110 to 119.
- Liu, Y. and Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, 34(6), 2689.
- MacLean, L. C., Thorp, E. O. and Ziemba, W. T. (2010). Long-term capital growth: the good and bad properties of the Kelly and fractional Kelly capital growth criteria. *Quantitative Finance*, 10(7), 681 to 687.
- Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4), 373 to 413.
- Platanakis, E. and Urquhart, A. (2020). Should investors include Bitcoin in their portfolios? A portfolio theory approach. *British Accounting Review*, 52(4).
- Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428), 1303 to 1313.
- Rotando, L. M. and Thorp, E. O. (1992). The Kelly criterion and the stock market. *American Mathematical Monthly*, 99(10), 922 to 931.
- Rujeerapaiboon, N., Kuhn, D. and Wiesemann, W. (2016). Robust growth-optimal portfolios. *Management Science*, 62(7), 2090 to 2109.
- Santostasi, G. (n.d.). *The Physics of Bitcoin* (book). thephysicsofbitcoin.com, accessed 2026-11-06.
- Sun, Q. and Boyd, S. (2018). Distributional robust Kelly gambling. arXiv:1812.10371.
- Thorp, E. O. (2006). The Kelly criterion in blackjack, sports betting and the stock market. In *Handbook of Asset and Liability Management*, Vol. 1. Elsevier.
- Tütüncü, R. H. and Koenig, M. (2004). Robust asset allocation. *Annals of Operations Research*, 132, 157 to 187.
- Wachter, J. A. (2002). Portfolio and consumption decisions under mean-reverting returns. *Journal of Financial and Quantitative Analysis*, 37(1), 63 to 91.
- Coin Metrics. Community Network Data, series CapMVRVCur (MVRV ratio), retrieved 2026-11-06.
- Bitcoin Power Law Observatory. Paper 12: Formal Verification (out-of-sample R^2, S097).
- Bitcoin Power Law Observatory. Paper 18: The Formula, Month by Month. 184 Real-Time Refits of Bitcoin's Power Law, 2011 to 2026, v1.0.
- Bitcoin Power Law Observatory. Paper 20: Does a Frozen Power-Law Floor Hold? A Pre-Registered Out-of-Sample Test of Four Bitcoin Floor Definitions, 2011 to 2026, v1.2.
- Bitcoin Power Law Observatory. Research Claims Registry v2.7 (`research/claims.md`).

Data and Reproducibility

Data:

- `bitcoin-powerlaw-site/shared/datasets/btc_historical.json`: 5,951 daily closes, 2010-07-18 to 2026-11-01. The file is appended weekly; every analysis here stops at the cut-off.
- `datasets/gold_monthly.csv`: GLD month-end closes, 2009-12 to 2026-10; see `datasets/GOLD-DATA.md`.

- `datasets/tbill_3m_monthly_fred.csv`: FRED TB3MS.
- `datasets/btc_mvrv_coinmetrics.csv`: Coin Metrics community MVRV, 2010-07-18 to 2026-11-05 (alternative floor only).

Parameters: published power law logA = -16.493, β = 5.688, genesis 2009-01-03; official floor 0.42× trend (C1.3). The headline floor is re-estimated monthly from the data up to each decision date.

Code, all in `analysis/`:

Script	Produces
`f007_backtest.py`	Engine
`f007_research.py`	Strategies, floors, bootstrap, Monte Carlo
`f007_rising_floor.py`	Rising floor
`f007_adaptive_mc.py`, `f007_rolling_floor.py`	Re-estimated floors
`f007_slope_sensitivity.py`	Slope sensitivity
`f007_start_years.py`	Table 4
`p17_v11_experiments.py`	Frozen parameters, windows, Holm, sensitivity, moving-average floor, repricing and decay worlds; writes `p17_v11_results.json`
`p17_v12_experiments.py`	Configuration registry, quantile Kelly, alternative floors, regime worlds; writes `p17_v12_results.json`
`p17_v12_search_stats.py`	Tables 8 and 9 and the PBO: search statistics on distinct configurations
`p17_v12_dsr_grid.py`, `p17_v12_dsr_mix.py`, `p17_v12_gold_window.py`, `p17_v12_penalty.py`	Table 9; gold-window rows of Table 10; the penalty form in Section 2.3
`p17_figures.py`	Figures 1, 3 to 7, 9 to 11, 14 to 16, 18 and 19
`p17_v12_figures.py`	Figures 2, 8, 12, 13 and 17
`p17_prospective.py`	The prospective record (Section 7); spec in `research/p17_frozen_rule.json`
`p17_v13_experiments.py`	v1.3: Holm over all 12 full-Kelly floors, the floor refit inside each simulated path, and the Kelly comparison table; writes `p17_v13_results.json`
`test_f007.py`	Engine tests (5 pass)

Every figure has a CSV of its plotted data in `figures/p17/`.

Seeds: bootstrap and Monte Carlo use seed 20261106; the v1.2 regime draws use seed 777.

Blind replication: the replicator's report and code are archived in `research/p19_blind_replication/`.

Prospective criteria: `research/p19_prospective_criteria.md`.

Appendix: Claim cross-reference

Claim	Statement	Status	Sections
C1.1	Power-law fit and parameters	VERIFIED	2.1
C1.3	Official floor 0.42× trend (definition)	VERIFIED (definition)	2.1, 5
C1.4	Floor visits are rare and short (restated v2.0): 3.7% of days below since 2010, 23 visits, longest 86 days	VERIFIED (descriptive, in sample)	2.1, 5
C6.1	Floor-anchored growth	ANALYTICAL	2.2
C6.2	Full floor-Kelly ends with more Bitcoin than buy and hold (historical)	VERIFIED (backtest), MODERATE; not significant after multiple-testing corrections (Holm, deflated Sharpe)	3.1 to 3.5
C6.3	The edge is timing, not the rebalancing premium	VERIFIED (backtest), HIGH after Holm correction, as a property of the power-law family, not of the headline alone	3.3, 3.5
C6.4	Half floor-Kelly is a risk-control rule	VERIFIED (backtest), MODERATE	3.1, 3.6, 4.3
C6.5	The edge depends on cycles persisting	SUPPORTED (simulation)	3.8
C6.6	A committed rising floor helps; re-estimated floors do not	SUPPORTED	3.9
C6.7	The edge requires the power-law floor shape; four other walk-forward floors fail	VERIFIED (backtest), HIGH	3.7
C6.8	The floor input sits near the 14th percentile of a walk-forward conditional forecast, about one forecast standard deviation below its central growth rate	VERIFIED (computation), model-based, MODERATE	2.3, 3.7
C7.2	Quantile floors are not calibrated out of sample (Paper 20)	VERIFIED (backtest)	1.5, 2.3, 5
C7.3	The trend times its lowest-1% multiple is the most stable floor tested (Paper 20)	VERIFIED (backtest)	1.5, 3.9, 7