Bitcoin Retirement Playbook Optimization

Floor Definitions, Sell-Tier Architecture, Floor Leverage, and Volatility Stress Analysis

Authors: Scale Invariant Capital — Bitcoin Power Law Observatory **Paper:** P17 **Version:** 1.1 **Date:** 2026-05-22 **Path budget:** 12.4 million simulated futures **Builds on:** P16 — *Passive vs Active Retirement Under Power-Law Growth*

TL;DR

We extend the wealth playbook framework from P16 with three targeted refinements and one stress test, using 50 000-path Monte Carlo simulations across 248 distinct scenarios:

1. **Floor definition sensitivity.** The canonical 0.432× floor multiplier corresponds to approximately the empirical 5th percentile (P5 = 0.4295×, $22 030/yr/BTC). P1 = 0.3133× anchors a more conservative $16 068/yr withdrawal; P10 = 0.4598× supports $23 583/yr. Playbook survival is 100% across all three definitions at 6% inflation and 30- and 50-year horizons — floor choice determines initial withdrawal size, not long-run safety.
2. **BTC stack growth benchmark.** The playbook delivers a net Bitcoin CAGR of ≥10%/yr in **80% of 20-year paths and ~100% of 50-year paths** at historical volatility, with a median of 20–26% across horizons. This is the empirical answer to "how likely is it to grow the stack at 10%/yr net?" under the power-law model.
3. **Volatility stress.** The playbook's edge collapses sharply between 0.5–0.7× historical step-sigma. At 0.7× (30% vol compression), median CAGR drops to ~10% and P(≥10%) falls to 48%. At 0.5×, median CAGR is only 3% and P(≥10%) is below 2%. Survival stays ≥99.7% at all tested levels because the floor barrier and conservative withdrawal rule prevent ruin regardless of playbook performance.
4. **Floor leverage.** Borrowing 20% of current stack value when price falls below 0.5× trend — buying more BTC at floor, repaying at trend — adds approximately **+10% of stack value net per floor visit** (+0.10 BTC on a 1-BTC stack; scales proportionally with stack size). This rule engages even in low-volatility regimes where euphoria-phase sell tiers rarely fire, providing a residual growth engine independent of cycle amplitude.
5. **Optimal sell-tier architecture.** The optimal configuration across all tested volatility regimes is **seven tiers (1.1 / 1.2 / 1.3 / 1.4 / 1.5 / 2.0 / 3.0× trend) at a uniform 30% of remaining stack per tier.** This outperforms the P16 baseline (1.5 / 2.0 / 3.0× @ 50%) by 0.6–1.3 pp in median CAGR across all three tested volatility levels. Key mechanism: additional lower tiers capture premium in cycles where price never reaches 1.5× trend, while 30% per tier preserves enough stack to benefit from 2× and 3× events when they do materialise.

Abstract

We extend the power-law retirement framework of P16 with four complementary analyses. First, we compare three empirically-derived floor definitions — P1 (0.3133×), P5 (0.4295×), and P10 (0.4598×) of the historical log-residual distribution — as anchors for the floor-growth coverage rule, finding that active playbook strategies achieve near-universal survival regardless of floor definition, while the choice affects only initial withdrawal size. Second, we quantify the distribution of annualised Bitcoin stack growth (CAGR) across playbook survivors, finding that the commonly-cited "10%/year BTC net return" benchmark is achievable with 80–100% probability over 20–50 year horizons under historical volatility. Third, we stress-test the playbook's sell-tier mechanism against a range of volatility compression scenarios (σ-multiplier 0.1–1.0×), identifying a breakdown threshold between 0.5× and 0.7× historical step-sigma. Fourth, we optimise the sell-tier architecture and floor leverage sizing, finding a robust optimum at seven tiers (1.1–3.0× trend) at a uniform 30% of remaining stack, with floor leverage capped at 20% of collateral value below 0.5× trend.

Keywords: Bitcoin retirement, floor-growth coverage rule, wealth playbook, sell-tier optimisation, floor leverage, volatility stress, BTC CAGR, Ornstein–Uhlenbeck, power-law model.

---

1. Introduction

P16 established that an active wealth playbook — borrow below 1.6× floor, settle debt at trend, sell 50% at each of 1.5/2/3× trend, redeploy all fiat at trend — can sustain withdrawal rates approximately 50–75 percentage points higher than passive hold at comparable confidence levels. Four questions remained.

Q1 — Floor definition. The simulations in P16 used the canonical 0.432× floor multiplier. Three competing definitions are used in practice: the empirical 1st percentile of log-residuals (P1), the 5th percentile (P5), and the 10th percentile (P10). Which should anchor a retirement product, and does the choice materially affect survival?

Q2 — BTC growth benchmark. Active investors commonly assess strategy quality in *Bitcoin terms* rather than dollar terms. The informal benchmark "10% net BTC/year" is cited frequently as the threshold between good and excellent outcomes. What probability does the playbook actually assign to achieving that benchmark?

Q3 — Playbook robustness. Bitcoin's per-cycle volatility has compressed monotonically (P11: *Volatility Decay*). If future cycles are substantially quieter, the sell tiers at 1.5–3× trend may fire rarely. At what volatility level does the playbook stop working?

Q4 — Tier architecture. The P16 sell configuration (50% at each of three tiers starting at 1.5×) was motivated by simplicity. Is there a better tier structure, particularly for compressed-volatility environments?

This paper addresses all four questions with targeted Monte Carlo experiments on the same OU price model used in P16, extending it with per-tier sell percentages, floor leverage, and parameterised volatility scaling.

---

2. Model and Methods

2.1 Price process

The log-residual $r_t = \log_{10} P_t - \log_{10} P_{\text{trend}}(d_t)$ follows a monthly Ornstein–Uhlenbeck process:

$$r_{t+1/12} = \mu_c + (r_t - \mu_c)\,e^{-\kappa/12} + \sigma_c \sqrt{\tfrac{1-e^{-\kappa/6}}{2\kappa}}\; \epsilon_t,$$

with $\kappa = 1.0$, cycle-specific means μ_c and monthly step-sigmas σ_c, Student-t innovations ϵ_t with cycle-specific degrees of freedom, and annual cycle assignment drawn from the recency-weighted blend (cycle 4: 60%, cycle 3: 25%, cycle 2: 10%, cycle 1: 5%). Full parameter values are tabulated in Appendix A.

The floor boundary is fixed per scenario at a constant log-residual corresponding to the chosen floor percentile (P1, P5, or P10). When the OU process produces a residual below this level, it is reflected symmetrically: $r \leftarrow 2r_{\text{floor}} - r$. The trend follows $\log_{10} P_{\text{trend}}(d) = \log A + \beta \log_{10} d$, with $\log A = -16.493$, $\beta = 5.688$, d days since the genesis block. All simulations use start date 2026-05-22 ($d = 6{,}347$), giving $P_{\text{trend}} = \$136{,}899$.

2.2 Withdrawal rule

The initial annual withdrawal anchors to floor-growth coverage:

$$W_0 = f_{\text{def}} \cdot \bigl[P_{\text{trend}}(d_0 + 365) - P_{\text{trend}}(d_0)\bigr],$$

where f_{def} is the chosen floor definition multiplier. Subsequent years inflate at a fixed rate π: $W_t = W_0 (1+\pi)^t$. At the start date, floor-growth withdrawals per BTC per year are:

Floor	Multiplier	Log-residual	Initial withdrawal
P1	0.3133	−0.5041	$16,068 /yr
P5	0.4295	−0.3670	$22,030 /yr
P10	0.4598	−0.3374	$23,583 /yr

2.3 Playbook rules

The extended playbook operates in priority order each month:

1. **Debt accrual:** outstanding debt compounds at 10% APR monthly.
2. **Redeploy at trend:** when, after a sell cycle (at least one tier has fired), price returns to ≤ 1.0× trend, use fiat to pay debt then convert remaining fiat to BTC. Reset all sell-tier and floor-leverage flags. *This is the cycle-close trigger — rule 3 does not apply during an active sell cycle.*
3. **Settle debt at trend:** between cycles (when no tier has fired this cycle), pay any outstanding debt from fiat then BTC when price is ≥ 1.0× trend.
4. **Sell tiers:** fire each tier once per cycle when price first crosses the threshold; sell a fixed % of remaining stack.
5. **Floor leverage** *(new)*: fire once per floor visit when price first falls below 0.5× trend; borrow a fixed % of current stack value, buy BTC immediately. Capped at 50% LTV on current collateral. Resets at next redeploy.

6. **Withdrawal sourcing:** fiat first; then borrow (if below 0.8× trend and within 50% LTV); then sell BTC.
7. **Ruin check:** net worth (stack × price + fiat − debt) ≤ 0 triggers ruin.

2.4 Volatility scaling

To stress-test against future volatility compression, we multiply all per-cycle step-sigmas by a scalar $\sigma_{\times} \in \{0.1, 0.3, 0.5, 0.7, 1.0\}$ before running simulations. $\sigma_{\times} = 1.0$ corresponds to the historical calibration; $\sigma_{\times} = 0.1$ represents a near-flat-price regime with minimal cyclical movement.

2.5 BTC CAGR

For surviving paths at horizon T, we compute the annualised BTC stack growth:

$$g = \left(\frac{S_T}{S_0}\right)^{1/T} - 1,$$

where $S_0 = 1.0$ BTC and S_T is the terminal stack. We report the median, 10th and 90th percentiles, and the fraction of survivor paths with $g \geq 10\%$ and $g \geq 5\%$.

2.6 Simulation parameters

All experiments: 50 000 paths, P5 floor definition unless otherwise stated, 6% inflation, 1 BTC starting stack.

---

3. Floor Definition Sensitivity

3.1 Survival rates

Table 1 reports survival probabilities for passive hold (baseline) and the full playbook at 30- and 50-year horizons, 6% inflation, 1 BTC, across all three floor definitions.

Table 1 — Survival probability by floor definition (6% inflation, 1 BTC)

Floor	Horizon	Baseline	Playbook	Δ
P1	30y	99.9%	100.0%	+0.1 pp
P1	50y	99.8%	100.0%	+0.2 pp
P5	30y	97.4%	100.0%	+2.6 pp
P5	50y	94.4%	100.0%	+5.6 pp
P10	30y	95.2%	100.0%	+4.8 pp
P10	50y	90.7%	100.0%	+9.3 pp

Three observations stand out. First, even the baseline shows very high survival at P1 because the $16 068/yr withdrawal is conservative enough that the stack depletes slowly relative to price growth. Second, P10 at 50 years shows 9.3 pp of baseline ruin — meaningful for a product — but the playbook eliminates it. Third, 100% playbook survival holds across all floor definitions and tested horizons.

3.2 Practical floor choice

The floor definition determines only the initial withdrawal size, not long-run safety. From a product-design standpoint:

- **P1 (0.3133×, $16 068/yr)** is the maximally conservative anchor. Corresponds to the deepest observed floor over 15 years of data. Appropriate for a product targeting near-zero ruin probability with passive management.
- **P5 (0.4295×, $22 030/yr)** is the practical middle ground, closely matching the previously-published 0.432× convention. Appropriate as the default for active-playbook products.
- **P10 (0.4598×, $23 583/yr)** supports a modestly higher initial lifestyle but requires the playbook to perform. Not recommended for passive strategies at 50-year horizons.

All subsequent analysis uses P5 as the default floor definition.

3.3 Inflation and lifestyle sensitivity

The playbook dominates inflation risk (Table 2). At 10% annual inflation — the most aggressive tested — the 50-year baseline survival collapses to 17.9%, but the playbook holds at 99.98%.

Table 2 — Inflation sensitivity (P5, 1 BTC, playbook vs baseline)

Inflation	30y Baseline	30y Playbook	50y Baseline	50y Playbook
6%	97.2%	100.0%	94.4%	100.0%
7%	93.3%	100.0%	85.2%	100.0%
8%	85.4%	100.0%	68.0%	100.0%
9%	73.2%	100.0%	43.0%	100.0%
10%	56.2%	100.0%	17.9%	100.0%

Lifestyle bump scenarios (two-phase early-high rates; glide-path decay models) show the same pattern: baseline survival degrades sharply under aggressive lifestyle growth, but the playbook maintains ≥99.9% survival across all 30-year scenarios and ≥99.9% at 50 years across all lifestyle scenarios tested.

---

4. BTC Stack Growth Distribution

4.1 The 10%/year benchmark

Bitcoin investors commonly assess strategy quality in *Bitcoin terms* rather than dollar terms: a strategy that grows the BTC stack at 10%/year net would roughly double every 7 years, independently of the dollar price. We ask: at what probability does the playbook achieve this?

Table 3 reports the CAGR distribution of the BTC stack across playbook survivors at horizons 20–50 years, P5 floor, 6% inflation, historical volatility. All values are generated from the current codebase (the full 7-tier, floor-leverage playbook).

Table 3 — BTC CAGR distribution (playbook, P5, 1 BTC, 6% inflation, historical volatility)

Horizon	Survival	p10 CAGR	Median CAGR	p90 CAGR	P(≥10%/yr)	P(≥5%/yr)	P(grows)
20y	100.0%	5.6%	19.6%	35.1%	80.4%	91.0%	97.0%
30y	100.0%	11.8%	23.0%	35.3%	93.4%	98.4%	99.8%
40y	100.0%	15.6%	24.9%	35.4%	98.2%	99.8%	100.0%
50y	100.0%	18.0%	26.3%	35.6%	99.5%	100.0%	100.0%

The 10%/yr benchmark is achievable in 80% of 20-year paths and essentially all 50-year paths. The median outcome substantially exceeds the benchmark: 20–26% CAGR across horizons. The convergence of the p10 toward the median over time reflects the averaging of multiple Bitcoin cycles: any single bad cycle is increasingly diluted as the horizon extends.

The primary risk at shorter horizons (20y) is a run of cycles where price oscillates only modestly above trend, so lower sell tiers (1.1–1.5×) fire but 2× and 3× are rarely reached. In that scenario the BTC CAGR relies more heavily on floor leverage, which provides incremental but more consistent gains.

4.2 Interpretation

A median CAGR of 23% over 30 years from a 1 BTC starting position, while withdrawing $22 030/yr (rising at 6% annually), implies that the playbook is not merely a defensive tool. It converts Bitcoin's inherent volatility premium into net BTC accumulation even while funding lifestyle withdrawals. The strategy is simultaneously a retirement income product and a wealth compounding mechanism.

---

5. Volatility Stress Analysis

5.1 Motivation

Bitcoin's per-cycle residual volatility has declined monotonically (P11: *Volatility Decay Analysis*). The monthly step-sigma in the model calibration for cycle 1 is approximately 53% higher than for cycle 4 (0.150 vs 0.098; see Appendix A). P11 documents that the peak-to-trough amplitude of log-residuals across full cycles has compressed by a larger factor, as early cycles combined high step-sigma with higher mean residuals. Extrapolating to cycles 5–7, the sell tiers at 1.5/2/3× trend may fire less frequently. We stress-test by multiplying all OU step-sigmas by a scalar $\sigma_{\times}$, preserving the mean-reversion structure but compressing residual amplitude.

5.2 Results

Table 4 reports the playbook's performance under vol compression, 30- and 50-year horizons, P5 floor, 6% inflation.

Table 4 — Volatility stress grid (playbook, P5, 1 BTC, 6% inflation)

σ-mult	Horizon	Survival	p10 CAGR	Median CAGR	p90 CAGR	P(≥10%/yr)	P(grows)
1.0×	30y	100.0%	11.9%	23.0%	35.4%	93.4%	99.8%
1.0×	50y	100.0%	18.0%	26.3%	35.6%	99.6%	100.0%
0.7×	30y	100.0%	2.5%	9.7%	16.8%	47.7%	95.9%
0.7×	50y	100.0%	7.0%	12.2%	17.6%	70.6%	99.9%
0.5×	30y	100.0%	−2.1%	2.6%	6.9%	1.9%	77.5%
0.5×	50y	100.0%	1.3%	4.6%	8.0%	2.3%	96.0%
0.3×	30y	100.0%	−5.0%	−2.9%	−0.8%	0.0%	4.4%
0.3×	50y	100.0%	−2.7%	−1.1%	0.5%	0.0%	18.2%
0.1×	30y	100.0%	−5.2%	−4.8%	−4.4%	0.0%	0.0%
0.1×	50y	100.0%	−3.6%	−3.3%	−3.0%	0.0%	0.0%

5.3 Critical threshold

The playbook's compounding edge essentially disappears between **0.5× and 0.7× historical volatility**. At 0.7×, approximately half of 30-year paths still achieve ≥10% CAGR — the strategy remains useful. At 0.5×, the 2× tier fires in ~94% of paths and the 3× tier fires in only ~42% of paths (vs ~99.8% at 1.0×), and the strategy's median outcome is barely positive.

Crucially, however, **survival stays at ~100% across all tested sigma levels.** At 0.1× sigma — essentially a bond-like BTC with minimal cyclical movement — the playbook user loses ~4.8%/yr from their BTC stack, simply because they are selling BTC to fund withdrawals with no playbook gains to offset. But they do not go bankrupt, because the conservative P5 withdrawal rule ensures the stack depletes slowly enough to last the full horizon.

5.4 The survival-vs-growth distinction

This result establishes a clean distinction between two separate risks:

- **Ruin risk** (stack goes to zero): controlled by the withdrawal rule, not the playbook. Robust to full volatility collapse.
- **BTC accumulation**: controlled by the playbook. Breaks at moderate volatility compression (0.5–0.7×).

Investors should be aware that the playbook's retirement-safety benefit (near-zero ruin) is more reliable than its wealth-accumulation benefit (≥10% CAGR), and that the two benefits operate through different mechanisms.

5.5 Adoption signals as early-warning indicators

The primary driver of the volatility threshold is S-curve saturation: as Bitcoin adoption approaches its total addressable market, speculative cycle amplitude compresses. The following metrics are *plausible* early indicators — they have not been formally validated in this paper as quantitative predictors — but bear monitoring:

- New-wallet growth rate deceleration relative to the S-curve model

- Layer-2 / settlement-layer transaction ratio: declining L1 speculative activity
- Institutional net flows normalised to circulating supply
- Hash-rate growth vs price premium: widening wedge signals miner-behaviour shift

If these metrics simultaneously plateau, beta drift and/or volatility compression may arrive 1–2 cycles ahead. At that point, the playbook should be recalibrated — specifically, sell-tier thresholds lowered and floor leverage sized upward.

---

6. Floor Leverage Optimisation

6.1 Mechanism

When Bitcoin's price falls below 0.5× trend ($r \leq \log_{10}(0.5) = -0.301$), the playbook enters a *floor leverage* episode: it borrows a fixed percentage λ of the current BTC stack value and uses the proceeds to purchase additional BTC. The episode fires once per floor visit and resets at the next trend redeploy.

Per-episode gain (simplified, ignoring interest): Starting with S BTC at entry price $0.5\,P_{\text{tr}}$:

- Borrow $\lambda \cdot S \cdot 0.5\,P_{\text{tr}}$ dollars; buy λS additional BTC.
- Stack becomes $S(1+\lambda)$; debt = $\lambda S \cdot 0.5\,P_{\text{tr}}$.
- At trend (P_{tr}), repay: sell $0.5\lambda S$ BTC to cover debt.
- **Net gain: $+0.5\lambda S$ BTC, or $+50\%$ of the borrowed fraction of the original stack.**

At $\lambda = 0.20$: net gain = $+0.10\,S$, i.e., **+10% of the current stack per floor visit**, before interest. Ten-percent APR for a 6-month visit costs approximately $0.005\,S$ BTC — negligible against the gain. Note that the absolute BTC gain scales with stack size; a 5-BTC position gains ~0.5 BTC per visit under the same mechanics.

This mechanism provides compounding independent of whether euphoria-phase sell tiers ever fire. In a volatility-compressed world where price oscillates between floor and trend, floor leverage alone delivers meaningful BTC accumulation.

6.2 Sweep results

Table 5 reports the effect of varying λ from 0–50% at historical volatility and half-sigma. All runs use the 3-tier P16 baseline sell configuration (1.5/2.0/3.0× @ 50%) to isolate the effect of floor leverage alone from the tier-architecture changes introduced in Section 7.

Table 5 — Floor leverage sweep (P5, 1 BTC, 6% infl, 30y, 3-tier P16 baseline)

λ	σ=1.0× med CAGR	P(≥10%)	σ=0.5× med CAGR	P(≥5%)
0%	20.1%	89.0%	1.6%	15.1%
10%	21.5%	91.4%	2.1%	19.0%
15%	22.2%	92.7%	2.3%	21.2%
20%	**22.9%**	**93.5%**	**2.6%**	**23.2%**
30%	24.2%	94.8%	3.0%	27.9%
40%	25.5%	96.0%	3.4%	31.7%
50%	26.6%	96.7%	3.7%	35.3%

Under the model's 50% LTV hard cap, survival is 100% at all tested λ values. The gains are monotonically increasing with λ. Comparing λ=0 to λ=20% at σ=0.5×: median CAGR rises from 1.6% to 2.6%, meaning floor leverage contributes approximately 1.0 pp out of 2.6 pp total — **roughly 39% of the σ=0.5× median outcome**. The choice of $\lambda = 0.20$ is therefore a **conservatism parameter, not a model optimum**, appropriate for two reasons not captured by the model:

1. **Lender behaviour:** in real floor events, collateral has declined sharply; lenders may enforce margin calls before the theoretical LTV limit.
2. **Execution uncertainty:** the model fires the borrow at the exact month the floor is crossed; in practice, borrowing capacity may be constrained in distressed markets.

At $\lambda = 20\%$, debt from floor leverage is small enough that even a further 30–40% price decline after entry would not threaten solvency. Investors with higher-conviction credit access and active monitoring can scale toward 30–40%.

---

7. Sell-Tier Architecture Optimisation

7.1 Strategy space

Table 6 evaluates seven sell-tier configurations across three volatility regimes, all with 20% floor leverage at the P5 floor, P5, 1 BTC, 6% inflation, 30y horizon.

Table 6 — Sell-tier strategy comparison (P5, 1 BTC, 6% infl, 30y, σ values)

Strategy	Tiers	Per-tier %	σ=1.0× CAGR	σ=0.7× CAGR	σ=0.5× CAGR
wide-50% (P16 baseline)	1.5 / 2.0 / 3.0	50%	22.9%	9.8%	2.6%
full-30%	1.25 / 1.5 / 1.75 / 2.0 / 3.0	30%	23.2%	10.2%	3.0%
mid-30%	1.25 / 1.5 / 2.0	30%	19.5%	8.5%	2.3%
tight-25%	1.25 / 1.5	25%	13.8%	5.7%	1.0%
tiny-20%	1.15 / 1.25 / 1.5	20%	15.2%	6.4%	1.3%
ladder-v1 (graduated)	1.1–1.5 / 2.0 / 3.0	40/30/25/20/20/20/20%	23.5%	10.3%	2.8%
ladder-v2 (tapered)	1.1–1.5 / 2.0 / 3.0	40/30/25/20/15/10/10%	22.5%	9.8%	2.6%
7-tier 30% (optimal)	**1.1 / 1.2 / 1.3 / 1.4 / 1.5 / 2.0 / 3.0**	**30%**	**24.1%**	**10.8%**	**3.2%**

The **7-tier uniform 30%** configuration consistently leads across all volatility levels. The P16 baseline (wide-50%) is ranked 5th at σ=1.0× and falls further behind as volatility compresses. Two structural failures are evident:

- **Tight two-tier grids** (tight-25%, tiny-20%) substantially underperform despite firing at lower price levels, because the small per-tier sell amounts generate insufficient fiat for meaningful buyback, regardless of how often the tiers fire.
- **Tapered ladders** (ladder-v2) that aggressively reduce the sell percentage at higher tiers sacrifice too much premium at the rare but high-value 2× and 3× events.

7.2 Uniform sell-percentage sweep on the 7-tier ladder

Table 7 fixes the tier grid (1.1 / 1.2 / 1.3 / 1.4 / 1.5 / 2.0 / 3.0×) and sweeps a single uniform sell percentage from 5% to 50%.

Table 7 — Uniform sell% sweep on 7-tier ladder (P5, 1 BTC, 6% infl, 30y)

Sell%	σ=1.0× CAGR	P(≥10%)	σ=0.7× CAGR	P(≥10%)	σ=0.5× CAGR	P(≥5%)
5%	10.4%	55.0%	3.7%	0.6%	−0.2%	0.1%
10%	16.6%	89.9%	6.9%	18.2%	1.4%	5.4%
15%	20.4%	94.1%	8.9%	39.6%	2.3%	17.8%
20%	22.6%	94.8%	10.0%	49.8%	2.9%	27.6%
25%	23.9%	94.3%	10.5%	53.4%	3.1%	32.7%
30%	**24.1%**	**92.6%**	**10.8%**	**54.5%**	**3.2%**	**35.9%**
35%	24.0%	90.6%	10.6%	53.1%	3.1%	37.0%
40%	23.4%	87.5%	10.3%	51.2%	2.9%	36.7%
50%	21.6%	78.8%	9.4%	47.9%	2.4%	36.0%

30% is the median-CAGR optimum at all tested volatility levels. The objective function is concave in sell% with a broad plateau from 25–35% and peak at 30%. One nuance worth noting: by the P(≥5%) criterion at σ=0.5×, the optimum shifts to 35% (37.0% vs 35.9% at 30%). This reflects a fat-tail effect in compressed-vol regimes — selling slightly more early generates more fiat that benefits from the rare 2× or 3× events, widening the right tail of outcomes at a modest median cost. Investors who prioritise a higher probability of any meaningful accumulation in a compressed-vol world could reasonably favour 35%.

7.3 Why 30% is optimal: cycle arithmetic

After 7 tiers at 30% each, the remaining stack is $(0.70)^7 \approx 8.2\%$ of the starting position. The fiat accumulated from selling across all tiers, expressed in BTC-equivalent at trend, is:

$$\text{BTC-equiv fiat} = \sum_{i=0}^{6} 0.30 \times (0.70)^i \times m_i,$$

where m_i is the tier multiple (1.1, 1.2, 1.3, 1.4, 1.5, 2.0, 3.0). For a complete cycle, this sums to approximately 1.23 BTC-equivalent, giving a total of $0.082 + 1.232 = 1.31$ BTC on buyback — a **+31% BTC gain per complete cycle**.

The 30% optimum reflects a balance: lower sell% (e.g., 10%) generates only 0.41 BTC-equivalent fiat per complete cycle; higher sell% (e.g., 50%) generates 1.61 BTC-equivalent fiat but the remaining stack at completion is only $(0.50)^7 \approx 0.8\%$, so the early-tier premium (10–20% above trend) dominates and the high-premium late tiers contribute very little.

7.4 Recommended configuration

The empirically optimal playbook is:

$$\text{Sell tiers: } {1.1,\; 1.2,\; 1.3,\; 1.4,\; 1.5,\; 2.0,\; 3.0} \times P_{\text{trend}}, \quad 30\%\text{ per tier}$$

$$\text{Floor leverage: } \lambda = 20\%\text{ of stack value when price} < 0.5 \times P_{\text{trend}}$$

$$\text{Borrow for withdrawals below } 0.8 \times P_{\text{trend}},\text{ capped at 50\% LTV}$$

$$\text{Settle debt and redeploy at } P_{\text{trend}}$$

This configuration delivers median BTC CAGR of 24.1% at historical volatility, 10.8% at 0.7× compression, and 3.2% at 0.5× compression — 0.6–1.3 pp above the P16 baseline (wide-50%) across all regimes.

---

8. Discussion

8.1 The two-source growth model

The playbook's BTC growth comes from two independent sources:

1. **Cycle premium capture:** selling BTC at euphoria-phase multiples (1.1–3.0×) and rebuying at trend (1.0×). This source dominates at historical volatility and remains active at 0.7× compression.
2. **Floor leverage:** borrowing at distressed floor prices (0.5×) and repaying at trend (1.0×). At σ=0.5×, floor leverage at λ=20% contributes approximately 1.0 pp to the 2.6 pp median CAGR — roughly 39% of the total. Its importance grows as cycle amplitude diminishes.

The two sources are structurally complementary: cycle premium capture works when price moves *up*, floor leverage works when price moves *down*. Together they exploit the mean-reverting character of the OU residual from both sides.

8.2 Model limitations and optimism

These results carry known optimism:

1. **Fixed reflecting floor barrier.** The floor prevents sustained below-floor prices, which has held empirically but is not guaranteed.
2. **Perfect execution.** Tiers fire at the exact month of threshold crossing; slippage from estimation error in trend, order execution, and liquidity is not modelled.
3. **Credit availability.** Borrowing is assumed available at 10% APR whenever the LTV constraint permits. In real floor events, credit markets may restrict or reprice lending.
4. **Beta stability.** The power-law trend is assumed constant. Observed beta drift of 0.001/yr would shift USD figures but preserves the BTC CAGR structure.

The qualitative findings — optimal sell% around 30%, floor leverage adds ~10%/episode, volatility threshold at 0.5–0.7× — are robust directionally but absolute CAGR numbers are model-maximum estimates.

8.3 Practical implementation

For an investor managing a Bitcoin retirement position, the findings suggest a hierarchy of priorities:

1. **Choose P5 as the floor anchor.** P1 is unnecessarily conservative; P10 requires the playbook to perform reliably.
2. **Use the 7-tier sell grid at 30% per tier.** Prefer 35% if the primary concern is P(≥5% CAGR) under low-volatility scenarios.
3. **Size floor leverage at 20% of stack value below 0.5× trend.** Scale toward 30% with reliable credit access and active monitoring.
4. **Treat volatility compression as the primary model risk.** Monitor adoption signals; if S-curve metrics plateau, migrate toward lower sell-tier thresholds and higher floor-leverage allocations.
5. **Do not rely on ≥10% BTC CAGR in any 20-year window.** The 80% probability is strong but not certain; roughly 1 in 5 twenty-year retirements may fall short of that benchmark even under the model's assumptions.

---

9. Conclusions

This paper extends the P16 playbook framework with four quantitative additions:

1. **P5 (0.4295×, $22 030/yr/BTC) is the practical floor anchor.** It provides 97.4% passive survival at 30 years and 100% playbook survival across all tested horizons and inflation levels up to 10%.

2. **The "10%/year BTC" benchmark is achievable with 80–100% probability** over 20–50 year horizons at historical volatility. The median playbook outcome is 20–26% CAGR, substantially above the benchmark.

3. **The playbook breaks between 0.5× and 0.7× historical volatility**, transitioning from a dominant strategy to a marginal one. Survival remains ≥99.7% regardless of volatility level, because ruin and accumulation are driven by separate mechanisms.
4. **The optimal sell-tier configuration is seven tiers (1.1–3.0× trend) at 30% per tier, plus 20% floor leverage below 0.5× trend.** This structure outperforms the P16 baseline (wide-50%) by 0.6–1.3 pp median CAGR across all tested volatility regimes, and provides residual compounding independent of whether euphoria-phase tiers fire.

Appendix A — Simulation Parameters

OU Process Parameters

Cycle	Mean μ_c	Step-sigma σ_c	Student-t d.o.f.	Recency weight
1	−0.1116	0.1499	4.71	5%
2	0.0517	0.1504	4.57	10%
3	0.0066	0.1009	5.08	25%
4	−0.0099	0.0981	5.95	60%

Cycle 1 step-sigma is ~53% higher than cycle 4 (0.150 vs 0.098). The OU process evolves as: $r_{t+1/12} = \mu_c + (r_t - \mu_c)e^{-\kappa/12} + \sigma_c\sqrt{(1-e^{-\kappa/6})/(2\kappa)}\,\epsilon_t$ with $\kappa=1.0$ and $\epsilon_t \sim t_{\nu_c}$.

Fixed Parameters

Parameter	Value
Paths per scenario	50,000
OU mean reversion κ	1.0
Floor boundary	Fixed at chosen percentile log-residual (P1/P5/P10)
Debt APR	10% (monthly compounding)
Borrow LTV cap	50%
Borrow trigger (withdrawals)	price < 0.8× trend
Floor leverage trigger	price < 0.5× trend
Floor leverage borrow cap λ	20% of stack value (Table 5 sweeps 0–50%)
Start date	2026-05-22
Starting stack	1 BTC, at trend
Inflation	6% (unless stated)

Appendix B — Code Availability

All simulation code is in the `analysis/` directory of the Bitcoin Power Law Observatory repository:

- `mc_floor_definitions_comparison.py` — floor definition calibration and baseline survival
- `mc_lifestyle_inflation.py` — inflation and lifestyle sensitivity
- `mc_playbook_floor_definitions.py` — playbook CAGR, volatility stress, floor leverage, tier optimisation
- `mc_playbook_floor_definitions.json` — full numerical results for all 248 scenarios

Results are reproducible with `python analysis/mc_playbook_floor_definitions.py` (requires numpy).