Sell Into Euphoria, Buy Back at the Trend

A Forecast-Free Sell Ladder for Bitcoin, Tested in Twelve Futures

Scale Invariant Capital Research | 2026-11-18 | Paper 22 | v1.3

Research on simulated worlds and walk-forward history. Not financial advice; no result here is a guarantee.

v1.3 corrects the statistical worlds. Their 12-month blocks of shocks, taken from Paper 21, under-sampled the first and last 12 months of the sample, both of which had below-average shocks; that set simulated prices about 10% higher against the trend than the fitted model. With circular blocks the recommendation is unchanged (L2-R1), but L2-R3 is now statistically level with it, over 30 years L2-RP now scores clearly higher, a 10% never-sold core no longer lifts L2-R1 above the random-walk stress line (0.78×), and L1.5-R3 becomes admissible (Revision notes).

Claims: C9.1 to C9.6 (Research Claims Registry). Drafted as research report F010. Builds on Paper 19 (Kelly Sizing With Floor-Anchored Growth) and Paper 21 (Floor-Kelly With Money In and Money Out, drafted as F009). Plan: research proposal `research/F010-proposal.md` v1.4 (five rounds of independent review, accepted at 9.1/10) and pre-registration `research/F010-preregistration.md` (commit 415dd52), both written before any simulation or history test ran. What was known when is in Section 12. A later study, F011 (`report/spend_floor_growth.md`, forthcoming as Paper 23), tested the recommended ladder for a retiree who spends from the stack (Section 8, Conclusion).

Abstract

Paper 21 found that floor-Kelly cuts its Bitcoin as the price recovers toward the trend, and a stylized test suggested it sells too early when Bitcoin swings between its floor and its trend. The Observatory's owner proposed a simpler rule:

- hold through the band;
- sell half the Bitcoin held at a first rung (1.5× or 2× the power-law trend), then half of what is left at each higher rung (2×, 3×, 4.5× and every ×1.5);
- buy back later.

The rule needs no forecast of how high the next peak goes, and it does nothing if no euphoria comes. This report tests eight versions (first rung 1.5× or 2×; four ways to buy back) against holding, floor-Kelly, trend-anchored Kelly and a Kelly fed the cycle-weighted histogram of trend multiples. The evidence comes from walk-forward history (2012 to 2026) and twelve simulated worlds of 2,000 paths each:

- seven cycle worlds built from the four observed cycle shapes (calm or euphoric, floor reached or missed, peaks decaying, a random mix, multiples ratcheting up);
- Paper 21's five statistical worlds.

Ladders fill on simulated daily closes resampled from history's intramonth paths; Kelly rules rebalance at month-end. The decision rule was fixed in advance:

- **Admissible:** a rule keeps at least 95% of a holder's Bitcoin on average (geometric mean, months 120 to 168) and 80% at the 10th percentile, in each of nine core worlds.
- **Score:** the family-weighted average log ratio.
- **Recommendation:** only if the rule beats holding in both families and survives month-end fills.

Result. Under the registered rule the recommendation is **L2-R1**: sell half at 2×, 3×, 4.5× and so on, and buy everything back when the price returns to the trend.

- **Score:** it scored highest (0.463) of seven admissible rules that also beat holding in both families. One rule was statistically level with it: L2-R3, which buys back in thirds (0.459). The registered rule then recommends the simpler of the two co-winners, L2-R1. A score of 0 means the same Bitcoin as holding; 0.463 means about 1.6 times as much, on a weighted geometric average across the worlds.
- **No euphoria (peaks of 1.1× to 1.6× trend):** it traded on under 1% of paths and kept a holder's Bitcoin (ratio 1.00).
- **One more euphoria:** it kept 1.52× a holder's Bitcoin if the floor followed and 1.47× if it did not.

Four qualifications matter as much as the headline:

1. **Its lead over trend-anchored Kelly depends on acting on daily prices.** With month-end fills only, or with smaller simulated intramonth moves, trend-anchored Kelly (0.374) scores higher.
2. **Over 30 years the price-anchored ladder L2-RP scores higher (0.821 against 0.763) and is far more robust under stress.** When multiples ratchet up, L2-R1's ratio falls to 0.19× by year 30 while L2-RP's is 1.49×. In a random-walk world it is 0.52× at 10 to 14 years, against L2-RP's 0.83×.
3. **The registered intramonth check flagged every world.** In the largest monthly moves, simulated dips run deeper than history's, which helps buy orders fill, so the ladders' simulated gains may be optimistic.
4. **The statistical worlds swing wider than history** (residual sd 0.346 against 0.283) and supply nearly three quarters of L2-R1's score (0.335 of 0.463).

Buying back only at the floor gained most when the floor came (2.53×) and lost most when it did not (0.55×). A 2× first rung beat 1.5× overall, through the statistical worlds and the admissibility gates; inside the cycle worlds, 1.5× scored slightly higher for three of four buy-backs. Floor-Kelly and the histogram Kelly were not admissible.

History.

- Since 2016, every 2× ladder beat holding (1.11× to 2.09×), and the floor buy-back did best.
- From 2012 every ladder lost on walk-forward lines, mainly because the young walk-forward trend rose from 0.69× to 2.33× the published trend between the 2013 sales and the 2014 buy-back.
- With month-end fills, the floor buy-back's result since 2016 fell from 2.09× to 0.44×, a loss. History since 2012 holds three euphoric cycles: it describes, it does not decide.

Keywords: Bitcoin, power law, sell ladder, take-profit, rebalancing, Kelly criterion, mean reversion, Monte Carlo, walk-forward.

Terms

- **Trend, multiple, floor.** The trend is the published power-law fit of price against time; the multiple is price ÷ trend (0.489 on 2026-10-31); the floor is 0.42× the trend, a level price rarely and briefly goes below.
- **Rung, pool, buy-back.** A rung is a multiple at which a ladder sells half its Bitcoin; the proceeds wait in a cash pool until the buy-back rule spends them.
- **Rule names.** L2 or L1.5: first rung at 2× or 1.5× trend. Buy-backs: R1, all at 1.0× trend; R3, in thirds at 1.0×, 0.75× and 0.5×; RF, at 1.10× the floor; RP, each sale at half its price or a floor-Kelly backstop. TA, FK and HKw: trend-anchored, floor and histogram-centre Kelly; 1 = full, ½ = half.
- **G.** Bitcoin kept (cash counted at the day's price) as a multiple of a holder's Bitcoin, averaged geometrically over years 10 to 14 and over 2,000 paths. **S** is the weighted mean of ln G across worlds (0 = holding). **S_cyc, S_stat:** the score over the cycle or the statistical family alone.
- **Core and stress worlds; families.** Core worlds decide admissibility; stress worlds (UPSHIFT, BROKEN, NO_REVERSION) are reported but do not gate. The cycle family is seven worlds built from history's four cycles; the statistical family is Paper 21's five model-based worlds.
- **Admissible, co-winner.** Admissible: G ≥ 0.95 and 10th percentile ≥ 0.80 in every core world. Co-winner: an admissible rule whose score cannot be told apart from the best by the bootstrap.
- **Daily and month-end fills; intramonth label.** Daily: orders act on simulated daily closes; month-end: only on the last close of the month. The intramonth label marks results that depend on the simulated daily prices, which the registered check found unlike history's.
- **Walk-forward lines.** Trend and floor fitted only to data known at the time.
- **SWR95.** The highest fixed real withdrawal rate that 95% of paths survive for 30 years.
- **The owner.** The Observatory's owner, who proposed the ladder.

Key findings

1. **Recommended under the registered rule: L2-R1.** It sells half the Bitcoin held at 2× trend, half again at 3×, then at 4.5×, 6.75× and so on. It buys all of it back at 1.0× trend, and re-arms once the cash is spent and the price is below trend (Section 3, Table 4).
    - **Score:** 0.463, against 0.459 (L2-R3), 0.437 (L2-RP) and 0.374 (TA 1).
    - **Co-winner with L2-R3.** The bootstrap interval for L2-R3 minus L2-R1 was -0.009 to +0.001, so the two cannot be told apart; the registered rule recommends the simpler one, L2-R1. Every other rule's interval lies below zero. These intervals measure Monte Carlo precision inside fixed worlds, not doubt about which world is true.
    - **Not unique on the gates.** Seven rules passed every core-world gate and beat holding in both families: L2-R1, L2-R3, L2-RP, L1.5-R1, L1.5-R3, TA ½ and TA 1. L2-R1 has the highest point score (within simulation precision of L2-R3) and the simplest buy-back.
2. **No euphoria costs the 2× ladders nothing.** With peaks below 1.6× trend (calm and decay worlds), they traded on under 1% of paths and kept a holder's Bitcoin (G 1.00) (Section 4).
3. **A euphoria pays whether or not the floor follows.** L2-R1 kept 1.52× a holder's Bitcoin when the floor came after the euphoria and 1.47× when it did not. Among the rules that trade, it is the least sensitive to the floor (Table 6, Figure 6).

4. **Waiting for the floor is a bet on the floor.** Buying back only at 1.10× the floor (L2-RF) gave 2.53× when the floor came and 0.55× when it did not, with cash unspent on 98% of paths after 14 years. It is not admissible. Thirds down to 0.5× (L2-R3: 1.93× and 1.29×) and the price-anchored rule with a backstop (L2-RP: 1.76× and 1.42×) sat in between (Section 5).
5. **2× beats 1.5× as the first rung overall, but not inside the cycle worlds** (Section 5):
    - **Overall:** 2× scored higher for all four buy-backs (P2).
    - **Where the advantage comes from:** the statistical worlds (S_stat 0.669 against 0.531 for R1), the admissibility gates (two of four 1.5× ladders fail a core gate, RF and RP, against one at 2×, RF) and the stress worlds.
    - **Inside the cycle worlds:** 1.5× scored slightly higher for R1, R3 and RP.
6. **Trend-anchored Kelly is the strongest rule that is not a ladder, and overtakes L2-R1 when execution changes** (Section 9):
    - **Daily execution:** L2-R1 beats TA 1 when ladders act on daily closes with the simulation's intramonth moves (resampled from history, but with deeper dips than history's in the biggest months).
    - **Month-end fills:** every 2× ladder's score falls by 41% to 48%. L2-R1 drops to 0.250, below TA 1's 0.374. Among the ladders, L2-R3 (0.272) then edges past L2-R1.
    - **Smaller simulated intramonth moves:** with shapes × 0.75 (0.359) or a Gaussian bridge (0.342), TA 1 also scores higher.
7. **Floor-Kelly and the histogram Kelly are not robust.**
    - **Floor-Kelly** gained where the floor came (CALM_FLOOR 1.48×), lost where it was missed (0.68× and 0.63×) and when multiples ratcheted up (0.21×).
    - **The histogram Kelly** scored second in the cycle family (S_cyc 0.285), but its 10th percentile fell below 0.80 in BASE, DAMPED and FADING.
    - Neither is admissible (Section 6).
8. **Stress and long horizons favour L2-RP** (Sections 7 and 9):
    - **Random-walk world (NO_REVERSION):** L2-R1 kept 0.52× at 10 to 14 years, and 10% of paths kept under 3%. L2-RP kept 0.83×.
    - **Multiples ratcheting up (UPSHIFT):** L2-R1 kept 1.38× at 10 to 14 years (L2-RP 1.09×), but its cash stayed stuck and its ratio fell to 0.19× by year 30, against L2-RP's 1.49×.
    - **The 30-year secondary choice:** L2-RP 0.821, ahead of L2-R1's 0.763.
9. **Retirees and savers** (Section 8):
    - **Retirees:** L2-R1 never lowered the 30-year safe withdrawal rate (SWR95, month-end fills) against holding, and raised it by 5.0 points after a euphoria where the floor came.
    - **Savers:** contributions wait in the pool while it holds cash, which hurt badly where the pool stays stuck (UPSHIFT 0.41×, NO_REVERSION 0.32×).
10. **History is description, not a test** (Section 10):
    - since 2016, every 2× ladder beat holding (1.11× to 2.09×), and the floor buy-back did best;
    - from 2012 every ladder lost on walk-forward lines, mostly because the young walk-forward trend rose from 0.69× to 2.33× the published trend between the 2013 sales and the 2014 buy-back (on the published trend, L2-R1 from 2012 was 1.01×);
    - with month-end fills, the floor buy-back's result since 2016 fell from 2.09× to 0.44×, a loss.

11. **A 10% never-sold core is cheap partial insurance against the stress case** (Section 9.3). Keeping 10% of the Bitcoin out of the ladder:
    - costs a little in a euphoria (1.52× → 1.47×);
    - lifts L2-R1 in the random-walk world from 0.52× to 0.78×, just short of the 0.80 stress line, and the worst 10% of paths there still keep under 13%;
    - raises its overall score slightly (0.463 → 0.472).

    This is a secondary analysis and does not change the registered choice.
12. **Against the pre-registered predictions,** five of seven are supported (P2, P3, P4, P6, P7) and two are not (P1, P5). Both failures were informed predictions:
    - floor-Kelly did far better in calm floor-reaching worlds than a stylized test suggested;
    - L2-R3 and L2-R1, level with each other, not trend-anchored Kelly, led the statistical family (Section 11).

---

1. Introduction

1.1 The question

The owner's question: how far above trend will Bitcoin go in its next euphoria, and how do you build a system that is ready for a 3× trend without needing one? It must work whether or not a euphoria comes, and whether or not the price returns to the floor.

Paper 21 tested floor-Kelly with money in and out. Floor-Kelly sizes Bitcoin from the growth to the floor four years ahead. In a stylized smooth band from 0.42× to 1.25× trend (an informal calculation made before the proposal, not in the repository), it ended with 0.69× a holder's Bitcoin after 12 years: it sold on the way up and was back in long before the next low. The owner proposed a ladder instead:

- hold through the band;
- sell half at 1.5×, half again at 2×, then at 3×, and so on;
- buy back near the floor.

Selling rich and buying back near a mean is the logic of mean-reversion trading rules (Bertram 2010; Leung and Li 2015); here the mean is the power-law trend.

1.2 How euphorias have shrunk

On daily closes against the published trend, the cycle peaks were 15.8×, 11.0×, 6.2×, 3.3× and 1.25×. Each peak was 0.38 to 0.70 times the one before (Figure 1). The lowest close in each halving cycle since 2012 was 0.36× to 0.44× the trend. In the 2016 to 2020 cycle the 2018 to 2019 bear market bottomed at 0.54×, and the price came back near the floor (0.44×) only in the COVID crash of March 2020. The 2022 and 2026 lows (0.41×) closed slightly below the published floor (0.42×), briefly: since 2016 the price has spent 0.3% of days below the floor, the longest visit 7 days. If the shrinking continues, the next peak may not reach 2×. If the shrinking stops, it may. The ladder is designed for that uncertainty.



Figure 1. Each euphoria has been smaller than the last; since 2015 the lows have stayed near the floor.

1.3 The trend-growth clock

A rule that sells Bitcoin must buy it back cheaper. Bitcoin sold at multiple M is bought back cheaper only if the multiple later falls below M divided by the trend's growth since the sale. The trend grows $((\text{age} + 1)/\text{age})^{5.688}$ a year, where age is years since 3 January 2009: 3.1× in mid-2013, 1.36× in October 2026 and 1.16× in 2046.

Table 1. The trend-growth clock from October 2026 (exact published trend).

Sold at	Break-even buy-back multiple after 1 year	After 2 years	After 3 years
1.5×	1.10×	0.82×	0.62×
2×	1.47×	1.09×	0.83×
3×	2.20×	1.64×	1.24×

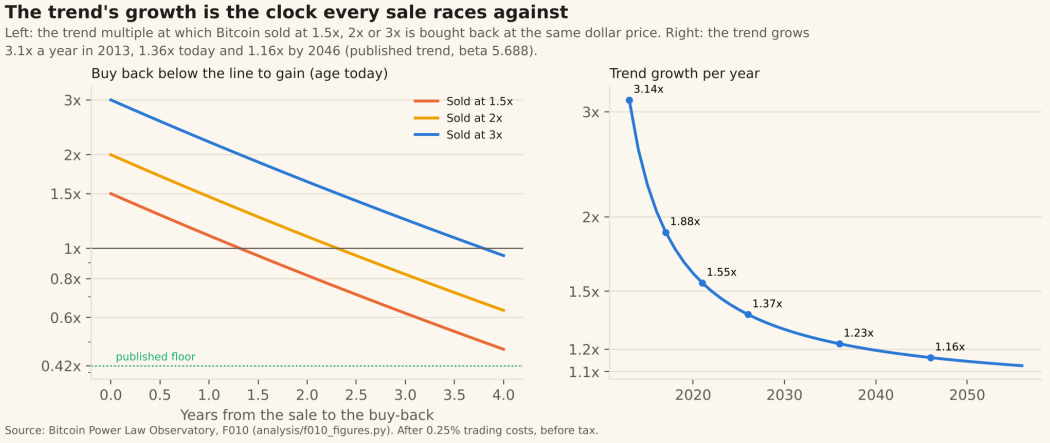


Figure 2. The trend's growth is the clock every sale races against. Left: the multiple at which Bitcoin sold at 1.5×, 2× or 3× is bought back at the same dollar price. Right: trend growth per year.

1.4 Questions

1. What does a sell ladder cost when no euphoria comes?
2. What does it gain when one does?
3. Which buy-back works whether or not the price returns to the floor?

4. Should the first rung be 1.5× or 2×?
5. Does a Kelly rule beat the ladder?
6. Is any rule acceptable in every core world and better than holding in both families of worlds?

2. Rules, Worlds and Decision Rule

This section summarises the design. The full specification is in the proposal (Sections 6.1 to 6.6).

2.1 The rules

Ladders. Each ladder starts 100% in Bitcoin. At the first rung (1.5× or 2× trend), then at 2×, 3×, 4.5×, 6.75× and so on, it sells half the Bitcoin held. The proceeds go into a cash pool earning the T-bill rate.

Table 2. The four buy-back rules.

Code	Buy-back
R1	The whole pool at 1.0× trend
R3	Three orders at 1.0×, 0.75× and 0.5× trend; each buys the pool divided by the orders still open
RF	The whole pool at 1.10× the floor (0.462× trend)
RP	Each sale separately, at half its sale price or at the floor-Kelly backstop, whichever comes first. The backstop is the price at which full floor-Kelly would be 100% Bitcoin, about 0.74× trend at today's volatility

How the ladders trade:

- **Fills.** Orders are limit orders checked on daily closes, and a crossing fills at the level.
- **Costs.** Each fill costs 0.25%.
- **Re-arming.** The rungs re-arm once the pool is empty and a close is below 1.0× trend.
- **Names.** L2-R1 means first rung 2×, buy-back R1.

Kelly rules. These hold Bitcoin and cash and rebalance at month-end closes, at full (c = 1) and half Kelly (Kelly 1956; on fractional Kelly, Rotando and Thorp 1992; MacLean, Thorp and Ziemba 2010), with $w = c(g + \sigma^2/2 - r)/\sigma^2$ capped to 0 to 1. The growth input g is:

- **FK:** growth to the floor four years ahead (floor-Kelly).
- **TA:** growth to the trend four years ahead (trend-anchored Kelly).
- **HKw:** the trend's growth plus reversion toward the cycle-weighted centre of past trend multiples (the owner's histogram idea).

Lines. In simulation every rule uses the published trend ($\log_{10} P = -16.493 + 5.688 \log_{10}$ days) and floor (0.42× trend); in history every rule uses walk-forward lines. A matched variant puts every rule on lines refit inside each path (Section 9).

2.2 The worlds

All worlds start at the October 2026 close (0.489× trend) and run 2,000 paths of 30 years, with cash at 4.05%.

Cycle worlds. Each path chains cycles. Each cycle takes the shape of one of the four observed cycles (2013, 2017, 2021, 2024), stretched in time and rescaled in log space to drawn peak and trough levels, plus small noise.

Table 3. Cycle worlds. Levels are monthly closes against the published trend. All seven share each path's shapes and random numbers (common random numbers; Glasserman 2004).

World	Peaks	Troughs	Role
CALM_FLOOR	1.1× to 1.6×	0.38× to 0.48×	Core
CALM_MISS	1.1× to 1.6×	0.55× to 0.75×	Core
EUPH_FLOOR	First peak 2.5× to 4×, then 1.1× to 1.6×	0.38× to 0.48×	Core
EUPH_MISS	First peak 2.5× to 4×, then 1.1× to 1.6×	0.55× to 0.75×	Core
DECAY	The excess over 1× keeps 40% to 70% of itself each cycle (shrinks by 30% to 60%), from 1.16×	0.38× to 0.48×	Core
WILD	1.0× to 4× each cycle	0.35× to 0.70× each cycle	Core
UPSHIFT	Troughs × 2.5 to 4	Troughs rise 25% to 50% per cycle	Stress

Statistical worlds (Paper 21's paths):

- **BASE, DAMPED and FADING** (core): AR(1) reversion toward 0.92× the published trend, with block-bootstrapped shocks. DAMPED halves the shocks; FADING shrinks them over time.
- **BROKEN and NO_REVERSION** (stress): BROKEN halves the trend's growth; NO_REVERSION is a random walk.
- **BASE is wider than history:** residual sd 0.346 against 0.283, and 16% of months below the floor against 2%.

Daily closes. Each simulated month's 30 daily closes follow the straight line between its monthly closes, plus a shape resampled from history's intramonth paths:

- **Cycle worlds:** from the 12 historical months nearest in standardised move, rescaled per size band.
- **Statistical worlds:** the shape of the month whose shock was drawn.

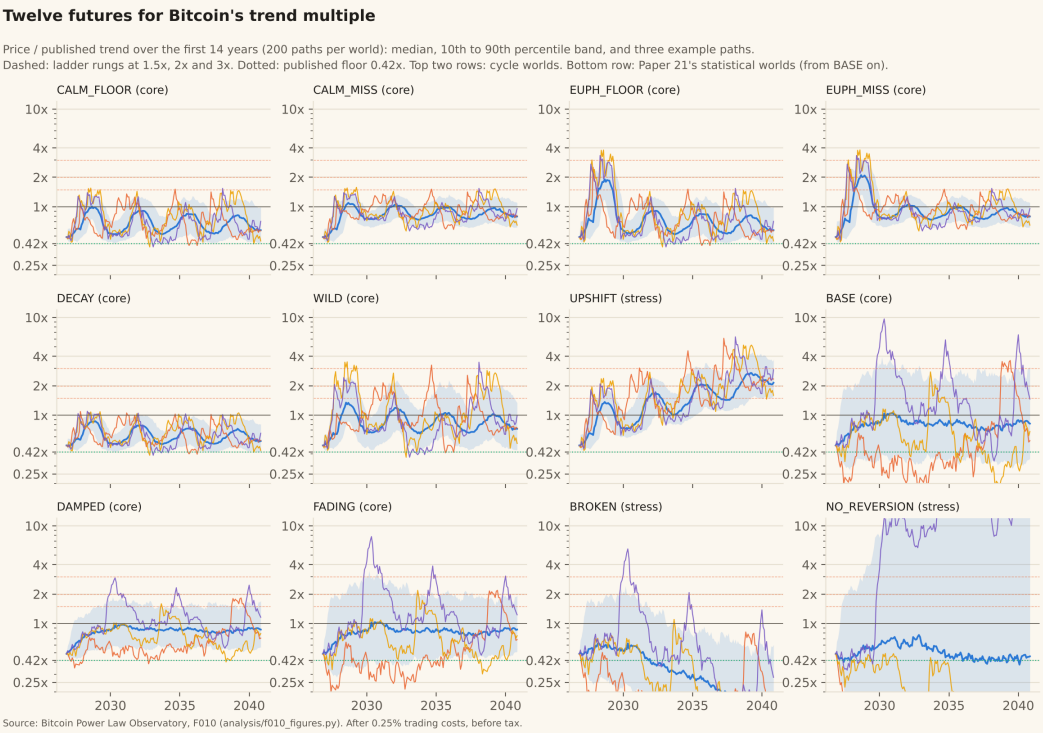


Figure 3. Twelve futures for Bitcoin's trend multiple: median, 10th to 90th percentile band and three example paths over the first 14 years, from 200 simulated paths per world (the results use 2,000).

2.3 Metric and decision rule

- **Stack ratio:** a rule's Bitcoin, with cash counted at the day's price, divided by a holder's Bitcoin.
- **Per path:** the mean of ln(stack ratio) over month-ends 120 to 168. This is about one cycle, ten to fourteen years out, so no single cycle phase dominates.
- **Per world:** G = exp(mean over paths), the geometric-mean ratio.

The rule was fixed before any run:

1. **Admissible:** G ≥ 0.95 and 10th percentile ≥ 0.80 in each of the nine core worlds.
2. **Scores:** S_cyc is the mean ln G over the 7 cycle worlds; S_stat is the mean over the 5 statistical worlds; S = ½ S_cyc + ½ S_stat.
3. **Winner and co-winners:** the winner is the admissible rule with the highest S. Co-winners are admissible rules whose 90% bootstrap interval for (S - S of the winner) includes 0, from 1,000 joint path resamples (Efron and Tibshirani 1993).
4. **Recommendation:** the simplest co-winner with S_cyc ≥ 0 and S_stat ≥ 0. A ladder must also stay admissible, with S ≥ 0, when filled only at month-end closes. Stress worlds give a label but do not gate.

3. The Answer

Table 4. G, the geometric-mean ratio to a holder's Bitcoin over months 120 to 168, by rule and world. * stress world. † not admissible (fails a core-world gate). 10th percentiles are in Appendix Table A1. All twelve worlds carry the registered intramonth label (Section 12): in the largest-move months, simulated dips run deeper than history's.

Rule	Calm floor	Calm miss	Euph. floor	Euph. miss	Decay	Wild	Upshift*	Base	Damped	Fading	Broken*	No rev.*	S_cyc	S_stat	S
HOLD	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.000	0.000	0.000
L2-R1	**1.00**	**1.00**	**1.52**	**1.47**	**1.00**	**1.96**	**1.38**	**3.94**	**1.52**	**3.01**	**3.01**	**0.52**	**0.256**	**0.669**	**0.463**
L2-RF†	1.00	1.00	2.53	0.55	1.00	2.03	0.57	3.02	0.98	1.86	4.19	0.50	0.069	0.488	0.279
L2-R3	1.00	1.00	1.93	1.29	1.00	2.25	1.01	4.07	1.27	2.63	3.72	0.56	0.248	0.670	0.459
L2-RP	1.00	1.00	1.76	1.42	1.00	2.24	1.09	3.07	1.34	2.39	2.66	0.83	0.259	0.615	0.437
L1.5-R1	1.24	1.15	1.58	1.42	1.00	1.88	1.05	2.92	1.56	2.46	2.92	0.44	0.263	0.531	0.397
L1.5-RF†	1.65	0.82	3.25	0.39	1.00	1.83	0.33	2.24	0.77	1.46	4.71	0.42	0.007	0.322	0.164
L1.5-R3	1.39	1.04	2.21	1.09	1.00	2.28	0.77	3.23	1.18	2.23	3.96	0.50	0.258	0.566	0.412
L1.5-RP†	1.48	1.18	2.27	1.18	1.00	2.15	0.81	2.30	1.17	1.84	3.25	0.72	0.301	0.491	0.396
TA ½	1.43	1.12	1.63	1.15	1.09	1.46	0.61	1.71	1.27	1.43	3.96	0.41	0.154	0.325	0.239
TA 1	1.06	1.01	1.59	1.37	1.00	1.77	1.18	3.01	1.36	2.20	2.79	0.54	0.226	0.522	0.374
FK ½†	0.64	0.29	0.53	0.25	0.84	0.36	0.09	0.57	0.31	0.43	2.76	0.27	-1.048	-0.569	-0.809
FK 1†	1.48	0.68	1.44	0.63	1.42	1.01	0.21	1.73	0.67	1.19	3.65	0.37	-0.183	0.123	-0.030
HKw ½†	1.48	1.18	1.46	1.04	1.38	1.18	0.47	1.33	0.99	1.09	3.79	0.27	0.099	0.077	0.088
HKw 1†	1.63	1.29	1.77	1.30	1.31	1.49	0.78	2.15	1.23	1.68	3.62	0.33	0.285	0.334	0.310

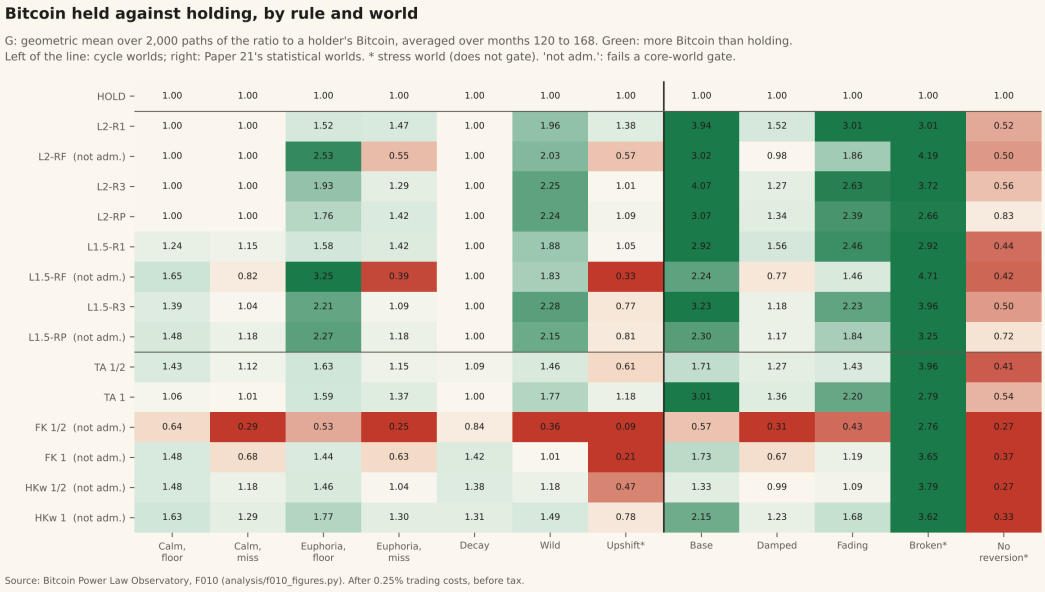


Figure 4. Bitcoin held against holding, by rule and world.

Admissibility. Eight rules passed: HOLD, L2-R1, L2-R3, L2-RP, L1.5-R1, L1.5-R3, TA ½ and TA 1. The rest failed a core-world gate:

Table 5. The core-world gates failed by the rules that are not admissible.

Rule	Failing core worlds
L2-RF	EUPH_MISS, DAMPED, FADING
L1.5-RF	CALM_MISS, EUPH_MISS, WILD, BASE, DAMPED, FADING
L1.5-RP	DAMPED (10th percentile)
FK ½	All nine
FK 1	CALM_MISS, EUPH_MISS, WILD, BASE, DAMPED, FADING
HKw ½	EUPH_MISS, BASE, DAMPED, FADING
HKw 1	BASE, DAMPED, FADING (10th percentile)

Every admissible rule except HOLD also beats holding in both families.

Winner.

- **L2-R1 has the highest score (0.463), and L2-R3 (0.459) is a co-winner.** L2-R3's 90% interval against L2-R1 is -0.009 to +0.001 and includes zero. Every other admissible rule's interval lies below zero; the next is L2-RP (-0.031 to -0.020).
- **The intervals are tight** because the bootstrap resamples paths inside fixed worlds. L2-R1's 90% half-widths are 1.4% or less of G in nine worlds, 2.6% to 3.6% in FADING, BASE and BROKEN, and 6.8% in NO_REVERSION. They measure simulation precision, not uncertainty about which world we live in.
- **The recommendation goes to the simpler co-winner.** The registered rule takes the simplest co-winner that passes the checks. L2-R1, one buy order at the trend, is simpler than L2-R3, three orders at 1.0×, 0.75× and 0.5×, and it passes: it beats holding in both families, and with month-end fills it stays admissible and scores 0.250 (≥ 0).
- **So L2-R1 is the recommended rule, with the stress label NO_REVERSION (G 0.52). L2-R3 is statistically level with it and carries the same stress label (0.56).**

Robustness of the choice. Between L2-R1 and L2-R3 the order is fragile, as their overlapping interval says; against every other rule it is not.

- **Equal world weights:** L2-R1 still wins.
- **Leaving out one world:** L2-R1 wins in 8 of 12 cases; without EUPH_MISS, UPSHIFT, DAMPED or FADING, L2-R3 wins.
- **Small weight changes swap the two co-winners.** The registered weights are 7.1% (cycle worlds) and 10% (statistical worlds). L2-R3 becomes the winner if EUPH_FLOOR's weight rises to 8.7% of the total, WILD's to 9.8%, BROKEN's to 11.7%, NO_REVERSION's to 14.7% or BASE's to 19.9%, or if EUPH_MISS's falls to 4.2%, UPSHIFT's to 5.9%, FADING's to 7.2% or DAMPED's to 7.9%.
- **Other rules need large weight changes:** CALM_FLOOR at 19.7% (toward L1.5-R3), CALM_MISS at 37.1% (toward L1.5-R1), DECAY at 73.7% (toward TA ½) and DAMPED at 77.2% (toward L1.5-R1). Table A3 in `analysis/f010_tables.md` lists them all.

Within each family:

- **Statistical family:** L2-R3 leads (S_stat 0.670) and L2-R1 (0.669) is its co-winner.
- **Cycle family:** L1.5-RP leads (S_cyc 0.301), ahead of HKw 1 (0.285) and L1.5-R1 (0.263). L2-R1 is sixth (0.256), and its interval against L1.5-RP lies below zero.

So L2-R1 is level with L2-R3 at the top of the statistical family and sixth in the cycle family. It wins overall because it does well in both.

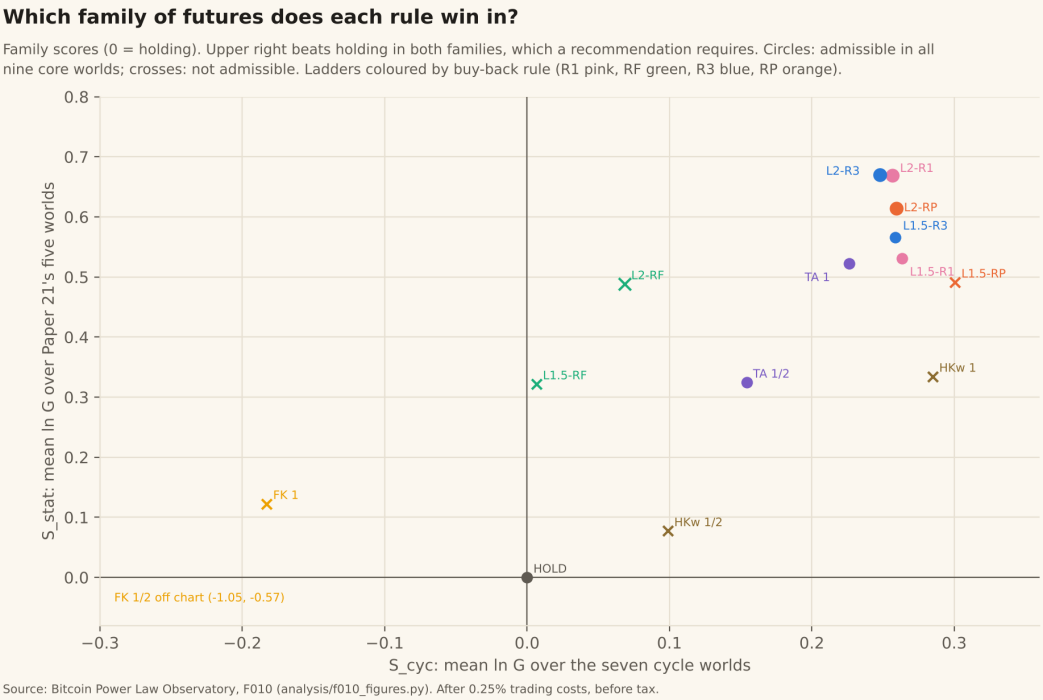


Figure 5. Family scores: upper right beats holding in both families.

4. The Owner's Four Cases

Table 6. The 2 × 2 of cycle worlds. Intramonth label applies (Section 12).

Rule	Calm, floor	Calm, miss	Euphoria, floor	Euphoria, miss	Floor minus miss (ln G)	Euphoria minus calm (ln G)
L2-R1	1.00	1.00	1.52	1.47	0.016	0.400
L2-RF	1.00	1.00	2.53	0.55	0.762	0.168
L2-R3	1.00	1.00	1.93	1.29	0.203	0.455
L2-RP	1.00	1.00	1.76	1.42	0.106	0.458
L1.5-R1	1.24	1.15	1.58	1.42	0.092	0.225
L1.5-RP	1.48	1.18	2.27	1.18	0.440	0.213
TA 1	1.06	1.01	1.59	1.37	0.098	0.355
FK 1	1.48	0.68	1.44	0.63	0.805	-0.046
HKw 1	1.63	1.29	1.77	1.30	0.273	0.045

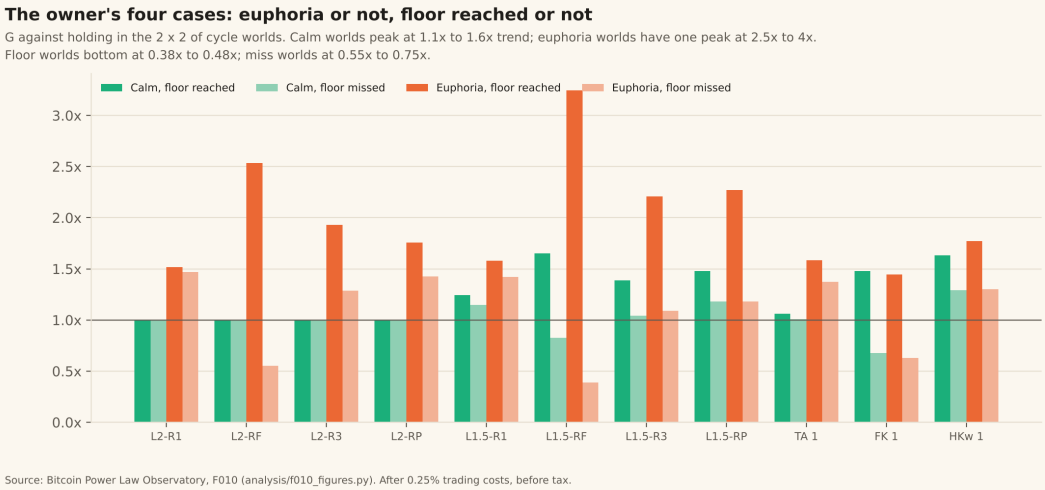


Figure 6. The owner's four cases: euphoria or not, floor reached or not.

No euphoria (Q1). With peaks between 1.1× and 1.6× trend, the 2× ladders traded on under 1% of paths and kept a holder's Bitcoin (G 1.00). That is the design: insurance that costs nothing when it is not needed.

The 1.5× ladders did trade in calm worlds and, against expectation, gained:

- 1.24× to 1.65× in CALM_FLOOR;
- 1.04× to 1.18× in CALM_MISS, except RF (0.82×).

Floor-Kelly (1.48×) and the histogram Kelly (1.63×) also gained where the floor came.

Interpretation, not tested separately: the historical cycle shapes swing up and down inside a calm band (the 2019 rally, the 2024 double top). A rule that sells the swings and buys the dips gains from that. The smooth stylized band that motivated P1 had no such swings, which is the likely reason P1's floor-Kelly threshold failed.

One more euphoria (Q2). Every ladder gained in EUPH_FLOOR (1.52× to 3.25×). L2-R1 gained about the same whether the floor came (1.52×) or not (1.47×). It sold at 2× and 3× on the way up and bought everything back at the trend on the way down, before the floor question arose.

Floor reached or missed (Q3).

- **R1 hardly depends on the floor** (floor minus miss 0.016, the smallest of any rule that trades).
- **RF is the floor bet.** It gained most when the floor came (2.53×; 3.25× at a 1.5× first rung) and lost most when it did not (0.55×; 0.39×).
- **R3 and RP sit in between.**

Figure 7 shows why RF fails: in EUPH_MISS, 98% of L2-RF paths still held unspent cash after 14 years.

5. Which Buy-Back, and Which First Rung?

First rung (Q4). Overall, the 2× version scored higher for all four buy-backs:

Table 7. Overall score S by first rung, for each buy-back.

Buy-back	S at 2×	S at 1.5×
R1	0.463	0.397
RF	0.279	0.164
R3	0.459	0.412
RP	0.437	0.396

The advantage does not come from the calm cycle worlds, where the 1.5× ladders gained (Section 4). Inside the cycle family, 1.5× scored slightly higher for R1 (0.263 against 0.256), R3 (0.258 against 0.248) and RP (0.301 against 0.259). The 2× advantage comes from three places:

- **The statistical family** (S_stat for R1: 0.669 against 0.531). *Interpretation:* its wider swings make a sale at 1.5× more often a sale in the middle of a rally.
- **Admissibility:** two of four 1.5× ladders failed a core-world gate (RF and RP), against one of four at 2× (RF).
- **The stress worlds:** UPSHIFT 1.05× against 1.38× for R1; NO_REVERSION 0.44× against 0.52×.

The trend-growth clock (Table 1) explains the historical result (Section 10). In the simulation it is at most part of the story.

Buy-back.

- **Depth pays only if the price gets there.** RF and R3 sell high and buy deep. That wins big when the floor comes and strands cash when it does not.
- **R1 gives up the last leg of every crash,** because it buys back at the trend, often well above the eventual low. In return it rarely waits long in the cycle worlds: its longest run more than half in cash, over the full 30 years, was 16 to 23 months at the 90th percentile in the euphoria and WILD worlds, against 129 months for RF in EUPH_MISS. In the statistical worlds it waits longer: in BASE, a median of 33 months and a 90th percentile of 59, with cash still in the pool at month 168 on 29% of paths.
- **RP is the most robust under stress.** It is the only ladder above 0.80 in NO_REVERSION (0.83). Over 30 years it leads the secondary choice (Section 9). At 10 to 14 years it scored 0.437, against R1's 0.463.

Table 8. Intramonth label applies (Section 12). In each cell: the share of paths with unspent cash in the pool at month 168, then the 90th percentile of the longest run more than half in cash, in months over the full 30 years.

Ladder	Euphoria, floor	Euphoria, miss	Wild	Upshift	Base
L2-R1	0% / 16	0% / 17	9% / 23	100% / 256	29% / 59
L2-RF	0% / 24	98% / 129	47% / 111	100% / 215	54% / 145
L2-R3	0% / 17	91% / 22	38% / 28	100% / 194	50% / 81
L2-RP	0% / 16	1% / 23	14% / 24	100% / 57	49% / 27

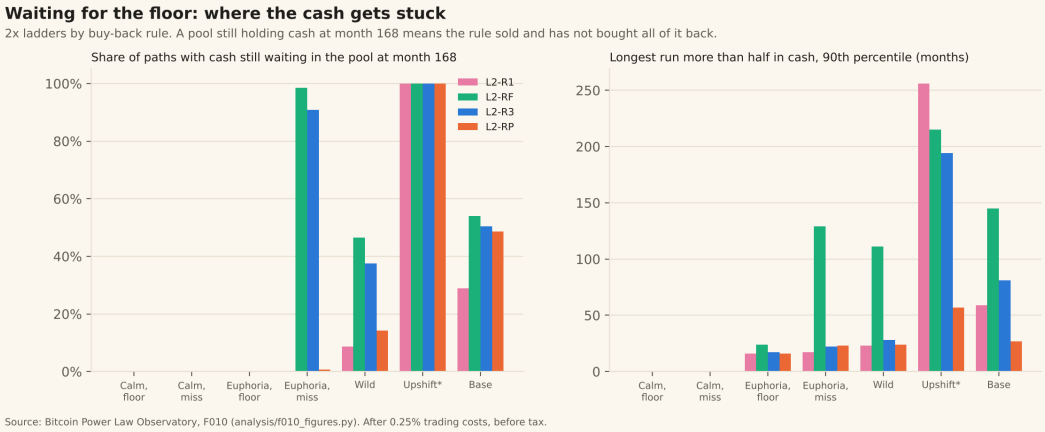


Figure 7. Waiting for the floor: where the cash gets stuck.

Timing, not exposure. L2-R1 beat a constant mix holding the same average Bitcoin:

- **Against a mix rebalanced monthly:** 1.62× in EUPH_FLOOR, 2.28× in WILD and 4.18× in BASE.
- **Against a mix never rebalanced:** 1.60×, 2.18× and 5.11×.

So its gain comes from when it sells and buys, not from holding less Bitcoin on average (Table A6 in `analysis/f010_tables.md`).

6. The Kelly Rules

Floor-Kelly (FK) is not admissible.

- **Where it gains:** where the floor came (CALM_FLOOR 1.48×, DECAY 1.42×).
- **Where it loses:** where the floor was missed (CALM_MISS 0.68×, EUPH_MISS 0.63×) and in UPSHIFT (0.21×). There it held little Bitcoin, because the price stayed far above the floor it uses as its growth anchor.
- **Half floor-Kelly** lost in every core world.
- **A 10% minimum weight**, the owner's earlier idea, changed almost nothing (Table A13 in `analysis/f010_tables.md`).

Histogram-centre Kelly (HKw) comes closest to the owner's histogram idea.

- **Cycle family:** it scored second best (S_cyc 0.285).
- **Statistical worlds:** it failed admissibility on the 10th percentile (BASE 0.77, DAMPED 0.75, FADING 0.72). That was a surprise: the proposal expected a built-in edge for HKw there, because it uses the published trend and the AR(1)'s own reversion speed.
- **Completed cycles only (sensitivity):** it did better in the cycle worlds (EUPH_FLOOR 2.15× for HKw 1). Its 10th percentile rose to 0.87 in BASE but stayed below 0.80 in DAMPED (0.74) and FADING (0.79), so it would still not be admissible.

Trend-anchored Kelly (TA 1) is admissible and scored 0.374.

- **Calm worlds:** it hardly trades (1.01× to 1.06×), because it holds 100% Bitcoin below about 1.75× trend, much like a ladder whose first rung sits at 1.75×.
- **Against L2-R1:** it was slightly ahead in the calm worlds (1.06 and 1.01 against 1.00), in EUPH_FLOOR (1.59 against 1.52) and in NO_REVERSION (0.54 against 0.52). It trailed in EUPH_MISS, WILD, UPSHIFT and the other four statistical worlds, by up to 0.93 in BASE.

- **The comparison depends on execution.** TA rebalances once a month; the ladders fill on daily closes. With month-end fills, TA 1 scores higher than every ladder (Section 9).
- **P5 failed:** L2-R3 and L2-R1, not TA, led the statistical family.

7. Stress Worlds

Table 9. G in the stress worlds over two windows. Below 0.80 means the rule kept less than 80% of a holder's Bitcoin. Intramonth label applies (Section 12).

Rule	NO_REVERSION, months 120 to 168	NO_REVERSION, months 312 to 360	BROKEN, months 120 to 168	UPSHIFT, months 120 to 168	UPSHIFT, months 312 to 360
L2-R1	0.52	0.37	3.01	1.38	0.19
L2-R3	0.56	0.46	3.72	1.01	0.20
L2-RP	0.83	0.99	2.66	1.09	1.49
TA 1	0.54	0.27	2.79	1.18	0.08
FK 1	0.37	0.16	3.65	0.21	0.01
HKw 1	0.33	0.10	3.62	0.78	0.25

NO_REVERSION is a random walk around the trend with no pull back. Every timing rule lost there: selling into a rise that never reverses strands the cash.

- **L2-R1:** its median path averaged level with holding (1.00), but 30% of paths fell below 0.80 and 10% kept under 3% of a holder's Bitcoin (Figure 12).
- **L2-RP lost least.** Its per-sale buy-back at half the sale price, or the backstop, fills sooner: its median longest run more than half in cash was 27 months, against L2-R1's 62.

BROKEN halves the trend's growth. Holding Bitcoin does badly there, and every rule that ever holds cash wins.

UPSHIFT ratchets the troughs up by 25% to 50% a cycle.

- **L2-R1:** its buy level at the trend soon stops being reached. Its median longest run more than half in cash was 211 months, and its pool still held unspent cash at month 168 on 99.9% of paths. At 10 to 14 years it still kept 1.38×, because it sold at the rising peaks before getting stuck. By year 30, its stuck cash had fallen to 0.19× a holder's Bitcoin.
- **L2-RP:** buys back at half its sale prices. It kept 1.09× at 10 to 14 years and 1.49× at year 30. Its median longest run more than half in cash was 32 months, against L2-R1's 211.

Read. L2-R1 is insurance against euphoria, not a hedge against the power law failing. If Bitcoin's multiples wander, or ratchet up for decades, a rule that waits for the trend can wait forever. L2-RP is the safer choice for an investor who weights those worlds heavily.

8. Money Out and Money In

Retiree ($1M, fixed real withdrawals, 30 years; ladders fill at month-end here and take withdrawals from their pool first). L2-R1's SWR95 was never below holding's, and it was higher where euphorias come:

Table 10. Retiree SWR95 (month-end fills, so the intramonth label does not apply). * At the top of the 1% to 40% grid, so censored; the UPSHIFT values (40.0% for most rules) are censored too.

World	Holding	L2-R1
EUPH_FLOOR	30.5%	35.5%
EUPH_MISS	36.5%	40.0%*
WILD	30.5%	32.5%
BASE	19.5%	21.0%

Other retiree results:

- **Half floor-Kelly** was the worst rule for retirees in every world except BROKEN, where it was among the best, and NO_REVERSION, where no rule sustained any rate at 95% survival.
- **The rates are high** because the simulated trend keeps compounding. They compare rules; they are not advice.

Saver (20 years of monthly contributions; month-end fills, as for the retiree). Contributions go into the pool while it holds cash; otherwise they buy Bitcoin. L2-R1's final value against Bitcoin-only DCA:

- **Gains:** 1.17× in EUPH_FLOOR, 2.39× in WILD and 4.16× in BASE.
- **Losses:** 0.41× in UPSHIFT and 0.32× in NO_REVERSION, where the pool stays stuck and every contribution waits in cash with it.

A saver using a ladder would do better to send new money into Bitcoin than into a pool that has waited for years. That variant was not tested; it is open question OQ18 in the Observatory's registry, together with tracking L2-R1 prospectively from November 2026 and a tax-aware version.

Later work: spending from the stack. A later study (F011, `report/spend_floor_growth.md`, forthcoming as Paper 23) tested L2-R1 for a retiree who spends a share of the floor's growth, in the same twelve worlds. Its pre-registered rule again picked L2-R1 among the eight rules it tested, with a small gain in safe spending (+4.8 coverage points on average). It also found that L2-R1's lead is largest with daily execution and comes mostly from euphorias and the wide statistical worlds, that its score falls with month-end fills (0.773 to 0.500), and that if multiples ratchet up (UPSHIFT) its cash can sit unspent for decades. L2-R3, L2-RP and trend-anchored Kelly were not in that field; a registered addendum (F011-R, `research/F011R-preregistration.md`) tests them under spending, and its results are not yet available.

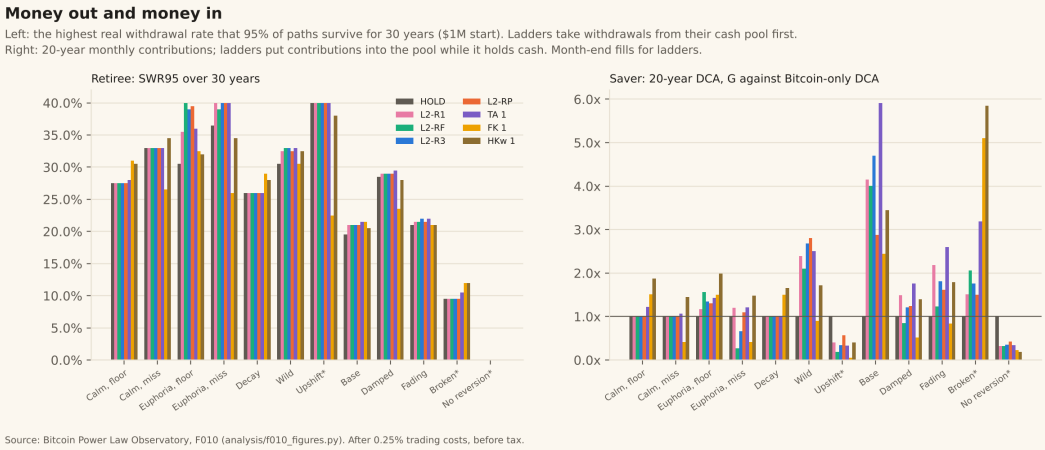


Figure 8. Money out and money in.

9. Sensitivities and Secondary Choices

Every registered sensitivity and secondary choice was run after the primary.

9.1 Execution and intramonth moves

Table 11. Overall score S by sensitivity; the intramonth label applies to every row except month-end fills. † not admissible in that variant. TA 1 rebalances at month-end, so its score does not change in the fill and intramonth variants. L1.5-R3, admissible in every variant except matched refit lines, scores higher than L1.5-R1 in the primary (0.412), with month-end fills (0.203), shapes × 0.75 (0.312), the Gaussian bridge (0.304) and cash at 2.65% (0.381); it is below every admissible 2× ladder in every row.

Variant	L2-R1	L2-R3	L2-RP	L2-RF†	L1.5-R1	TA 1	Highest-scoring admissible rule
Primary	0.463	0.459	0.437	0.279	0.397	0.374	L2-R1
Month-end fills only	0.250	0.272	0.226	0.158	0.141	0.374	**TA 1**
Intramonth shapes × 1.25	0.601	0.565	0.555	0.362	0.573	0.374	L2-R1
Intramonth shapes × 0.75	0.359	0.369	0.342	0.215	0.267	0.374	**TA 1**
Gaussian bridge instead of history's shapes	0.342	0.352	0.320	0.209	0.275	0.374	**TA 1**
Cash at 2.65% (about 0% real)	0.443	0.435	0.422	0.243	0.371	0.350	L2-R1
Matched refit lines (all 15 rules)	0.491	0.442	0.452	0.264	0.381	0.396	L2-R1

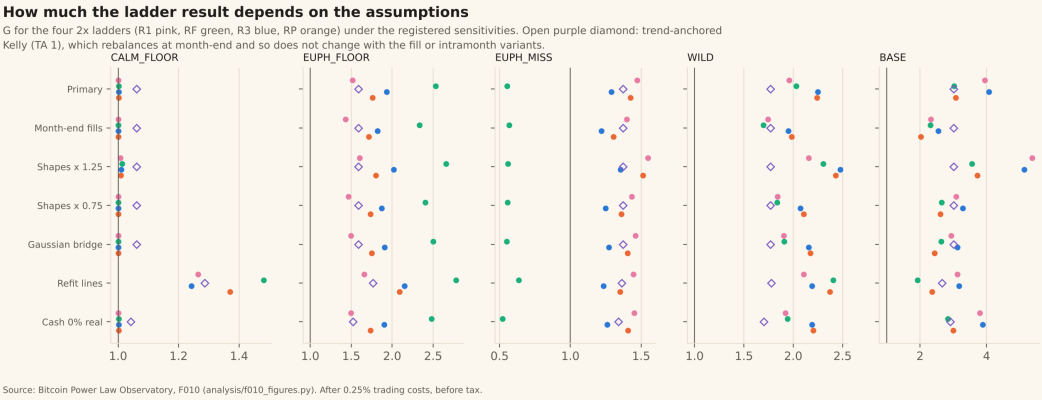


Figure 9. How much the ladder result depends on the assumptions. Open diamond: trend-anchored Kelly (TA 1).

What the variants show:

- **Daily execution drives L2-R1's lead over trend-anchored Kelly.**
  - With month-end fills only, every 2× ladder's score falls by 41% to 48% (1.5× ladders by more). L2-R1 drops to 0.250, below TA 1's 0.374.
  - With smaller intramonth moves (shapes × 0.75: 0.359) or a Gaussian bridge (0.342), TA 1 also scores higher, and among the ladders L2-R3 edges past L2-R1.
  - **In practice:** an investor who can keep standing limit orders at the rungs and at the trend gets the ladder's edge, if intramonth moves are as large as in the main simulation. With 25% smaller moves or a Gaussian bridge, TA 1 scores higher even with daily orders. An investor who acts once a month does better with trend-anchored Kelly. Rungs and the buy-back are

multiples of a trend that rises about 0.08% a day, so the order prices must be updated often; an order re-priced only monthly drifts by up to 2.6%.

- **The size of the gain depends on intramonth spikes.**
  - L2-R1's score ranges from 0.342 to 0.601 across the intramonth variants.
  - In the euphoria worlds the range is narrow: EUPH_FLOOR 1.47× to 1.60×.
  - In the statistical worlds it is wide: BASE 2.93× to 5.37×.
- **Refit lines and cash do not change the choice.** With all 15 rules on lines refit inside each path, L2-R1 is again the winner and the only co-winner, among admissible rules HOLD, L2-R1, L2-R3, L2-RP, L1.5-R1 and TA 1. With cash at 0% real, L2-R1 still leads (0.443 against TA 1's 0.350).
- **RP depth trades robustness for gain.**
  - At 30%: EUPH_FLOOR and EUPH_MISS both 1.33×, NO_REVERSION 0.89×.
  - At 70%: EUPH_FLOOR 1.96×, EUPH_MISS 1.22×, UPSHIFT 0.66×. L2-RP stays admissible at both depths, but at 70% only narrowly (10th percentile 0.85 in DAMPED).
- **Cycle-world noise changes little in the euphoria worlds:** L2-R1 kept 1.51× and 1.47× at noise 0, and 1.55× and 1.51× at 0.04 (1,000 paths each, as registered); UPSHIFT moves more (1.34× to 1.57×).

9.2 Other horizons (registered secondary choices)

Table 12. Secondary choices on other horizons, among the rules admissible in the primary. The 10th percentiles for these windows were not stored, so admissibility is taken from the primary. Intramonth label applies (Section 12).

Window	First	Second	Third	Fourth	Fifth
30 years (months 312 to 360), S over G	L2-RP 0.821	L2-R1 0.763	L1.5-R3 0.749	L1.5-R1 0.745	L2-R3 0.740
Month 144, S over the median point ratio	L2-R1 0.487	L2-R3 0.466	L1.5-R1 0.438	L1.5-R3 0.435	L2-RP 0.427

Over 30 years, L2-RP leads L2-R1 by 0.058, about ten times the primary horizon's interval half-widths of about 0.005 (no bootstrap was run for this window). L2-RP's lead comes from the stress worlds (UPSHIFT 1.49× against L2-R1's 0.19×; NO_REVERSION 0.99× against 0.37×). At the month-144 point, L2-R1 leads as in the primary.

9.3 Never-sold core, certainty equivalent and diagnostics

A never-sold core (registered, computed analytically per path from a deterministic rerun): a core of share k gives k + (1 − k) × the ladder's ratio on every path. The deterministic rerun reproduced the primary G exactly (`analysis/f010_core.py`).

Table 13. G with a never-sold core. Intramonth label applies (Section 12).

Rule and core	EUPH_FLOOR	EUPH_MISS	WILD	BASE	UPSHIFT	NO_REVERSION	S
L2-R1, no core	1.52	1.47	1.96	3.94	1.38	0.52	0.463
L2-R1, 10% never sold	1.47	1.42	1.87	3.69	1.35	0.78	0.472
L2-R1, 25% never sold	1.39	1.35	1.74	3.30	1.30	0.90	0.436
L2-RP, 10% never sold	1.68	1.38	2.13	2.89	1.09	0.96	0.424

A 10% core (the owner's earlier idea, applied to the ladder) lowers G by 3% to 6.5% in the euphoria and statistical worlds, about a tenth of the gain over holding. It lifts L2-R1 in NO_REVERSION from 0.52× to 0.78×, just short of the 0.80 line, so L2-R1 would still carry its stress label (a 25% core reaches 0.90×), and its 10th percentile there rises from 0.03 to 0.13. The overall score moves from 0.463 to 0.472, a small difference with no interval computed. This is a secondary analysis and does not change the registered choice.

Certainty equivalent (CRRA γ = 2, real wealth at month 144, relative to holding). L2-R1's results point the same way as its G, and are smaller in the wide worlds (BASE 2.88 against G 3.94; WILD 1.78 against 1.96):

Table 14. L2-R1's certainty equivalent relative to holding (CRRA γ = 2, real wealth at month 144). Intramonth label applies (Section 12).

World	EUPH_FLOOR	EUPH_MISS	WILD	BASE	NO_REVERSION
CE relative to holding	1.50	1.46	1.78	2.88	1.02

At month 144, in NO_REVERSION, the point measure and the cycle average differ (Section 7).

Diagnostics (registered):

- **Realized volatility:** the cycle worlds ran at 0.30 to 0.54 a year, the statistical worlds at 0.42 to 0.83, against history's 0.86 (2012 to 2026) and 0.47 (2022 to 2026).
- **Refit trend:** the refit OLS trend fell to 0.64× to 0.70× the published trend by month 144 in DECAY and CALM_FLOOR. In the refit-line variant, that puts the 2× rung within reach of calm peaks, which is why L2-R1 traded and gained there (CALM_FLOOR 1.26×).
- **Phase:** 55.6% of path-months in the 120 to 168 window were in a rising leg, in every cycle world.

10. History

History is description, not a test (proposal Section 2.2). History since 2012 holds three euphoric cycles, too few to rank rules. All rules use walk-forward lines and cash at the actual T-bill rate. Ladders fill on daily closes at the order level; Kelly rules rebalance monthly.

Table 15. History: Bitcoin against holding. A cycle runs from one halving month to the next; the 96-month windows are 72 overlapping lump sums from 2012-11. The last column runs the ladders on the in-sample published trend (registered sensitivity H4).

Rule	From 2012-11	From 2016-07	From 2020-05	2013 cycle	2017 cycle	2021 cycle	96-month windows: median / share ahead	Published trend: from 2012-11 / from 2016-07
L2-R1	0.25	1.11	1.07	0.23	1.04	1.07	1.11 / 82%	1.01 / 1.70
L2-RF	0.93	2.09	1.86	0.44	1.12	1.86	2.09 / 92%	1.97 / 2.38
L2-R3	0.52	1.98	1.39	0.26	1.42	1.39	1.98 / 89%	1.65 / 2.33
L2-RP	0.94	1.86	1.30	0.52	1.43	1.30	1.86 / 88%	2.53 / 1.60
L1.5-R1	0.12	0.79	0.86	0.16	0.92	0.86	0.79 / 10%	0.49 / 1.22
L1.5-RF	0.40	1.27	1.58	0.32	0.80	1.58	1.27 / 82%	0.85 / 1.53
L1.5-R3	0.30	1.64	1.16	0.18	1.41	1.16	1.64 / 83%	0.71 / 1.48
L1.5-RP	0.41	1.46	1.07	0.29	1.36	1.07	1.46 / 85%	1.45 / 1.74
FK 1	1.53	1.46	1.10	1.05	1.33	1.09	1.43 / 100%	n/a
TA 1	1.09	1.11	1.06	0.98	1.05	1.06	1.11 / 100%	n/a
HKw 1	1.53	1.27	1.14	1.21	1.11	1.14	1.27 / 99%	n/a

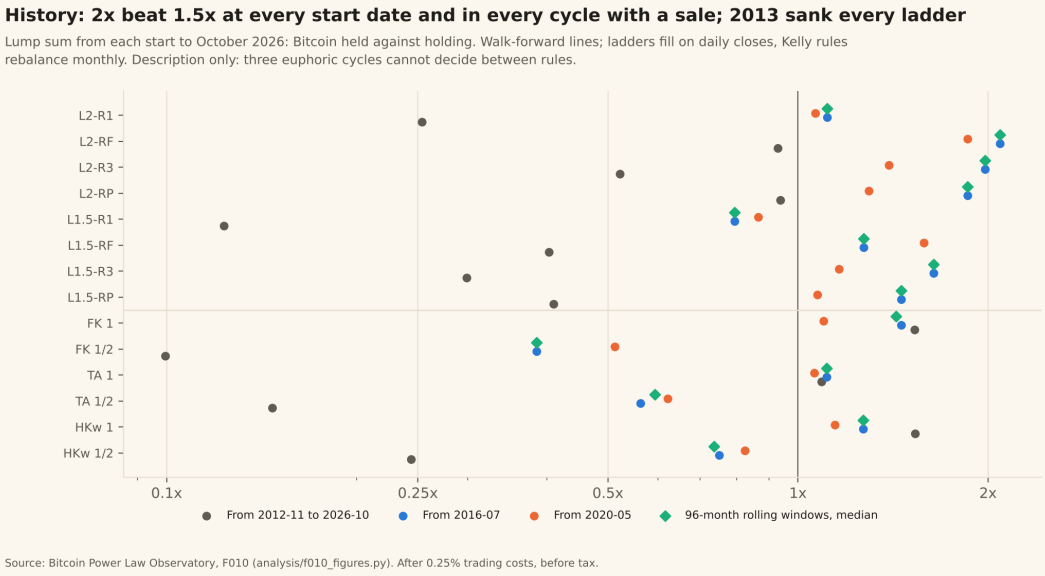


Figure 10. History: results by rule and start date.

What history shows:

- **2× beat 1.5× at every start date with a sale, in every cycle with a sale and in the median rolling window** on walk-forward lines. The 2024-04 start is a tie, because no ladder traded.
  - In individual rolling windows, 1.5× sometimes won (5 to 13 of 120 48-month windows per pair).
  - On the published trend from 2016, L1.5-RP (1.74) beat L2-RP (1.60).
- **2013 sank every ladder on walk-forward lines, but mostly because of the line.**
  - In March 2013 the walk-forward trend, fitted on less than three years of data, stood at 0.69× the published trend. The sales at 2×, 3×, 4.5× and 6.75× that line, on 19 and 30 March and 7 and 9 April 2013 ($59 to $230), were 1.4× to 5.0× the published trend.
  - Over the next 17 months the walk-forward trend rose about 17×, to 2.33× the published trend. R1 bought back on 14 August 2014 at about $510.
  - On the published trend, L2-R1's 2013 cycle was 0.59 rather than 0.23, and from 2012 it was 1.01 rather than 0.25.
  - The fast clock of 2013 (3.1× a year) adds to the loss, but the immature line is the larger part.
  - The trades, line ratios and per-cycle published-trend results are stored in `analysis/f010_history_2013.json` (`analysis/f010_history_2013.py`).
- **Since 2016 the floor came after every euphoria,** so the floor buy-back did best (2.09×). The simulation's MISS worlds, where it does not, are why R1 wins there.
- **History has had only a near-miss.** The 2018 to 2019 low was 0.54× the published trend (0.48× the walk-forward trend) and missed the floor. COVID then took the price to it in March 2020 (Figure 11). Without COVID, L2-RF would have held the cash from its late-2017 sales until the 2022 floor touch.
- **Daily fills matter in history too.** With month-end fills, L2-RF from 2016 fell from 2.09× to 0.44×, because the floor touches in March 2020 and June 2022 were brief. L2-R1 fell from 1.11× to 1.07×.
- **Full floor-Kelly was steadier than any ladder in history:** it was ahead in 100% of 96-month windows, with a 10th percentile of 1.23. In the simulated worlds that miss the floor, it was not.

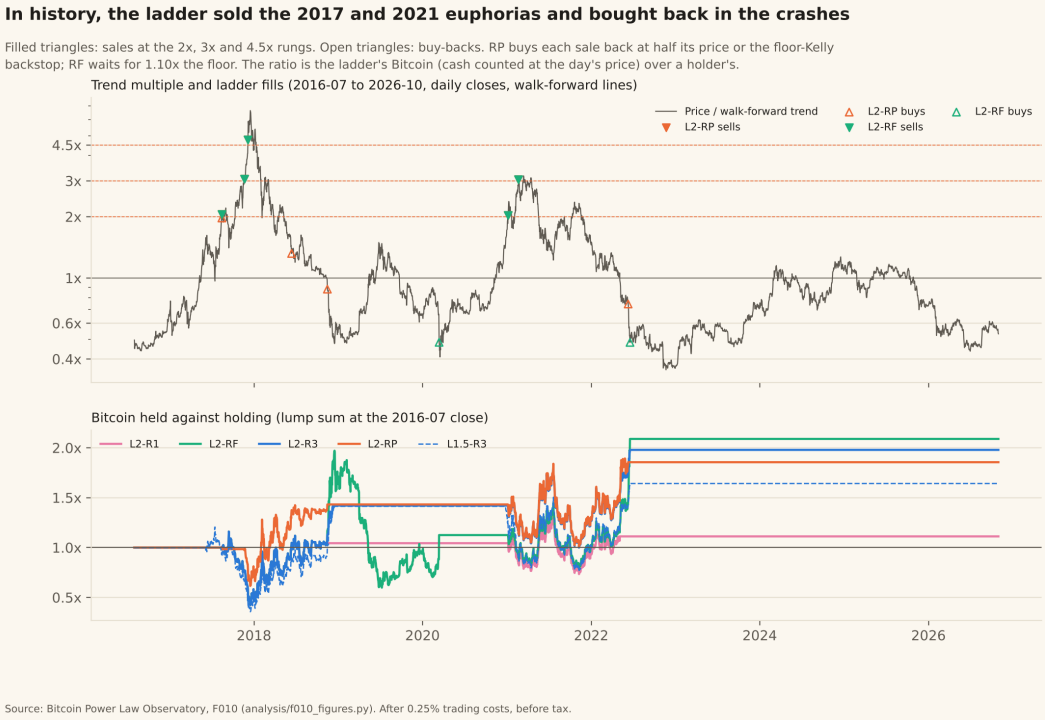


Figure 11. In history, the 2× ladders sold the 2017 and 2021 euphorias and bought back in the crashes. Top: L2-RP and L2-RF trades on walk-forward lines. Bottom: Bitcoin against holding from 2016-07, for the four 2× ladders and L1.5-R3.

11. Scorecard Against the Plan

Table 16. Pre-registered predictions.

#	Prediction	Informed or built in	Result	Verdict
P1	In CALM_FLOOR the best 1.5× ladder has G ≥ 0.97, and FK 1 has G ≤ 0.90	FK threshold informed (stylized band 0.69×)	Best 1.5× ladder 1.65 (L1.5-RF); FK 1 1.48	**Not supported** (FK part)
P2	2× beats 1.5× for at least 3 of 4 buy-backs	Informed (history)	4 of 4 overall; 1 of 4 inside the cycle family	**Supported**
P3	In EUPH_FLOOR all 8 ladders have G > 1.00, and the best ≥ 1.20	None	1.52 to 3.25	**Supported**
P4	In EUPH_MISS, L2-RP and L1.5-RP both have G ≥ 1.00	RF's failure built in	1.42 and 1.18 (RF 0.55 and 0.39)	**Supported**
P5	TA leads S_stat; a ladder leads S_cyc; a ladder is admissible	TA clause informed (Paper 21)	L2-R3 leads S_stat, with L2-R1 level. Among fully admissible rules, L1.5-R1 leads S_cyc (L1.5-RP within the cycle family). Five ladders admissible	**Not supported** (TA clause)
P6	HKw ½ is not admissible	Informed (stylized tests)	Fails EUPH_MISS, BASE, DAMPED, FADING	**Supported**
P7	In UPSHIFT, R1 has the highest G of the four buy-backs, for both first rungs	RF and R3 being stuck is built in	1.38 against 1.09 / 1.01 / 0.57; 1.05 against 0.81 / 0.77 / 0.33	**Supported**

12. What Was Known When, and Deviations

Before the proposal (proposal Section 2.1):

- Paper 21's results on the reused statistical paths;
- a stylized smooth band, an informal calculation made before the proposal and not in the repository (floor-Kelly 0.69×);
- a quick history check of three buy-back rules on the in-sample published trend.

During the five proposal reviews (6.8, 8.0, 8.4, 8.9 and 9.1 out of 10), the design changed:

- the core was removed (an exact blend);
- the worlds became a 2 × 2 plus stress worlds;
- the primary became a cycle-averaged geometric mean;
- daily paths were resampled from history;
- every rule was put on one set of lines;
- the stress worlds stopped gating.

All of this happened before any registered run. Code tests on 6 to 40 paths ran with outputs not inspected. **Registration** is commit 415dd52; the primary simulation, the history tests and the sensitivities then ran with the registered settings.

Added in v1.1, after the v1.0 review (all registered outputs that v1.0 had left out, computed from the stored runs or a deterministic rerun):

- the 30-year and month-144 secondary choices (Section 9.2);
- the never-sold core (Section 9.3);
- per-world intervals, the certainty equivalent and the registered diagnostics;
- the published-trend history (H4).

Exploratory analysis A14 (gains by rung) was not run.

Intramonth label. The registered per-world check compares each world's intramonth excursions, normalised by the month's sd, with history's by size of move. Every world fell outside the 0.85 to 1.15 tolerance in at least one bin:

Table 17. The registered intramonth check: bins outside the 0.85 to 1.15 tolerance, as ratios of simulated to historical excursions.

Worlds	Bin	Ratio to history	Bias
Cycle worlds	Lows in the largest-move months (|z| > 1.5, about 16% of months)	1.25 to 1.34	Dips too deep
Statistical worlds	Lows in the largest-move months	1.39 to 1.46	Dips too deep
Statistical worlds	Highs in the largest-move months	0.74 to 0.80	Spikes too small
BASE, DAMPED, BROKEN	Highs in the smallest-move months (|z| < 0.5)	1.15 to 1.16	Spikes slightly too large

- **Direction of bias:** in big months, simulated dips are deeper than history's, which helps buy orders fill. In the statistical worlds, big-month spikes are smaller, which means fewer sales.
- **Likely cause:** the check pools up and down months, and 23 of history's 29 largest-move months were up months.
- **What the rule says:** results stand with this label. The shape-scale and Gaussian-bridge sensitivities (Section 9.1) show how far they move, and with smaller intramonth moves trend-anchored Kelly overtakes L2-R1.

13. Limitations

- **The power law is the anchor.**
  - Every simulated world except BROKEN and UPSHIFT assumes the published trend holds. The model's out-of-sample R^2 is 0.546, against 0.956 in-sample (S097; Paper 12).
  - A declining slope (C7.4, OQ16) is represented only by BROKEN; UPSHIFT is the opposite case, with multiples ratcheting up.
  - R1's buy level is the trend itself: if the trend is wrong, so is R1's buy level.
- **The floor is a level, not a boundary** (`research/floor-definition.md`). Bitcoin has spent 3.7% of days since July 2010 below the published floor, in 23 visits, none longer than 86 days; since 2016, 0.3% of days, the longest visit 7 days. That record is in-sample: the floor was fitted on the

same history. The floor worlds here draw troughs of 0.38× to 0.48×, so the price goes below the floor in them too, and the noise can take it lower; BASE spends 16% of months below it.

- **The decay of peaks is uncertain.** DECAY and the calm worlds rest on it, and history has only four peak-to-peak transitions (Section 1.2). Paper 19 (Limitations 6) states that each single cycle-to-cycle transition of the Observatory's decay estimate has a confidence interval that spans zero; the interval analysis behind that statement is a planned scan (S017) and is not in this repository.
- **Execution.** The ladder's lead over trend-anchored Kelly depends on daily fills and on the intramonth model (Section 9.1), which the registered check flagged in every world (Section 12).
- **Taxes are out of scope** (OQ18). Each ladder sale may be taxable, which for many investors favours holding.
- **Behaviour.** Buying back after a 50% fall is hard. R1's single order at the trend is simpler to follow than a deep order that waits for years.
- **The worlds are beliefs.** Twelve worlds and equal family weights are a choice. The full table lets readers reweight, and the flip weights show how far the weights would have to move (Table A3 in `analysis/f010_tables.md`).
- **Four cycle shapes.** All cycle worlds are built from four shapes, so their swings within a cycle are history's. The calm-world results for floor-Kelly and the 1.5× ladders depend on that texture.
- **The statistical worlds are wider than history** (BASE residual sd 0.346 against 0.283), which favours selling high and buying low. L2-R1's large gains there (BASE 3.94×) should be read in that light. The statistical family carries half the weight, but it supplies nearly three quarters of L2-R1's score (0.335 of 0.463).

14. Conclusion

For an investor who holds Bitcoin and wants to be ready for another euphoria without forecasting one, the pre-registered test recommends a simple rule:

- sell half at 2× trend, half again at 3×, and so on;
- buy everything back when the price returns to the trend.

Of seven rules that passed the gates in every core world and beat holding in both families of futures, it had the highest score. One rule, L2-R3, which buys back in thirds at 1.0×, 0.75× and 0.5× trend, was level with it within simulation precision; the registered rule recommends the simpler of the two. It costs nothing when no euphoria comes, and it keeps about 1.5 times a holder's Bitcoin when one does, whether or not the floor follows.

The recommendation has limits:

- **Execution.** The rule needs standing limit orders that act on daily prices, and its lead over trend-anchored Kelly also depends on intramonth moves being as large as in the main simulation, whose dips in big months run deeper than history's. With smaller moves, or for an investor who acts once a month, trend-anchored Kelly scores higher.
- **Long horizons and stress.** Over thirty years, and in worlds where the multiples wander or ratchet up, the price-anchored variant L2-RP scores higher and is much safer.
- **The floor.** Waiting for the floor to buy back earns most when the floor comes and strands the cash when it does not.

- **The first rung.** 1.5× gained in the calm cycle worlds but lost overall through the statistical and stress worlds.
- **A 10% never-sold core** costs a little in a euphoria and lifts the random-walk world from 0.52× to 0.78×, though the worst 10% of paths there still keep under 13% of a holder's Bitcoin.
- **Scope.** The rule is insurance against euphoria, not protection against the power law failing.
- **Where the gain comes from.** Nearly three quarters of L2-R1's score comes from the statistical worlds, which swing wider than history, and simulated dips in big months run deeper than history's, which helps buy-backs fill. History agrees in direction since 2016, when every 2× ladder beat holding. From 2012, on walk-forward lines, every ladder lost, mostly because of the young trend line. History holds only three euphoric cycles.

What came later. F011 (forthcoming as Paper 23) found the same ladder best among eight rules for a retiree who spends from the stack, with the same dependence on daily execution and on euphorias, and cash that can stay idle for decades if multiples ratchet up. Whether L2-R3, L2-RP or trend-anchored Kelly does better under spending will be tested in a registered addendum (F011-R), not yet run.

Appendix

Table A1. 10th percentile of the per-path ratio (months 120 to 168). Admissibility needs ≥ 0.80 in every core world.

Rule	Calm floor	Calm miss	Euph. floor	Euph. miss	Decay	Wild	Upshift*	Base	Damped	Fading	Broken*	No rev.*
L2-R1	1.00	1.00	1.27	1.24	1.00	1.21	0.88	1.48	1.00	1.32	1.00	0.03
L2-RF	1.00	1.00	1.88	0.47	1.00	0.82	0.45	0.94	0.66	0.67	1.00	0.02
L2-R3	1.00	1.00	1.54	1.10	1.00	1.20	0.78	1.50	1.00	1.20	1.00	0.03
L2-RP	1.00	1.00	1.54	1.08	1.00	1.14	0.65	1.43	1.00	1.14	1.00	0.14
L1.5-R1	1.00	1.00	1.21	1.16	1.00	1.28	0.66	1.26	1.04	1.23	1.05	0.02
L1.5-RF	1.00	0.61	2.12	0.33	1.00	0.59	0.26	0.68	0.43	0.47	1.12	0.01
L1.5-R3	1.00	0.98	1.60	0.95	1.00	1.21	0.53	1.26	0.84	1.10	1.09	0.02
L1.5-RP	1.00	0.99	1.42	0.84	1.00	1.11	0.48	1.04	0.72	0.93	1.09	0.09
TA ½	1.17	1.02	1.27	0.96	1.02	1.08	0.44	0.90	1.00	0.88	1.77	0.01
TA 1	1.00	1.00	1.29	1.16	1.00	1.20	0.93	1.46	1.01	1.25	1.13	0.04
FK ½	0.50	0.22	0.42	0.20	0.66	0.24	0.07	0.25	0.17	0.21	1.56	0.00
FK 1	1.15	0.52	1.09	0.47	1.20	0.67	0.15	0.74	0.35	0.56	1.64	0.01
HKw ½	1.11	0.95	1.04	0.80	1.17	0.81	0.34	0.52	0.62	0.51	1.55	0.01
HKw 1	1.24	1.07	1.16	0.89	1.15	0.94	0.52	0.77	0.75	0.72	1.37	0.01

Further tables. `analysis/f010_tables.md`, generated by `analysis/f010_tables.py`, holds:

- scores with bootstrap intervals;
- flip weights and within-family scores (A3);
- the 2 × 2 contrasts (A4);

- cash and traded shares (A5, A7);
- timing value (A6);
- month-end fills (A9a);
- other horizons (A3b);
- retiree and saver results (A8);
- exploratory rules (A13);
- diagnostics (A12);
- every sensitivity;
- history H1 to H4.

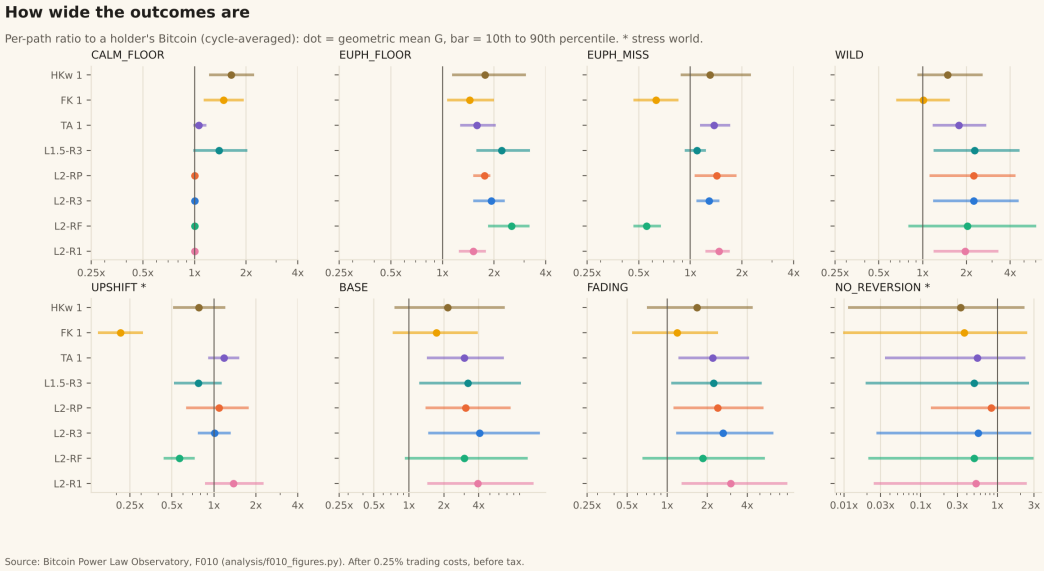


Figure 12. How wide the outcomes are: per-path ratio, geometric mean and 10th to 90th percentile.

Revision notes

Version 1.3 (2026-11-18) reruns the simulation after a correction to the block bootstrap inherited from Paper 21.

- **The defect.** The statistical worlds draw 12-month blocks of Paper 21's monthly shocks. A block could start only in the first 164 of the 176 months, so the first and last 12 months were drawn less often (the very last month never). Both stretches had below-average shocks, so the drawn shocks averaged about 0.003 ($\log_{10}$) a month above the pool's mean. That set the statistical worlds' prices about 10% higher against the trend than the fitted model says. The cycle worlds take only their inflation paths from the draw. The fix (circular blocks, found on 2026-11-17 while porting the engine to the site) is logged in `research/F010-preregistration.md`; `F009_BLOCKS=legacy` reproduces v1.2.
- **Rerun:** the primary, every sensitivity, the never-sold core and the core-sale exploration, with the registered seeds. History is unchanged.
- **The decision.** The registered rule still recommends L2-R1 (S 0.463, was 0.493). L2-R3 (0.459) is now a co-winner: its interval against L2-R1 is -0.009 to +0.001 (was -0.029 to -0.017). The rule picks the simpler co-winner, so the recommendation does not change, but the abstract, Key

findings, Sections 3 and 14 and C9.1 no longer call L2-R1 the only co-winner. The winner between the two now flips with small weight changes (Section 3).

- **Other changes:** L1.5-R3 becomes admissible (eight admissible rules, seven beating holding in both families); L2-RF now fails three core worlds instead of four; over 30 years L2-RP leads clearly (0.821 against 0.763, was 0.831 against 0.827); with a 10% never-sold core L2-R1 reaches 0.78× in NO_REVERSION, short of the 0.80 stress line (was 0.83×); with month-end fills L2-R1 scores 0.250 and TA 1 0.374 (were 0.260 and 0.400); the statistical family supplies 0.335 of L2-R1's 0.463 (was 0.364 of 0.493); RP at 70% depth is now admissible; NO_REVERSION's medium-month lows are now inside the intramonth tolerance, and the smallest-move highs of BASE, DAMPED and BROKEN are slightly outside it (Table 17); P5 still fails, now because L2-R3 and L2-R1 lead the statistical family; F011's figures are its rerun values.
- **Unchanged:** every cycle-world G (Table 4's cycle columns and Table 6), P1 to P7's verdicts, and all history results. The retiree and saver results move slightly in the cycle worlds too, because their inflation paths come from the same draw (Table 10).

Version 1.2 (2026-11-10) prepares the report for publication as Paper 22 and answers an independent review of v1.1. It:

- **Labels:** names the report Paper 22 and the earlier study Paper 21, lists the claims (C9.1 to C9.6) and corrects the review history of the proposal (five rounds).
- **Floor:** describes the floor as a level, not a boundary (`research/floor-definition.md`), with its record below the published floor; states that the 2022 and 2026 lows closed slightly below it.
- **Caveats in the abstract and conclusion:** the share of L2-R1's score from the wider statistical worlds, the direction of the intramonth bias, the floor buy-back's loss of its historical gain with month-end fills, and the three euphoric cycles in history.
- **Later work:** adds what F011 (forthcoming as Paper 23) found for a retiree who spends from the stack, and the registered addendum F011-R.
- **Clarity:** adds a Terms section, defines the score in the abstract, corrects "closest rival" for trend-anchored Kelly and gives the full flip weights.
- **Minor items:** the UPSHIFT pool share (99.9%), the 1.5× column of Table 11 (then Table 9), the walk-forward trend's rise in 2013 (0.69× to 2.33× the published trend), the UPSHIFT and floor-Kelly wording, month-end fills for the saver, OQ18, order re-pricing, the noise and RP-depth results, citations of the references, the informal stylized test, figure titles, table numbering (Table 10b became a numbered table) and the data path.
- **PDF:** fixes the running header, the lists and the width of Table 4.

A second review of v1.2 led to these further changes, also in v1.2:

- **Numbers:** the statistical family's share of L2-R1's score is 0.364 of 0.493 (it was printed as 0.365); L2-RP's UPSHIFT results are given at both horizons (1.09× at 10 to 14 years, 1.49× at year 30); the floor buy-back's month-end result since 2016 is given as a fall from 2.09× to 0.44×.
- **History:** the conclusion states both what history since 2016 shows and why every ladder lost from 2012; the L1.5-RF row is added to the history table.
- **Terms:** rule names and the two family scores are defined; "only co-winner" is explained in the abstract (superseded in v1.3, where L2-R3 is also a co-winner).
- **Sources:** the decay-uncertainty limitation now says where the statement comes from (Paper 19, Limitations 6) and that its interval analysis (S017) is a planned scan.

- **Wording:** Section 1.2 (“if the shrinking stops”), the trend-growth formula, and the conclusion’s “about 1.5 times” and “passed the gates”.
- **Tables and figures:** every table now has a numbered caption (the four unnumbered tables became Tables 5, 7, 14 and 17, and later tables were renumbered); column widths and Table 4’s spacing were adjusted; captions are kept with their figures and tables; the Figure 3 and Figure 11 captions say what they show; Figures 8 and 12 have distinct colours for every rule, readable axis labels and, in Figure 12, ticks below 0.25×.

Version 1.1 answered an independent review of v1.0. It:

- **Trend-anchored Kelly:** shows TA 1 next to the ladders in every sensitivity, and states that TA 1 leads L2-R1 under month-end fills or smaller intramonth moves (Table 11, Figure 9, abstract, conclusion).
- **Secondary choices:** reports the registered 30-year and month-144 secondary choices (L2-RP leads on 30 years), the 30-year stress results, the never-sold core, the certainty equivalent and the registered diagnostics.
- **First rung:** corrects the explanation. 1.5× wins inside the cycle family; 2× wins overall through the statistical family, admissibility and stress.
- **2013:** corrects the account. It was mostly the immature walk-forward line, with dates fixed and the published-trend history added.
- **Wording:** corrects the uniqueness, family and “most robust” statements.
- **Intramonth label:** reports it in full, with its direction of bias.
- **Minor items:** fixes the DECAY definition, “every window”, “never stuck”, the informed/built-in column, the 10th-percentile appendix table, figure and table numbering, citations (R043 removed; S097 source added), retiree censoring, the timing-value wording, the statistical worlds’ anchor (0.92× trend) and the peak ratios.

References

- Bertram, W. K. (2010). Analytic solutions for optimal statistical arbitrage trading. *Physica A* 389(11), 2234–2243.
- Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman and Hall.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer.
- Kelly, J. L. (1956). A new interpretation of information rate. *Bell System Technical Journal* 35(4), 917–926.
- Leung, T. and Li, X. (2015). Optimal mean reversion trading with transaction costs and stop-loss exit. *International Journal of Theoretical and Applied Finance* 18(3), 1550020.
- 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), 681–687.
- Rotando, L. M. and Thorp, E. O. (1992). The Kelly criterion and the stock market. *American Mathematical Monthly* 99(10), 922–931.
- Bitcoin Power Law Observatory: Paper 12 (Formal Verification; out-of-sample R^2, S097); Paper 19 (Limitations 6); Paper 21 (Floor-Kelly With Money In and Money Out, `report/kelly_cashflow.md`); F011 (`report/spend_floor_growth.md`, forthcoming as Paper 23) and its

addendum F011-R (`research/F011R-preregistration.md`); `research/floor-definition.md`; `research/claims.md` (C7.4, OQ16, OQ18); `research/pipeline.json` (S017, a planned scan).

Data and Reproducibility

Data:

- Bitcoin: 5,950 daily closes from 2010-07-18 to 2026-10-31 (`bitcoin-powerlaw-site/shared/datasets/btc_historical.json`);
- monthly T-bill rate (FRED TB3MS) and CPI-U (FRED CPIAUCSL), as in Paper 21.

Model:

- power-law parameters logA = -16.493, β = 5.688, genesis 2009-01-03;
- published floor 0.42× trend (the Observatory's floor_current of 0.432× is not used here);
- walk-forward lines as in Paper 21.

Commands:

```
python3 -W ignore analysis/f010_mc.py 2000 4 primary   # analysis/f010_mc_results.json (45 minutes on 4 cores)
python3 -W ignore analysis/f010_mc.py 2000 4 sens      # analysis/f010_mc_results_sens.json (31 minutes)
python3 -W ignore analysis/f010_ladder.py              # analysis/f010_history.json
python3 -W ignore analysis/f010_core.py                # analysis/f010_core.json (never-sold core)
python3 -W ignore analysis/f010_tables.py              # analysis/f010_tables.md
python3 -W ignore analysis/f010_figures.py             # figures/f010/
```

Seeds:

- cycle worlds 20261108;
- statistical worlds 20261107 (Paper 21's paths);
- bootstrap (20261108, 7).