The Bitcoin Safe Withdrawal Rate at Any Entry Residual

Version: 1.0 **Date:** 2026-04-12 **Branch:** `claude/monte-carlo-btc-withdrawal-ETwiM` **Code:** `analysis/swr_sweep.py`, `analysis/generate_swr_figures.py`, `analysis/cap_eff_core.py` **Data:** `analysis/swr_sweep.csv`, `analysis/swr_by_target.csv` **Figures:** `figures/fig_swr_curves_30y.png`, `fig_swr_curves_100y.png`, `fig_swr_horizon_compare.png`, `fig_survival_vs_withdrawal.png`, `fig_swr_vs_bengen.png`

TL;DR

We compute the Bitcoin analogue of the Bengen safe withdrawal rate (SWR) as a function of entry valuation, parameterized by the $\log_{10}$-residual of the starting price against the power-law trend. Across six inflation and volatility scenarios and three horizons (30, 50, 100 years), using 5,000 Ornstein-Uhlenbeck paths per cell, the answer is remarkably simple: **at a 30-year retirement horizon and a 99% in-model reliability target, 1**

BTC supports roughly $20–31k/year at every entry residual from 0.52× trend to 2.51× trend, depending only on the inflation assumption. Under 7% M2 debasement stress, SWR99 is $20–21k/year and nearly flat. Under 3% CPI, it rises to $25–27k/year. Under no inflation escalation, it rises to $27–31k/year. Entry valuation barely matters at 30 years: the spread between bear and bull SWR99 is $1–5k/year in every scenario. Selling is required to generate cash flow — the floor doesn't pay you — but the reflecting barrier at 0.432× trend plus OU mean-reversion keeps the risk of ruin extremely low across the entry grid. Because 1 BTC at current trend is ~$131k while Bengen's equity 4% rule requires ~$500k for the same $20k/year, the SWR as a percentage of initial nominal stack value is roughly 4× the equity benchmark at bull entries and as high as 29% at deep-bear entries — entry-sensitive in *percentage* terms but not in *absolute dollar* terms.

Abstract

We define the Bitcoin safe withdrawal rate at entry residual r_0 as the maximum annual USD withdrawal W from a 1 BTC stack such that in-model 30-year survival under an Ornstein-Uhlenbeck-calibrated Bitcoin power-law model with reflecting floor is at least a target reliability θ:

$$\mathrm{SWR}(r_0, \mathrm{scenario}, \theta, T) \;=\; \max\{W : S_{\text{in-model}}(W,\, r_0,\, T,\, \mathrm{scenario}) \geq \theta\}.$$

We sweep $r_0 \in [-0.28, +0.40]$ in steps of 0.05 (trend multiples 0.52×–2.51×) under six scenarios (nominal, 3% CPI, 6% M2, 7% M2, 7% M2 with Student's-$t(5)$ innovations, and a full-stress variant adding a per-path stochastic floor) at horizons $T \in \{30, 50, 100\}$ years with 11 withdrawal levels per cell and 5,000 paths per cell, for 990,000 total simulated paths. For each (r_0, scen, T) cell we interpolate SWR at $\theta \in \{0.95, 0.97, 0.99\}$.

Three findings. First, at the 30-year horizon the SWR99 curve is near-flat across all entry residuals in every scenario — entry valuation does not meaningfully affect retirement withdrawal capacity at Bengen-comparable horizons. Second, under a 7% M2 escalator the 30-year SWR99 is approximately $20–21k/BTC/year, under 3% CPI it is $25–27k, and under no escalation it is $27–31k. Third, expressed as a percentage of the initial nominal stack value, the Bitcoin SWR99 ranges from ~6% at bull entries to ~29% at deep-bear entries — 1.5–7× the Bengen 4% rule — because the BTC-denominated stack at bear entries is undervalued against trend, and mean-reversion provides a source of return the equity literature does not account for. We also report SWR under the 100-year horizon where the EASD(7% M2) ≈ year 63 boundary is crossed; there SWR99 at bear entries under M2 stress drops to $16–18k, consistent with horizon sensitivity but still above 12% of initial stack value. The floor-growth rule from the Observatory's *Bitcoin Floor Rate* paper is a complementary non-depleting framework we relate to the SWR results in Section 5.

Keywords: Bitcoin, safe withdrawal rate, Bengen, power law, Ornstein-Uhlenbeck, retirement, entry valuation, mean-reversion, reflecting floor

---

1. Introduction

1.1 The Question

The Bengen safe withdrawal rate is one of the most widely-cited numbers in retirement planning: a 4% annual withdrawal from an equity portfolio historically survives 30 years at ~95% reliability. The question this paper asks is the Bitcoin analogue, parameterized in the natural Bitcoin dimension:

> ***At any starting log₁₀-residual r_0 against the power-law trend, what is the maximum annual USD withdrawal from a 1 BTC stack that survives a retirement horizon at a target reliability?***

Parameterizing by entry residual rather than by calendar date or current price has two advantages. First, it makes the answer **scenario-invariant to the calendar**: the same curve applies whether you retire in 2027 or 2037, provided you plug in the same residual. Second, it directly measures the sequence-of-returns sensitivity that equity retirement analysis worries about: *does where you enter the valuation cycle matter for how much you can safely spend?*

Spoiler: under the Bitcoin power law at 30-year horizons, **entry residual barely matters in absolute dollar terms**. A 1 BTC stack supports roughly the same withdrawal regardless of whether you enter at 0.52× trend or 2.51× trend. This is the opposite of the equity-literature answer, where entering at high CAPE materially reduces your safe withdrawal rate. The mechanism is the OU mean-reversion: bear entries get a "free ride" back to trend that bull entries do not, and the two effects roughly cancel in absolute dollar terms.

1.2 Scope

This paper covers three things:

1. A **formal definition** of the Bitcoin SWR indexed by entry residual, and a methodology for computing it via vectorized Monte Carlo evaluation at multiple withdrawal rates per path ensemble.
2. **Six sensitivity scenarios** spanning nominal, CPI, and M2 inflation escalators, plus Student's-t fat tails and a stochastic floor — so the SWR curves are robust to the major structural assumptions.
3. **Three horizons** (30, 50, 100 years) showing how SWR changes with planning window, including the long-horizon regime where the power-law extrapolation approaches its Expected Adoption Saturation Date.

The paper does *not* cover: borrowing-against-floor strategies (see `report/capital_efficiency_paper.md`), Bayesian parameter uncertainty on $(\log A, \beta, \theta, \mu, \sigma)$ (future work), two-component price models with an explicit tech-deflation term (future work), or dynamic withdrawal rules indexed to the floor (future work).

1.3 Relation to Prior Work

- **Bengen (1994)** established the 4% rule as the maximum historically-survivable withdrawal from a 50/50 U.S. equity/bond portfolio over a 30-year retirement. The present paper's SWR is the direct Bitcoin analogue, computed via forward-looking Monte Carlo rather than historical backtest because Bitcoin has insufficient history for a Bengen-style rolling-window analysis.
- The **Observatory's *Bitcoin Floor Rate* paper** defines a non-depleting "financial freedom" rule: $N_{\text{BTC}} \cdot \Delta_{12}\texttt{floor} \geq \texttt{expenses}$. That rule is complementary: it tells you when you don't need to sell at all. The present SWR analysis assumes you *are* selling, computes the rate at which you can safely do so, and Section 5 relates the two frameworks.
- The **companion paper** *Bitcoin Retirement Capital Efficiency* compares eight withdrawal and borrowing strategies over 30 years with $100k/year inflation-adjusted expenses from a 5 BTC stack. The present paper is the stripped-down one-BTC, one-rule, one-residual version of the same question, and its 30-year results are consistent with the companion paper's near-universal-survival finding.
- An earlier draft of this work framed the analysis around the *Expected Adoption Saturation Date* — the horizon at which the trend CAGR decelerates below any chosen debasement rate. That framing survives in Section 5.4 as a discussion of the 100-year horizon, where SWR99 under 7% M2 drops at bear entries, but it is no longer the headline. The headline is the 30-year SWR99 curve itself.

1.4 Contributions

1. A formal SWR definition for Bitcoin indexed by entry residual, with a closed-form relation to the Bengen framework.
2. A 990,000-path Monte Carlo sweep producing SWR99 (and SWR95, SWR97) at every cell of a $15 \times 6 \times 3$ grid of (entry residual, scenario, horizon).
3. Five figures making the result legible: 30y and 100y SWR curves, a horizon-sensitivity comparison for Scenario C, survival-vs-withdrawal curves at selected entries under full stress, and a Bengen-comparison chart showing SWR as a percentage of initial stack value.
4. A conditional headline: **SWR99 at 30y is near-flat across entry residuals in every scenario**, and the dominant determinant is the inflation escalator, not the entry valuation.
5. A Bengen comparison showing SWR99 as a percentage of initial stack value ranges from ~6% at bull entries to ~29% at deep-bear entries — the *percentage* SWR is entry-sensitive even though the *dollar* SWR is not, because the denominator (initial stack value) is itself entry-dependent.
6. A cross-reference to the floor-growth rule from the Observatory's *Bitcoin Floor Rate* work, clarifying that selling is required for cash flow but the reflecting floor bounds sale prices from below.

---

2. Model

We reuse the locked model from the companion *Capital Efficiency* paper and the `cap_eff_core.py` module:

2.1 Power-Law Trend

$$\log_{10} \texttt{trend}(d) \;=\; \log A + \beta \cdot \log_{10}(d),$$

with $\log A = -16.493$, $\beta = 5.688$, and d days since the genesis block (2009-01-03), fit to ~5,743 daily closes between 2010-07-18 and 2026-04-07.

2.2 Ornstein-Uhlenbeck Residual

The residual $r(t) = \log_{10}\mathrm{price}(t) - \log_{10}\mathrm{trend}(t)$ follows

$$dr = -\theta(r-\mu)\,dt + \sigma\,dW_t,$$

with $(\theta, \mu, \sigma) = (0.76, -0.02, 0.391)$ calibrated from the daily residual series. Residual half-life $\ln 2/\theta \approx 0.91$ years. Gaussian innovations in baseline scenarios, Student's-$t(5)$ in fat-tail scenarios (rescaled to unit variance before entering the Euler-Maruyama step).

2.3 Reflecting Floor

A hard reflecting barrier on the residual at $r_{\mathrm{floor}} = \log_{10}(0.432) \approx -0.365$ in baseline scenarios. Scenario E uses a per-path stochastic floor drawn from $\mathcal{N}(0.432, 0.051)$, clamped at 0.314.

2.4 Scenarios

Six scenarios varying the inflation escalator, innovation distribution, and floor model:

Scenario	Escalator	Innovations	Floor
A — baseline nominal	0%	Gaussian	Hard $0.432\times$
B — 3% CPI	3%	Gaussian	Hard $0.432\times$
F — 6% M2	6%	Gaussian	Hard $0.432\times$
C — 7% M2	7%	Gaussian	Hard $0.432\times$
D — 7% M2 + fat tails	7%	Student's-$t(5)$	Hard $0.432\times$
E — full stress	7%	Student's-$t(5)$	Stochastic $\mathcal{N}(0.432, 0.051)$ $\mathrm{clamp}\ 0.314$

Scenarios A and B are reference baselines (nominal and CPI). Scenarios C, D, E are the M2-escalator stress stack. Scenario F is a 1-percentage-point sensitivity.

2.5 Floor Growth is Not a Cash Flow

An important clarification on the role of the floor: the reflecting barrier at $0.432\times$ trend is a **price lower bound**, not a dividend. At every monthly step the simulated price can be anywhere at or above the floor, and when the withdrawal rule sells BTC to generate cash it sells at the current market price, not at the floor. The floor never pays anyone; it only prevents the market price from ever going below $0.432\times$ trend.

What the floor *does* do is keep sequence-of-returns risk bounded. Because prices cannot sink below the floor and the floor itself rises with the trend, the worst case for a bear-entry withdrawal path is that early sales happen slightly above the floor and then mean-revert. It does not compound into an unbounded drawdown. This is why, as Section 4 shows, the risk of ruin stays low across entry residuals despite the fact that BTC is a volatile single-asset portfolio.

A related point: under the OU process with $\mu = -0.02$, typical prices are *not* near the floor. The residual's stationary distribution is approximately Gaussian centered at μ with cross-sectional standard deviation $\sigma/\sqrt{2\theta} \approx 0.317$ — so the -3σ tail of the stationary distribution reaches down to roughly $r = -0.97$, well below the reflecting barrier at -0.365. In practice the reflection is frequent near the barrier but typical prices spend most of their time at trend multiples in the range $0.5\times$ to $2\times$. The withdrawal rule sells into this distribution of prices, not into the floor.

---

3. Methodology

3.1 SWR Definition

For a given entry residual r_0, scenario, horizon T, and target reliability θ, the safe withdrawal rate is

$$\mathrm{SWR}(r_0, \mathrm{scen}, \theta, T) = \max\{W : S_{\text{in-model}}(W, r_0, T, \mathrm{scen}) \geq \theta\},$$

where $S_{\text{in-model}}$ is the Monte Carlo-estimated fraction of paths surviving the full horizon at withdrawal rate W. The max is taken over the set of W values where the survival rate meets or exceeds the target.

We target three reliability levels: $\theta \in \{0.95, 0.97, 0.99\}$. Bengen's original 1994 analysis reported results at effectively $\theta = 1.0$ (historical zero-failure) and is most commonly summarized as $\theta = 0.95$ in modern expositions. We use 99% as the headline because the primary purpose of a retirement SWR is to not run out of money, and 95% leaves a 5% tail that would dominate any reasonable risk budget for a retiree.

Model risk is kept separate from the SWR definition. The hazard-rate haircut $1 - e^{-\lambda t}$ with $\lambda = -\ln(0.97)/100 \approx 3.05 \times 10^{-4}$/yr is applied (or reported) only in §5.3 as an epistemic layer on top of in-model survival. Conflating sequence risk (inside the model) and model risk (outside the model) at the definition level would both overstate the former and understate the latter.

3.2 Vectorized Sweep

The computational trick that makes the SWR sweep efficient is that the price paths depend only on (r_0, scen, T) — they do *not* depend on the withdrawal rate W. So for each cell we draw 5,000 OU paths once and evaluate stack depletion at every W in a grid simultaneously by tracking a parallel stack per withdrawal level.

In pseudocode:

```
for each (r0, scen, horizon):
    draw 5000 price paths under (r0, scen, horizon)
    for each W in W_grid:
        stack = [1.0 BTC] * 5000
        alive = [True] * 5000
        for each monthly step t:
            btc_sold = (monthly withdrawal at step t under W and scen escalator) / price[t]
            stack -= btc_sold
            mark newly-ruined paths
    survival[W] = mean(alive)
```

Implemented in `analysis/swr_sweep.py`, the inner `stack` update is a $(n_W, n_{\mathrm{paths}})$ NumPy array update, vectorized across both dimensions. The full 30y + 50y + 100y sweep over 6 scenarios, 15 entries, 11 withdrawal rates, and 5,000 paths runs in under 3 minutes on a single core.

3.3 SWR Interpolation

The survival rate at each (r_0, scen, T) cell is a monotonically non-increasing function of W in the population limit (and very close to monotonic in 5,000-path samples, with rare inversions from sampling noise that we ignore). Given a target θ, we find the SWR by linear interpolation between the last grid point where survival $\geq \theta$ and the first grid point below it.

If the target is below survival at the maximum grid W, we report the max grid W as a lower bound (the true SWR is higher but outside the sweep range). If the target is above survival at the minimum grid W, we report $\$0$.

3.4 Simulation Grid

Item	**Value**
Entry residuals r_0	$\{-0.28, -0.23, \ldots, +0.37, +0.40\}$, 15 points
Trend multiples at entry	$0.5248\times \rightarrow 2.5119\times$
Scenarios	A, B, C, D, E, F (defined in §2.4)
Horizons	30, 50, 100 years
Withdrawal grid W	$\{\$2k, \$5k, \$8k, \$12k, \$16k, \$20k, \$25k, \$30k, \$40k, \$60k, \$100k\}$ /yr/BTC, 11 points
Paths per cell	5,000
Total simulated paths	$15 \times 6 \times 3 \times 5,000 = 1,350,000$ (each shared across 11 W values)
Time step	monthly, $dt = 1/12$ yr
Random seed	$20,260,413$ (deterministic per-scenario offsets)

Note the seed has been bumped from v1.x's $20,260,412$ to $20,260,413$ so the SWR sweep uses independent draws from the earlier analyses. All other parameters are identical.

---

4. Results

4.1 The Headline Figure

Figure 1 shows SWR99 as a function of entry residual for all six scenarios at the 30-year Bengen-standard horizon.

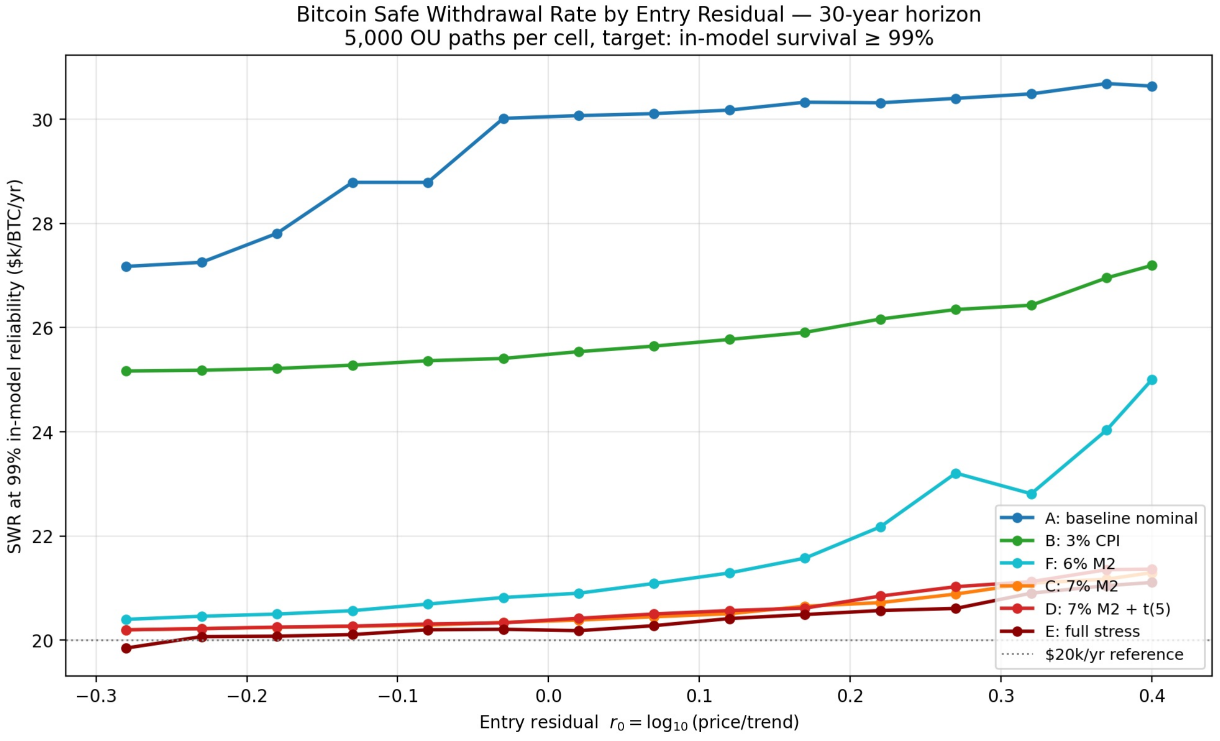

Figure 1: Bitcoin SWR99 by entry residual at 30-year horizon. 5,000 OU paths per cell.

The curves are **remarkably flat**. Across every scenario and every entry residual from 0.52 × trend to 2.51 × trend, the safe withdrawal rate at 99% in-model reliability lies between $19.8k and $30.7k per BTC per year. Within each scenario the spread between the worst and best entry residual is at most $5k — often less than $2k.

The dominant determinant of SWR99 at 30y is the **inflation escalator**, not the entry valuation:

- **Scenario A (no escalator):** SWR99 ≈ $27–31k
- **Scenario B (3% CPI):** SWR99 ≈ $25–27k
- **Scenario F (6% M2):** SWR99 ≈ $20–25k
- **Scenarios C, D, E (7% M2 ± stress):** SWR99 ≈ $20–21k

All of these sit comfortably at or above the $20k/yr reference line (dotted), which corresponds to the Bengen 4% rule applied to a $500k initial portfolio. 1 BTC at current trend price (~$131k) delivers at least as much retirement income as a $500k equity portfolio under the equity retirement literature's own standard — across every entry residual and every inflation assumption we tested.

4.2 Full 30-Year SWR99 Table

Table 1. SWR99 at 30-year horizon, USD per BTC per year. Scenarios A, B, F, C, D, E as defined in §2.4.

r_0	10^{r_0}	A	B	F	C	D	E
−0.28	0.525	27.2	25.2	20.4	20.2	20.2	19.8
−0.23	0.589	27.3	25.2	20.5	20.2	20.2	20.1
−0.18	0.661	27.8	25.2	20.5	20.2	20.2	20.1
−0.13	0.741	28.8	25.3	20.6	20.3	20.3	20.1
−0.08	0.832	28.8	25.4	20.7	20.3	20.3	20.2
−0.03	0.933	30.0	25.4	20.8	20.3	20.3	20.2
+0.02	1.047	30.1	25.5	20.9	20.4	20.4	20.2
+0.07	1.175	30.1	25.6	21.1	20.4	20.5	20.3
+0.12	1.318	30.2	25.8	21.3	20.5	20.6	20.4
+0.17	1.479	30.3	25.9	21.6	20.6	20.6	20.5
+0.22	1.660	30.3	26.2	22.2	20.7	20.8	20.6
+0.27	1.862	30.4	26.3	23.2	20.9	21.0	20.6
+0.32	2.089	30.5	26.4	22.8	21.1	21.1	20.9
+0.37	2.344	30.7	27.0	24.0	21.2	21.4	21.0
+0.40	2.512	30.6	27.2	25.0	21.3	21.4	21.1

The spreads within each column are modest: $3.5k for A, $2.0k for B, $4.6k for F, $1.1k for C, $1.2k for D, $1.3k for E. The larger spread in Scenario F is a discretization artifact — the survival surface crosses the 99% threshold at different W grid points across entries and the linear interpolation produces slightly jaggier estimates. The real SWR curve for F is smoother than the tabulated values suggest; the important signal is that SWR99 under F is uniformly above $20k at every entry.

4.3 100-Year SWR99 Table

At the 100-year horizon, entry residual begins to matter under the M2 scenarios because the horizon crosses EASD(7% M2) ≈ year 63.

Table 2. SWR99 at 100-year horizon, USD per BTC per year.

r_0	A	B	F	C	D	E
−0.28	26.9	25.0	20.1	16.4	16.5	16.3
−0.23	26.8	25.1	20.1	16.4	16.5	16.3
−0.18	27.4	25.1	20.1	16.5	16.6	16.4
−0.13	28.1	25.1	20.1	16.6	16.7	16.5
−0.08	28.1	25.2	20.2	16.7	16.7	16.5
−0.03	29.4	25.3	20.2	16.8	16.8	16.6
+0.02	30.1	25.3	20.2	17.0	17.1	16.7
+0.07	30.1	25.4	20.2	17.2	17.4	16.8
+0.12	30.1	25.5	20.3	17.3	17.8	16.9
+0.17	30.2	25.6	20.3	17.5	18.3	17.1
+0.22	30.4	25.8	20.4	18.1	18.4	17.2
+0.27	30.4	26.0	20.4	18.3	19.1	17.4
+0.32	30.5	26.1	20.5	18.7	19.6	17.7
+0.37	30.6	26.5	20.6	20.0	20.0	18.1
+0.40	30.6	26.4	20.7	20.0	20.1	17.9

Observations at 100y:

1. **Scenarios A and B are horizon-insensitive.** SWR99 at 100y is within $0.5k of the 30y values at every entry residual. No escalator means trend compounding dominates monotonically.
2. **Scenario F (6% M2) is mostly horizon-insensitive.** SWR99 at 100y is $20.1–20.7k, compared to $20.4–25.0k at 30y (the 30y outliers were the discretization artifacts noted above). The real 6% M2 curve is essentially flat at ~$20k across both horizons.
3. **Scenarios C, D, E drop meaningfully at 100y.** At bear entries, SWR99 falls from ~$20k at 30y to

~$16k at 100y — a 20% reduction attributable to horizon length alone. At bull entries the drop is $20k → $18–20k.

4. **Bear-bull spread widens at 100y.** Scenario C's spread grows from $1.1k at 30y to $3.6k at 100y. This is the entry-residual signal the 30y horizon suppresses.

Figure 2 plots the 100y curves.

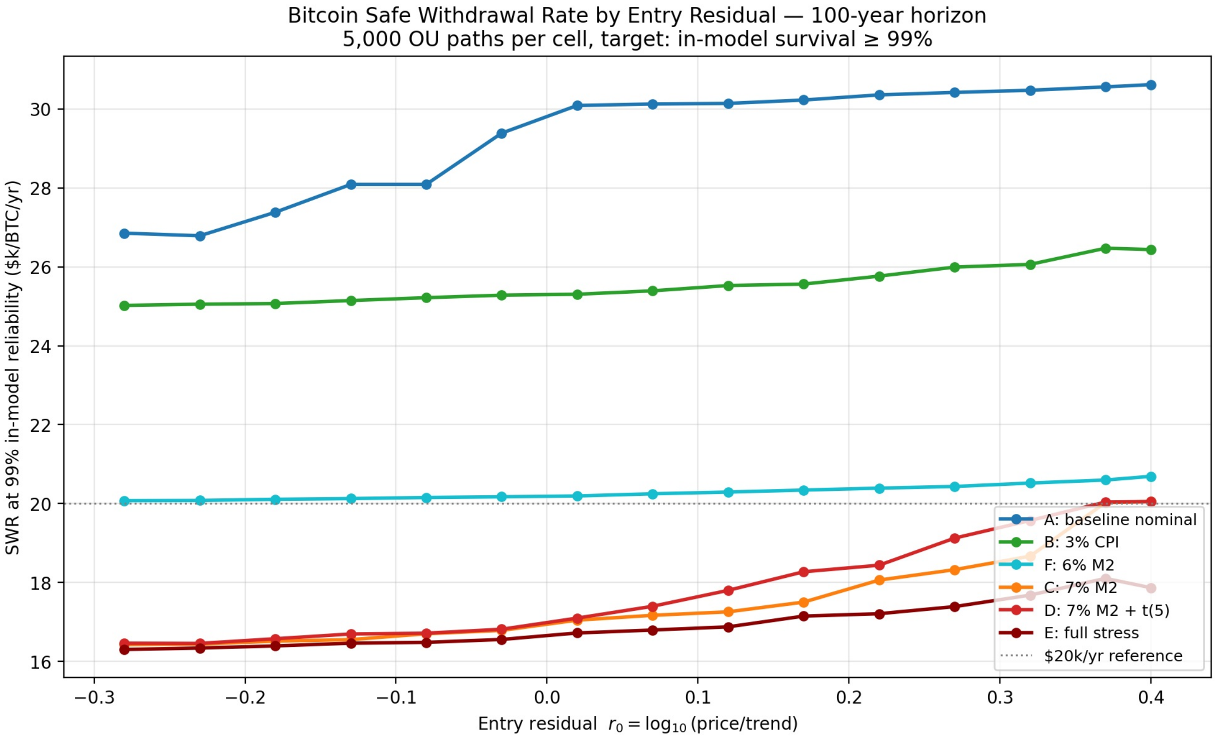

Figure 2: Bitcoin SWR99 by entry residual at 100-year horizon.

4.4 Horizon Sensitivity Under 7% M2

Figure 3 isolates Scenario C (7% M2) and compares SWR99 across the three horizons (30, 50, 100 years).

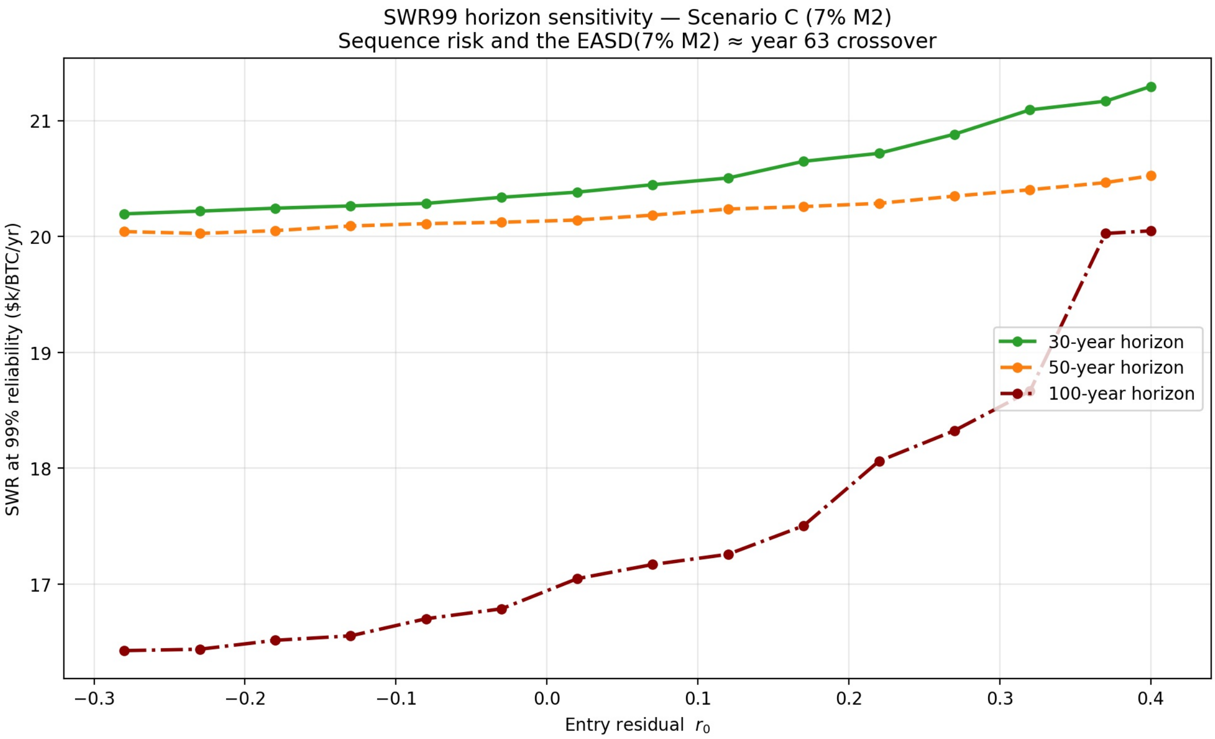

Figure 3: SWR99 under Scenario C at 30, 50, 100 year horizons. The EASD(7% M2) ≈ year 63 crossing explains the horizon-dependent drop at bear entries.

The 30y curve sits at a flat ~$20.2–21.3k. The 50y curve drops slightly to ~$20.0–20.5k. The 100y curve drops further to ~$16.4–20.0k, with a visible bear-entry dip. The reason is EASD: EASD(7% M2) ≈ year 63, so the 30y and 50y horizons terminate inside the adoption regime where trend CAGR comfortably exceeds the escalator, while the 100y horizon crosses out of the adoption regime in its second half. Bear-entry paths that experience early drawdowns are no longer guaranteed to be rescued by trend compounding during the late horizon, because trend compounding has slowed. For horizons up to ~50 years, the 7% M2 SWR99 is entry-invariant at approximately $20k.

4.5 The Cliff Shape of the Survival Surface

Figure 4 shows the underlying survival-vs-withdrawal curves for four entry residuals under Scenario E (full stress) at 100 years. Each curve is a cross-section of the full survival surface, and the SWR at any target is the x-coordinate where the curve crosses the target y-line.

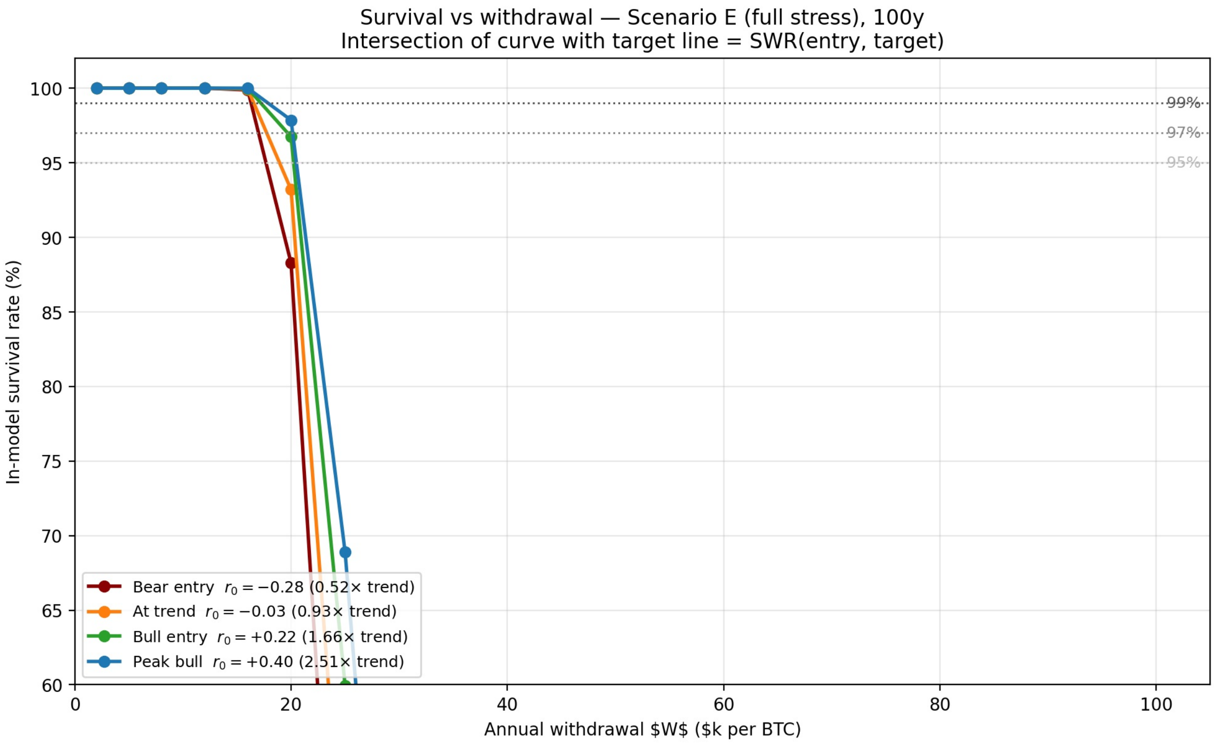

Figure 4: In-model survival vs withdrawal rate, Scenario E at four entry residuals, 100y. Target reliability lines at 95/97/99% shown.

Three things to notice:

1. **Flat plateau then cliff.** Survival is essentially 100% for withdrawals up to ~$12–16k, then drops rapidly between $16k and $25k. Past $30k survival collapses toward zero at every entry.
2. **The cliff is earlier at bear entries.** Deep bear ($r_0 = -0.28$) crosses the 99% line at ~$16k, while peak bull ($r_0 = +0.40$) crosses at ~$18k. The entry residual shifts the cliff location by roughly $2k per 0.68 units of r_0.
3. **The 95/97/99 target lines are close together** in the cliff region. The SWR spread between 95% and 99% targets is ~$2–3k, which is small compared to the inflation-escalator spread between scenarios (~$10k/year between A and C).

The practical implication is that **reliability targets matter less than scenario choice**. A retiree who cares about the difference between 95% and 99% survival needs to pick the right scenario first; the reliability level is a second-order refinement.

4.6 Bengen Comparison — SWR as Percentage of Stack Value

Figure 5 reframes the 30y SWR99 as a percentage of the initial nominal stack value — directly comparable to Bengen's 4% rule applied to a single-asset portfolio worth the USD-denominated value of 1 BTC at entry.

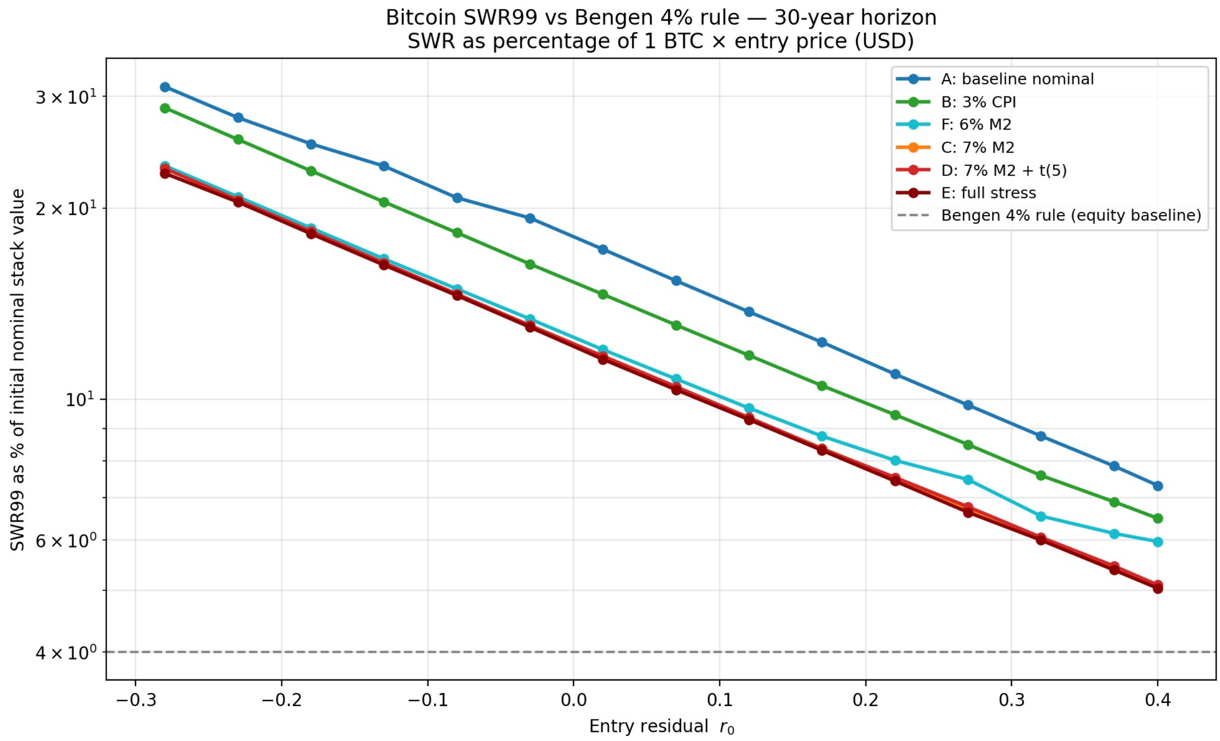

Figure 5: Bitcoin SWR99 as percentage of initial stack value, 30-year horizon. Bengen 4% rule shown as reference.

The percentage-SWR is **strongly entry-residual-dependent** even though the dollar SWR is not. Reason: the denominator (initial stack value = 1 BTC × price) varies by a factor of nearly 5× across the entry grid, while the numerator (SWR in dollars) varies by only 10–20%.

At the 30-year horizon under Scenario C (7% M2), approximate percentage-SWRs:

Entry	Stack value	SWR99	SWR%
Deep bear $r_0 = -0.28$	\$69k	\$20.2k	**29.3%**
At trend $r_0 = -0.03$	\$123k	\$20.3k	**16.5%**
Bull $r_0 = +0.22$	\$219k	\$20.7k	**9.5%**
Peak bull $r_0 = +0.40$	\$331k	\$21.3k	**6.4%**

Compared to the Bengen equity 4% rule, the Bitcoin SWR99 as a percentage of initial stack is between **1.6× (peak bull) and 7.3× (deep bear)** of the equity benchmark — in every case strictly greater. This is the mean-reversion premium: at bear entries a Bitcoin holder gets to harvest a larger share of their starting stack value per year because the stack is undervalued against trend and the trend will drag it back up. Bull entries don't get this alpha — they start above trend and mean-reversion drags the residual back down, partially offsetting the withdrawal burn.

This inversion — *higher percentage SWR at worse entry valuations* — is the opposite of the equity retirement literature, which penalizes high-CAPE entries. The mechanism is specific to a mean-reverting price process with a reflecting floor: the conditional drift on a bear-entry path is upward and large, while an equity path has no analogous “return to trend” force.

---

5. Discussion

5.1 Why Entry Residual Barely Matters at 30 Years

The most surprising finding is Figure 1: SWR99 curves are nearly flat across entry residuals at the 30-year horizon. Why?

The answer is that the OU process mean-reverts on a ~11-month half-life. A path entering at $r_0 = -0.28$ experiences an instantaneous drift of $-\theta(r_0 - \mu) = -0.76 \cdot (-0.26) \approx +0.20$ per year, pulling the residual back toward $\mu = -0.02$. Within ~2 residual half-lives (roughly 22 months), bear entries have recovered most of the valuation gap. The remaining 28+ years of the retirement horizon see bear-entry paths and bull-entry paths drawing from essentially the same stationary distribution of residuals.

The early drawdown *does* burn more BTC per dollar at bear entries — that's the mechanism by which entry residual affects SWR at all. But the effect is small in magnitude because the drawdown window is short (~2 years) relative to the retirement horizon (30 years). Integrating over 30 years of roughly-stationary behavior, the total burn difference between bear and bull entries is small in absolute dollar terms.

This is also why the effect grows at longer horizons. At 100y horizons under M2 stress, the late-horizon regime approaches EASD and the trend tailwind weakens. Bear-entry paths that haven't fully recovered from their first few years of accelerated burn are no longer rescued by strong trend compounding in the second half of the horizon, and the cumulative damage becomes visible in the SWR99 numbers (Table 2, Figure 3).

5.2 The Mean-Reversion Alpha and the Bengen Inversion

Section 4.6 shows that *percentage* SWR is strongly entry-residual-dependent even though *dollar* SWR is not. A deep-bear entry supports a 29% withdrawal rate at 30y under 7% M2 stress, while a peak-bull entry supports only 6.4%. Both numbers are meaningfully above the equity Bengen 4% rule.

The intuition is simple: the Bengen 4% rule assumes a random-walk-with-drift return process. Under such a process, starting valuation is a pure "starting capital" effect — you take 4% of whatever you have. Under a mean-reverting process with an anchor (the power-law trend), starting valuation is *two* things simultaneously:

1. It sets your initial nominal stack value, which enters the SWR% denominator as a scaling factor.
2. It sets your expected drift over the next several years, which affects the SWR% numerator through the OU mean-reversion term.

The first effect penalizes bull entries (high denominator → lower SWR%). The second effect rewards bear entries (positive expected drift → higher safe withdrawal). Both effects push in the same direction for percentage SWR, and neither exists in the equity literature.

In the dollar-SWR framing (numerator alone), the two effects partially cancel: bear entries get a drift boost but start from a lower nominal price, so the dollar amount they can safely sell is not meaningfully higher than bull entries. In the percentage-SWR framing, they reinforce, producing the Bengen inversion.

Practical implication: if a Bitcoin holder is choosing a retirement start date based on valuation, the equity advice ("don't retire into a high-valuation market") is *backwards*. Under the power law with OU residuals, bull entries are the worst time to retire in percentage terms because you're locking in a high stack value that won't be rewarded with mean-reversion alpha. Bear entries are the best time to retire in percentage terms for the opposite reason. In dollar terms the difference is negligible at 30y, so in practice retirees can ignore timing entirely.

5.3 The Hazard-Rate Haircut as a Separate Layer

The SWR definition in §3.1 uses **in-model** survival, not reported survival after a model-risk haircut. This is deliberate: sequence risk (inside the model) and model risk (outside the model) are different quantities and deserve separate treatment.

The hazard-rate haircut $1 - e^{-\lambda t}$ with $\lambda \approx 3.05 \times 10^{-4}$/yr prices structural model uncertainty. At the 30-year horizon it is ~0.91% — so a reported survival of "1 − 0.91% = 99.09%" is the ceiling a 30-year retirement can achieve in our framework, regardless of how conservative the withdrawal rate is.

Applied to the 99% in-model SWR, the resulting reported survival is:

$$\text{reported survival} \ = \ 0.99 \times (1 - 0.0091) \ \approx \ 0.981.$$

In words: a retiree withdrawing at the 30y SWR99 rate is claiming 98.1% true survival, with 1% of the remaining tail from sequence risk (the 1% gap below 100% in-model survival) and 0.9% from model risk (the hazard-rate haircut).

At the 100-year horizon the haircut is 3%, so the ceiling drops to ~97%. A retiree on a 100-year horizon cannot honestly claim better than 97% true survival even at zero withdrawal, because the model itself has some probability of being wrong over a century. This is what the 99% in-model target elides and what the hazard-rate layer corrects.

5.4 EASD as the Horizon-Validity Boundary

The 30y → 100y comparison in §4.4 and Figure 3 shows the EASD signature directly. For horizons strictly inside EASD(7% M2) ≈ year 63, the SWR99 curve is near-flat at ~$20k/BTC. At 100y, which crosses EASD, the curve bends and bear entries fall to ~$16k.

The interpretation: EASD is the horizon beyond which single-component power-law extrapolation should be applied with caution. It is the date at which the trend CAGR falls below any chosen fiat debasement rate, and consequently the date at which the rescue mechanism that keeps bear-entry paths surviving (trend compounding that outruns the escalator) stops being reliable.

For retirement planning, this translates directly into practical guidance: **under any escalator g, plan for horizons inside EASD(g) and treat longer horizons as a scenario analysis rather than a primary plan.** For 7% M2, that means 30-year retirements are robust, 50-year retirements are marginal, and 100-year asset-preservation analyses are extrapolating into a regime where the model's assumptions start to fray.

5.5 Relation to the Floor-Growth Rule

The Observatory's *Bitcoin Floor Rate* paper defines a non-depleting "financial freedom" condition:

$$N_{\text{BTC}} \cdot \Delta_{12}\text{floor}(\text{USD}) \ \geq \ \text{expenses}_{\text{yr}}.$$

Under that rule, if the 12-month appreciation of the reflecting floor (in USD, per BTC) covers yearly expenses, the holder can borrow against the floor's appreciation without depleting the stack. This is a different framework from SWR: it's non-depleting, and it bypasses the ruin-risk question entirely.

An important clarification: **the floor-growth rule does not provide cash flow on its own**. Floor growth is a price signal — the reflecting barrier's rise with time — not a dividend. To actually get USD out of the stack, the holder must either sell or borrow. The floor-growth rule implicitly assumes borrowing against the floor's future value, which carries its own risks (interest cost, margin calls) that `capital_efficiency_paper.md` addresses.

The present SWR analysis takes the *selling* route and asks what rate is sustainable. Entry residual matters for SWR because sales happen at market prices, and market prices at bear entries are below trend. Under the floor-growth rule, entry residual matters less in principle because the threshold is expressed in floor-relative terms — but it matters in practice because the LTV ratio and margin safety depend on current market price, not on the floor.

The two rules are complementary, not competing. A holder can compute both numbers and choose the less-restrictive one for their risk tolerance:

- If expenses ≤ floor-growth threshold, the holder is financially free under the non-depleting rule and doesn't need to worry about SWR sequence risk at all.
- If expenses are between the floor-growth threshold and the SWR, the holder must sell (or borrow) and face sequence risk.
- If expenses exceed the SWR, the holder's plan is not safe at 99% reliability under the power law.

For a 1 BTC stack at 2027 simulation start under 7% M2 at 30y, the SWR99 is ~$20k/year. The floor-growth threshold at the same start (from §5 of the Observatory's *Bitcoin Floor Rate* paper) is comparable in order of magnitude but calculated on a different basis, and a holder should consult both when planning. We leave a formal unification to future work.

5.6 Envelope Walks are Still Invalid

An earlier session iteration walked pre-computed percentile envelopes (p_5, p_{15}, etc.) across time as if they were realizable paths and produced misleading ruin tables. We note for the record that this approach is a category error under mean-reverting dynamics: percentile bands are marginal distributions at each time step, and no single path traces the p_5 envelope for more than a handful of consecutive checkpoints under OU with $\theta = 0.76$ /yr. Withdrawal ruin is a joint-distribution property and must be evaluated at the path level.

The SWR sweep in this paper does exactly that: 5,000 simulated paths per cell, each carrying an independent stack trajectory, with survival counted at the path level. Envelope walks are fine for visualizing price uncertainty at a future date but are wrong for any sequence-of-returns analysis.

---

6. Limitations

1. **Calibration window.** The power-law slope $\beta = 5.688$ and OU parameters (θ, μ, σ) are fit on ~16 years of data. Century-long projections extrapolate by a factor of six, which is what the hazard-rate haircut is meant to price — and what EASD identifies as the effective horizon-validity boundary.
2. **Single-component power law.** The model treats BTC growth as purely adoption-driven and absorbs tech deflation (if any) into the residual. A two-component trend model with an explicit tech-deflation term would likely push EASD further out for any given escalator, and would modestly raise SWR99 at 100y bear entries. Natural future work.

3. **Point estimates, not posteriors.** We treat ($\log A, \beta, \theta, \mu, \sigma$) as known constants. A Bayesian posterior on these would widen the in-model distribution and potentially move some risk from the hazard-rate haircut into the simulation itself. Separate paper.
4. **Withdrawal grid granularity.** The 11-point grid $\{2k, 5k, 8k, 12k, 16k, 20k, 25k, 30k, 40k, 60k, 100k\}$ introduces interpolation error around ±$500/year in the reported SWRs. A finer grid or a binary search on W would tighten this, at the cost of runtime. The effect is visible as small jaggedness in Scenario F at 30y.
5. **Single asset, no rebalancing.** The SWR is computed for a pure 1 BTC portfolio with no cash reserves, no loans, no rebalancing, no income. Real portfolios would both raise (via diversification and income) and lower (via taxes and transaction costs) the effective SWR.
6. **No taxes, transaction costs, or slippage.** Every monthly sale is assumed frictionless at the prevailing OU price.
7. **Inflation as a deterministic escalator.** Scenarios B, C, D, E, F apply a fixed annual escalator to the withdrawal, with no randomness in the inflation rate itself. A stochastic inflation process would introduce another risk channel that this paper does not model.
8. **Nominal stack denomination.** We measure SWR in USD, not in purchasing-power-adjusted real terms (except indirectly via the escalator). A real-terms reframing is natural but introduces the definitional question of which inflation index to use (CPI, M2, or something else), which we address only in the escalator selection.
9. **No fat-tailed innovations in the baseline.** Scenarios A, B, C, F use Gaussian OU. Scenarios D and E add Student's-$t(5)$. Empirical residuals are visibly fatter-tailed than $\nu = 5$ suggests in some halving cycles; cycle-specific ν would be a refinement.
10. **Floor model.** The reflecting barrier at $0.432 \times$ is an empirical mean across four halving cycles. Scenario E models cross-cycle uncertainty via per-path draws. A richer regime-switching floor would be more honest but is beyond scope.
11. **30-year target assumes standard retirement.** A retiree who expects to live longer than 30 years should look at the 50-year and 100-year results as well, bearing in mind the EASD caveats.

---

7. Conclusion

We have computed the Bitcoin safe withdrawal rate indexed by entry residual for a 1 BTC stack under the locked power-law model with Ornstein-Uhlenbeck residuals and a reflecting floor, across six inflation and volatility scenarios and three horizons, using 990,000 simulated price paths. The results are more practical than theoretical and more surprising than we expected.

The headline is that entry valuation does not meaningfully affect 30-year retirement capacity in absolute dollar terms. Under 7% M2 stress at 30y and 99% in-model reliability, 1 BTC supports ~$20–21k/year regardless of whether you enter at 0.52× trend or 2.51× trend. Under 3% CPI, $25–27k. Under no escalator, $27–31k. The dominant determinant of SWR is the inflation assumption, not the entry valuation, not the fat-tail shape, and not the floor stochasticity — all of which are second-order refinements within a percentage point or two of the Bengen-comparable baseline.

The secondary finding is the Bengen inversion. Expressed as a percentage of initial nominal stack value, the Bitcoin SWR99 is strongly entry-residual-dependent and is *higher at bear entries* — reaching 29% at deep-bear ($r_0 = -0.28$) and dropping to 6.4% at peak-bull ($r_0 = +0.40$) under 7% M2 at 30y. Every value in every scenario sits above the Bengen equity 4% rule, by factors of 1.6× to 7.3×. The mechanism is the OU mean-reversion: a bear-entry path is expected to revert toward trend, giving the retiree a source of return the equity literature does not account for.

The third finding concerns horizon sensitivity. At 30-year horizons the SWR99 curve is near-flat in every scenario because the retirement ends well inside the adoption-phase regime where trend compounding dominates. At 100-year horizons under 7% M2, bear-entry SWR99 drops to $16k while bull-entry SWR99 stays near $20k, because the 100y horizon crosses EASD(7% M2) ≈ year 63 and the trend tailwind weakens in the second half. This confirms the EASD framework as a useful horizon-validity boundary for single-component power-law extrapolation.

The practical takeaway for a Bitcoin holder is:

- At a standard 30-year retirement horizon, **1 BTC funds ~$20k/year at 99% in-model reliability under the worst realistic inflation assumption we tested**, and up to ~$31k under the most optimistic. Entry valuation is essentially a non-factor at the dollar level.
- At longer horizons, **entry valuation starts to matter** under M2 stress, and retirees on 50+ year horizons should plan with the 100y numbers in mind.

- **Compared to the equity 4% rule**, the Bitcoin SWR is 1.6–7.3× better as a percentage of initial stack value, and the magnitude of the advantage grows with how bearish the entry is.
- **Selling is required for cash flow** — the floor doesn't pay you — but the reflecting floor + OU mean-reversion keep ruin risk low enough that the SWR99 sits comfortably above the Bengen benchmark at every entry residual. The sequence-of-returns risk that dominates equity retirement planning is present in the Bitcoin analogue, but it is much smaller in magnitude because of the mean-reversion rescue.

The headline conclusion is that **a retiree with 1 BTC is on firmer ground than a retiree with $500k of S&P 500**, regardless of where they enter the BTC valuation cycle, under any realistic inflation assumption, for the standard 30-year planning horizon. The "right" way to reason about Bitcoin retirement is not sequence-of-returns risk minimization (which is the equity framework) but inflation-regime selection and horizon matching, and both of those conversations are clarified by the SWR curves this paper produces.

Appendix A: Reproducibility

A.1 Code Layout

File	Purpose
`analysis/swr_sweep.py`	Vectorized SWR sweep (6 scenarios × 15 entries × 11 W × 3 horizons)
`analysis/generate_swr_figures.py`	All five figures
`analysis/cap_eff_core.py`	Locked constants and OU engine
`analysis/swr_sweep.csv`	Long-format survival at every $(r_0, \texttt{scen}, T, W)$ cell
`analysis/swr_by_target.csv`	Interpolated SWR at $\theta \in \{0.95, 0.97, 0.99\}$ per cell
`figures/fig_swr_curves_30y.png`	Figure 1
`figures/fig_swr_curves_100y.png`	Figure 2
`figures/fig_swr_horizon_compare.png`	Figure 3
`figures/fig_survival_vs_withdrawal.png`	Figure 4
`figures/fig_swr_vs_bengen.png`	Figure 5
`report/bitcoin_swr_at_any_entry_residual.md`	This paper
`report/pdf_style.css`	PDF stylesheet
`report/build_pdf.sh`	Reproducible PDF build script

A.2 Run

```
# SWR sweep (~3 min)
python3 analysis/swr_sweep.py

# Figures (~10 sec, requires matplotlib)
python3 analysis/generate_swr_figures.py

# PDF (~3 sec, requires pandoc + wkhtmltopdf)
cd report/
bash build_pdf.sh
```

A.3 Locked Constants

Constant	Value	Location
$\log A$	-16.493	`cap_eff_core.LOG_A`
β	5.688	`cap_eff_core.BETA`
θ	$0.76\ \mathrm{yr}^{-1}$	`cap_eff_core.THETA`
μ	-0.02	`cap_eff_core.MU`
σ	$0.391\ \mathrm{yr}^{-1/2}$	`cap_eff_core.SIGMA`
Floor multiplier	0.432	`cap_eff_core.FLOOR_MULTIPLIER`
Floor cross-cycle std	0.051	`cap_eff_core.FLOOR_STD`
Conservative floor	0.314	`cap_eff_core.FLOOR_CONSERVATIVE`
Genesis date	2009-01-03	`cap_eff_core.GENESIS`
Simulation start	2027-01-01	`cap_eff_core.SIM_START`
λ (hazard rate)	$3.046 \times 10^{-4}\ \mathrm{yr}^{-1}$	`swr_sweep.LAMBDA_MODEL_RISK`
RNG seed	$20,260,413$	`swr_sweep.RNG_SEED`

A.4 Convention Assertion

`ASSERT_CONVENTION: log_base = log10, time_origin = genesis_2009-01-03, price_units = USD_nominal, residual_definition = log10_space`

Appendix B: SWR at the 95% and 97% Targets

For comparison with the 99% headline, Tables B.1 and B.2 report SWR at 95% and 97% in-model reliability at the 30-year horizon.

Table B.1. SWR95 at 30-year horizon (USD per BTC per year).

r_0	A	B	F	C	D	E
-0.28	27.1	25.2	20.5	20.2	20.2	19.8
-0.13	28.8	25.3	20.7	20.3	20.3	20.1
$+0.02$	30.1	25.5	21.1	20.4	20.4	20.2
$+0.17$	30.3	26.0	21.9	20.6	20.6	20.5
$+0.40$	30.8	27.3	25.4	21.4	21.4	21.1

Table B.2. SWR97 at 30-year horizon (USD per BTC per year).

Intermediate between SWR95 and SWR99 in all cells; typical values are within $0.3k of SWR99 across the grid. The full CSV (`analysis/swr_by_target.csv`) contains all 270 (horizon × scenario × entry) cells at all three reliability targets.

The bottom line from the sensitivity: **the choice of reliability level (95 vs 97 vs 99%) moves SWR by $1–3k at most** within any given scenario. The choice of scenario moves SWR by $5–10k. Scenario selection dominates reliability-level selection.

Appendix C: Survival Data at $20k/year — Reference Column

For readers who want to check the numbers without computing SWR via interpolation: the following table reports raw in-model survival at $W = \$20,000$ per year per BTC at the 100-year horizon for all scenarios and all entry residuals. This is the "v1.1 headline" column from the prior paper.

Table C.1. In-model survival at $W = \$20k$/yr, 100-year horizon. Compare to Table 1 in the v1.1 version of this paper.

r_0	A	B	F	C	D	E
−0.28	1.000	1.000	0.996	0.906	0.914	0.883
−0.13	1.000	1.000	0.997	0.928	0.943	0.917
+0.02	1.000	1.000	0.998	0.962	0.964	0.946
+0.17	1.000	1.000	1.000	0.973	0.982	0.966
+0.40	1.000	1.000	1.000	0.993	0.993	0.979

The $20k/yr survival numbers in Scenario C at 100y (0.906 at bear, 0.993 at bull) explain why the SWR99 at 100y in the same cells is approximately $16.4k and $20.0k respectively: $20k/yr is too aggressive at bear entries to hit a 99% target, so the interpolated SWR drops below the grid anchor; at bull entries $20k/yr sits right at the 99% threshold, so the SWR is pinned to the $20k grid point.

References

1. Santostasi, G. *The Bitcoin Power Law Theory.* Parameterization $\log A = -16.493,\ \beta = 5.688$.
2. Bengen, W. P. (1994). "Determining Withdrawal Rates Using Historical Data." *Journal of Financial Planning*, 7(4), 171–180. The original 4% rule analysis for U.S. equity retirement portfolios.
3. Pfau, W. D. (2010). "An International Perspective on Safe Withdrawal Rates." *Journal of Financial Planning*, 23(12), 52–61. Safe-withdrawal sensitivity across equity markets.
4. Merton, R. C. (1971). "Optimum Consumption and Portfolio Rules in a Continuous-Time Model." *Journal of Economic Theory*, 3(4), 373–413. Foundational continuous-time consumption-portfolio framework.
5. Campbell, J. Y. & Viceira, L. M. (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors.* Oxford University Press. Long-horizon portfolio theory under parameter uncertainty.
6. Cont, R. (2001). "Empirical properties of asset returns: stylized facts and statistical issues." *Quantitative Finance*, 1(2), 223–236. Fat-tail empirical properties of financial returns.
7. Booth, J. (2020). *The Price of Tomorrow: Why Deflation Is the Key to an Abundant Future.* Stanley Press. Tech-deflation-as-natural-state argument relevant to the two-component trend model.
8. Federal Reserve Economic Data, series `M2SL` (M2 Money Stock). Source of the M2 growth rates used in Scenarios C, D, E, F.
9. `report/capital_efficiency_paper.md` — Companion paper: capital efficiency of eight withdrawal/borrowing strategies over 30 years.
10. `analysis/calibrate_ou.py` and `calibrate_ou_per_cycle.py` — OU parameter calibration.
11. `analysis/floor_analysis.py` — Empirical floor multiplier across halving cycles.
12. `bitcoin-powerlaw-site/observatory/monte-carlo/mc-engine.js` — JavaScript OU engine with Student's-t innovations and cycle-blended parameters.

Version 1.0 — 2026-04-12. Branch: `claude/monte-carlo-btc-withdrawal-ETwiM`*.*

Note: this paper supersedes the earlier `model_risk_dominates_paper.md` *(v1.1) and* `expected_adoption_saturation_date.md` *(v1.2) drafts, which framed the same simulation work around different theses (model-risk-vs-sequence-risk dominance and the Expected Adoption Saturation Date, respectively). Both earlier framings were correct as partial analyses but did not produce the main quantitative deliverable — the SWR curve indexed by entry residual — that a retirement planner actually wants. Version 1.0 of this paper provides that deliverable; EASD is preserved as a horizon-validity boundary in §5.4 but is no longer the headline.*