Durable Leverage

Volatility Segmentation for Buy-and-Hold Leverage

Fragments Research

February 2026

Abstract

Durable leverage remains underdeveloped relative to the scale of long-duration capital in global markets. Moderate leverage (1–2×) can enhance long-run compound growth for assets with positive drift, yet today's leverage stack forces a tradeoff between tactical, timing-dependent leverage and structurally expensive set-and-hold leverage.

Margin and perpetual futures both require ongoing margin management and expose holders to liquidation, making them operationally unsuitable for strategic buy-and-hold capital allocation. Leveraged exchange-traded funds (LETFs) come closest to the desired user experience—automated, passively rebalanced exposure—but finance leverage through external hedging counterparties and therefore embed financing costs that compound against the holder over time.

We call leverage durable when an asset's long-run drift dominates its drag: the sum of financing costs and geometric compounding penalties. Existing automated leverage products often fail this test for two structural reasons: (1) leverage is intermediated through external hedging counterparties whose pricing ultimately reflects short-side capital costs, balance-sheet usage, and liquidity/volatility risk—embedding high financing drag; and (2) industry-standard multiples (>2×) push quadratic volatility drag past the point where drift can compensate.

We introduce volatility segmentation, an on-chain protocol that partitions asset volatility into a protected senior tranche and a leveraged junior tranche. This capital structure compresses the financing premium by sourcing leverage from yield-seeking depositors rather than counterparties who are hedging-constrained. The combination of a structural financing advantage and moderate leverage makes the structure compatible with buy-and-hold capital—aligning leveraged exposure with long-duration directional conviction rather than short-term market timing.

We demonstrate the approach on Bitcoin with a 1.33× leverage ratio. **Over 2015–2025, it outperforms 1× spot; by contrast, 2× underperforms and 3× collapses to near zero**. The architecture is extensible to any sufficiently liquid tokenized asset with a reliable price feed, and is strategically compelling where the underlying exhibits positive long-run drift.

---

1. Introduction

Leverage is easy to obtain but hard to keep. A position that doubles your exposure for a day is trivial to construct; one that compounds in your favor over a decade is not. We call leverage durable when, over long horizons, the incremental return from leverage—relative to holding the underlying unlevered—exceeds its costs: volatility drag and financing drag combined, formalized in Section 2 as the Durability Condition.

One might treat 1× exposure—one dollar of exposure per dollar of capital—as the natural baseline. But there is nothing inherently optimal about it. In fact, half a century of portfolio theory [1, 2, 3] implies that for assets with positive long-run returns, the growth-maximizing allocation is often greater than 1×: roughly 1.5–2× for Bitcoin, Gold, and the S&P 500 at current volatilities. Ayres and Nalebuff [4] argue that young investors should lever equity exposure precisely because time diversification favors it; Gayed and Bilello [5] show that systematic leverage on equities can magnify long-run returns when risk is managed. Yet no existing instrument delivers leveraged exposure you can hold passively for years.

A leveraged position left alone drifts from its target: up moves dilute leverage (reducing upside participation), while down moves concentrate it (increasing risk). Maintaining constant leverage therefore requires periodic rebalancing. Two forces then compound against long-horizon leveraged returns:

- **Volatility drag**—the geometric compounding penalty. Like any volatile asset, a rebalanced leverage strategy is subject to compounding drag; at leverage L, this penalty scales approximately with L^2, making high multiples especially punitive.
- **Financing drag**—the cost of borrowed capital, proportional to the amount leveraged.

Together, these forces compound against the position, and too much leverage is far worse than too little [6].

Existing automated leverage products sit on the wrong side of both forces. LETFs (∼$250B AUM [7]) embed swap financing and implementation costs directly into realized returns; these costs do not appear in the stated expense ratio and, in some cases, can exceed it by an order of magnitude—showing up instead as systematic underperformance versus the target multiple. Perpetual futures also embed high financing drag—they have carried funding rates averaging 10–15% annualized for Bitcoin [8]—and maintaining a target leverage ratio requires ongoing margin management under liquidation constraints. Margin borrowing can sometimes be cheap in calm markets, but it is liquidation-constrained and therefore operationally unsuitable for buy-and-hold allocation; rates can range from roughly 4–20% depending on venue and conditions. Each product works for short-term trading; none is built for leverage held passively for years.

At a high level, durable leverage comes down to two questions: *how volatile is the underlying*, and *how expensive is the leverage*. Market conditions set the underlying's drift (μ), volatility (σ), and the risk-free rate (r_f). Product design sets the leverage target (L) and the financing premium (s). Volatility determines how quickly compounding drag accelerates as leverage increases; the financing premium captures who supplies the leverage and what they require to do so. Leverage is durable when the additional drift earned by levering up exceeds the additional drag it introduces (equivalently, when $R(L) > R(1)$). Achieving durability therefore requires both disciplined leverage targeting and a capital structure that compresses the financing premium.[1]

2. The Anatomy of Leverage Drag

Long-term leveraged returns are eroded by two distinct forces that are often conflated.

2.1 Volatility Drag

Volatility drag is mathematical: the geometric compounding penalty that arises whenever a position is periodically rebalanced to a target leverage ratio. For a daily-rebalanced position with leverage L on an underlying asset with annualized arithmetic mean return μ and annualized standard deviation σ, the annualized growth rate R is approximated by:[2]

$$R \approx L\mu - \tfrac{1}{2} L^2 \sigma^2 \tag{1}$$

The second term, $\frac{1}{2}L^2\sigma^2$, is the volatility drag. It scales *quadratically* with leverage: doubling L quadruples the drag. This relationship is intrinsic to geometric compounding and cannot be eliminated by design.

For Bitcoin ($\sigma \approx 60\%$), a 1× position already loses 18% annually to volatility drag; at 2× the drag reaches 72%; at 3×, 162%—exceeding any plausible drift.[3] Lower-volatility assets such as the S&P 500 ($\sigma \approx 19\%$) can sustain much higher leverage before the penalty becomes prohibitive.

2.2 Financing Drag

Financing drag is institutional: the cost of borrowed capital used to produce leveraged exposure, including any premium required by the leverage provider to warehouse hedging, liquidity, and residual risk. It has two components:

- **Amount borrowed (set by leverage ratio L):** a position with leverage L borrows $(L-1)$ dollars per dollar of equity.
- **Cost per dollar borrowed:** the all-in financing rate r, which varies by instrument and market structure.

[1] Here $R(L)$ denotes the long-run compound growth rate of a periodically rebalanced leverage-L position (net of financing and fees), and $R(1)$ the corresponding 1× growth rate.

[2] Avellaneda and Zhang [9] derive this expression in continuous time; Cheng and Madhavan [10] give the discrete-time analogue.

[3] We use $L = 1.33$ as a running example throughout; it lies near the growth-rate optimum for Bitcoin-class volatility (Section 2.3), and Section 3 derives it from the protocol's capital structure.

Total annual financing drag is:

$$\text{Financing drag} \;=\; (L-1)\times r \tag{2}$$

where:

- $r = r_f + s$
- r_f is the risk-free rate
- s is the financing premium above risk-free.

We treat s as a *reduced-form, all-in premium* that subsumes the compensation required for hedging frictions, balance-sheet usage, liquidity and market-impact costs, and any residual risk that cannot be perfectly hedged. In practice, s is determined by market structure and can vary over time, but its average level is the key design variable.

Unlike volatility drag, financing drag is reducible. Lower leverage shrinks the $(L-1)$ term, and redesigning the capital source can compress s. Because volatility drag is intrinsic and cannot be designed away, s is the primary designable cost parameter at a given leverage ratio—making it the main target for protocol optimization.

2.3 The Durability Condition

Combining Equations (1) and (2), the compound growth rate of a leveraged position with financing cost r is:

$$R(L) \;\approx\; L\,\mu \;-\; \tfrac{1}{2}\,L^2\,\sigma^2 \;-\; (L-1)\,r, \tag{3}$$

where $r = r_f + s$.

The first term—drift—grows linearly with leverage. The second and third terms—volatility drag and financing drag—grow quadratically and linearly, respectively. At moderate leverage, drift can dominate. But as L increases, quadratic drag accelerates and eventually overtakes drift. Beyond that crossing point, additional leverage destroys value.

At representative Bitcoin parameters ($\mu = 50\%$, $\sigma = 60\%$),[4] leverage above roughly 1.5–1.6× already underperforms holding the underlying unlevered, and the threshold drops as financing costs rise.

Excess compound growth. Define the *excess compound growth* of leverage L relative to a 1× position as:

$$\Delta R(L) \;:=\; R(L) - R(1). \tag{4}$$

[4]These are round figures used for analytical illustration. Empirical values over 2015–2025 are $\mu \approx 71\%$ and $\sigma \approx 66\%$; the supplementary material uses the empirical values throughout.

Leverage is *durable* when it adds value versus the 1× baseline, i.e. when $\Delta R(L) > 0$.

The growth-maximizing leverage is the continuous-time Kelly fraction [1, 2, 11]:

$$L^* = \frac{\mu - r}{\sigma^2}. \tag{5}$$

Durability Condition. A leveraged position is durable if and only if it adds positive excess compound growth:

$$\Delta R(L) = R(L) - R(1) > 0.$$

Equivalently,

$$L\,\mu - \tfrac{1}{2}L^2\sigma^2 - (L-1)(r_f + s) > \mu - \tfrac{1}{2}\sigma^2.$$

Intuitively, the incremental drift gained from leverage must exceed the incremental drag it introduces.

The condition takes drift μ as an input. When drift is high, the productive band extends well above 1×; when drift is low or negative, L^* falls below 1×—the framework correctly says to de-lever.

Solving the Durability Condition across leverage ratios defines a *productive band*—the range where $\Delta R(L) > 0$. Figure 1 evaluates this relationship across representative (μ, σ) regimes.

These are closed-form evaluations of Equation (3); Section 4 tests them against actual returns. For Bitcoin-class parameters, $L = 1.33$ sits near the growth-rate optimum and comfortably inside the productive band at $s = 1\%$.

3. Volatility Segmentation

We introduce an on-chain protocol that provides durable, buy-and-hold leverage through volatility segmentation.[5] The structure partitions asset volatility into two perpetual derivatives:

- **Perpetual Senior (Sr)**—a yield-bearing position with volatility protection, provided by Junior first-loss capital.
- **Perpetual Junior (Jr)**—capital-efficient leveraged exposure absorbing first-loss volatility.

Junior capital absorbs underlying price movements first, giving holders leveraged exposure. Senior exposure remains protected until Junior capital is completely depleted. Because leverage

[5] Volatility segmentation generalizes the tranche-based approach first proposed in the SPOT protocol [12], extending it from a stablecoin mechanism to a general-purpose leverage architecture. Both are built on the on-chain tranching protocol described by ButtonTranche [13].

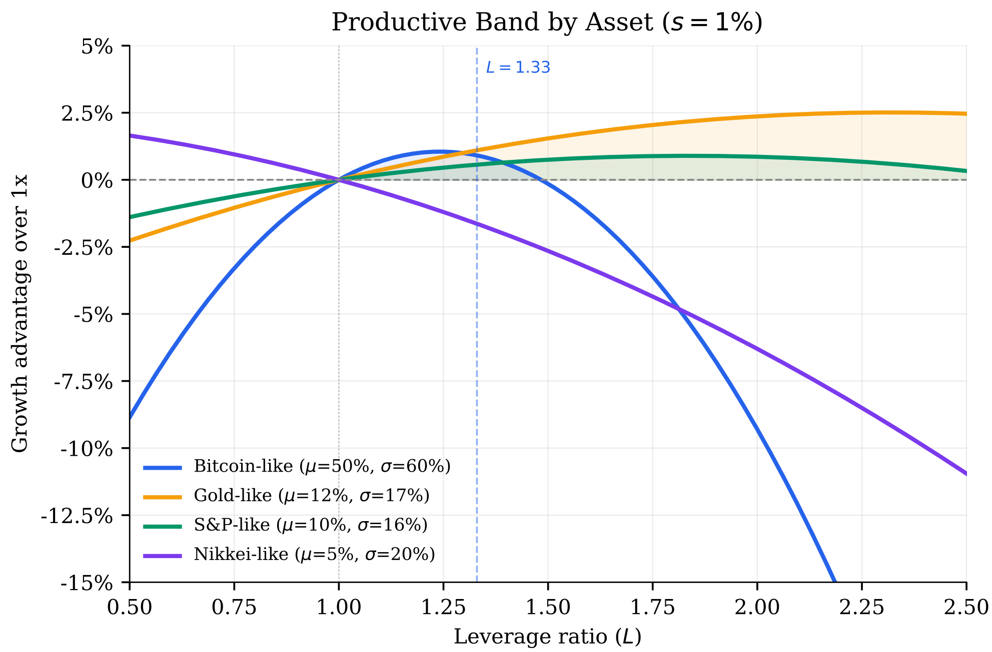

Figure 1: *Excess compound growth rate $R(L) - R(1)$ at $s = 1\%$ for four (μ, σ) regimes representative of different asset classes (see text). The shaded region above zero is the* productive band*—leverage adds value only inside it. The dashed line shows $L = 1.33$.*

arises from the capital structure itself, no external counterparty is needed to take the opposite side of a trade.

Applied to Bitcoin, the structure produces two perpetual derivatives: **Bitcoin Senior (BTC-Sr)**—downside-protected BTC exposure with yield—and **Bitcoin Junior (BTC-Jr)**—moderately leveraged BTC exposure. The remainder of this section walks through how these are configured; Section 4 tests them empirically.

3.1 Capital Structure

The protocol creates time-bounded bonds backed by an underlying asset (e.g. BTC). Each bond is split into a Senior tranche and a Junior tranche according to a configurable *tranche ratio* $\alpha \in (0, 1)$. Senior holds a priority claim; Junior absorbs residual price movement. A concrete example makes the mechanics clear.

Example. Consider a single bond created when BTC trades at \$100,000, with $\alpha = 0.50$. The bond holds 1 BTC of underlying, so the Senior tranche has a \$50,000 claim and the Junior tranche has a \$50,000 claim.

- **BTC rises 10% to \$110,000.** The Senior claim stays at \$50,000 (capped at par). The Junior receives the remainder: \$110,000 − \$50,000 = \$60,000, a 20% gain on its initial

$50,000—exactly 2× the underlying move.

- **BTC falls 10% to $90,000.** The Senior claim stays at $50,000 (protected). The Junior receives $90,000 − $50,000 = $40,000, a 20% loss—again 2× the underlying move.

Junior absorbs all price movement on half the capital, producing ≈2× leverage without borrowing.

General form. For tranche ratio α, Senior holds a par claim of $\alpha \cdot P_{\text{open}}$ per unit of underlying, where P_{open} is the underlying price at bond inception. At maturity, Senior is paid first up to par; Junior receives the remainder. This gives Junior an effective leverage of $1/(1-\alpha)$ at inception.

Bonds have a finite maturity. Perpetual exposure is achieved by automatically rolling maturing bonds into fresh ones; holders of either tranche maintain uninterrupted exposure without manual intervention.

The target leverage ratio L is configurable per underlying asset by adjusting the *target system ratio*—the ratio of Junior capital to Senior capital in the reserve. The appropriate L is determined by the asset's productive band (Section 2.3): the (μ, σ) regime pins down L^*, and the protocol targets leverage near that optimum. For BTC with $\alpha = 0.50$, the leverage ceiling is 2×; the protocol targets $L = 1.33\times$ (system ratio 3:1 Junior-to-Senior), which falls near L^* and well inside the productive band for BTC-class volatility.

3.2 Leverage Target Selection

The cost of getting leverage wrong is radically asymmetric. Undershooting the optimal ratio—holding 1.33× when L^* is 1.5×—sacrifices roughly a percentage point of annual return. Overshooting it—holding 2× or 3× when the productive band closes at 1.8×—is catastrophic: Section 4.1 shows that 3× BTC is effectively annihilated in the first major drawdown and never recovers. The penalty curve is not symmetric around the optimum; it is gentle on the left and vertical on the right.

This asymmetry explains why every major crypto leverage product destroys long-run value. The industry defaults to 2–3× because those ratios work in equities, where $\sigma \approx 19\%$ and the productive band extends past 3×. But Bitcoin's $\sigma \approx 60\%$ collapses the band to roughly 1.0–1.6×—and 2× is already outside it. The problem is not leverage itself but the uncritical transplant of equity-market intuitions onto crypto-volatility assets.

The protocol targets 1.33× for Bitcoin—a ratio that remains inside the productive band across every non-stress (μ, σ) regime tested in Section 4, from current Bitcoin down to equity-like parameters. The priority is regime-robustness, not point-optimality.

3.3 Funding

Like a perpetual future, the protocol uses a funding transfer to anchor the system to its target. Perp funding settles every eight hours based on the premium to spot; here, Junior pays Senior once daily based on the deviation from the target system ratio. The rate is market-determined: it emerges from the supply–demand balance between Senior and Junior capital, just as perp funding emerges from the balance between long and short interest.

We use a mature-deployment estimate of the BTC-Sr equilibrium yield of $\approx$4.5% (roughly $r_f \approx 3.5\%$ plus an $\sim$1% spread). The key point is *not* that 4.5% is a precise forecast; it is a reasonable long-run level once the protocol is mature, where the incremental spread above r_f is dominated by protocol risk rather than expected collateral impairment at the baseline configuration. Appendix A derives this from tranche payoffs (option equivalence) and shows why Senior financial risk over a single 28-day bond window is pushed far into the left tail. Empirically, the main results are not fragile to this point estimate: Jr durability is driven by being inside the productive band, and financing drag scales linearly with the assumed Sr yield.

At a 3:1 system ratio, Seniors earn 4.5% on \$100 of capital—\$4.50 per year. Spread across \$300 of Junior equity, this amounts to an effective financing cost of $\approx$ 1.5%.

System fees are governed by asymmetric, convex curves tied to the system's deviation from its target ratio. Minting or redeeming when the system is balanced costs nothing or near-nothing; transacting in a way that increases imbalance incurs escalating fees. Collected fees go to Junior holders, not to external parties, partially offsetting their financing cost below the $\sim$2.0% headline estimate. For the BTC deployment, the protocol charges a 0.50% AUM fee on Jr capital. The savings from cheaper financing are so large that Jr could charge several times this fee and still undercut most LETFs on total cost.

4. Empirical Analysis

We backtest the framework against daily prices from January 2015 to December 2025 for Bitcoin modeled using daily rebalancing. Three headline results:

1. **Moderate leverage is optimal for Bitcoin.** Beyond a narrow band, higher leverage destroys wealth (Section 4.1).
2. **The strategy survives bear markets.** Durability holds across full market cycles—bull surpluses more than compensate for bear drawdowns (Section 4.2).
3. **Lower financing drag compounds to materially higher wealth.** BTC-Jr's cost advantage over LETFs and perpetual futures widens with every year of compounding.

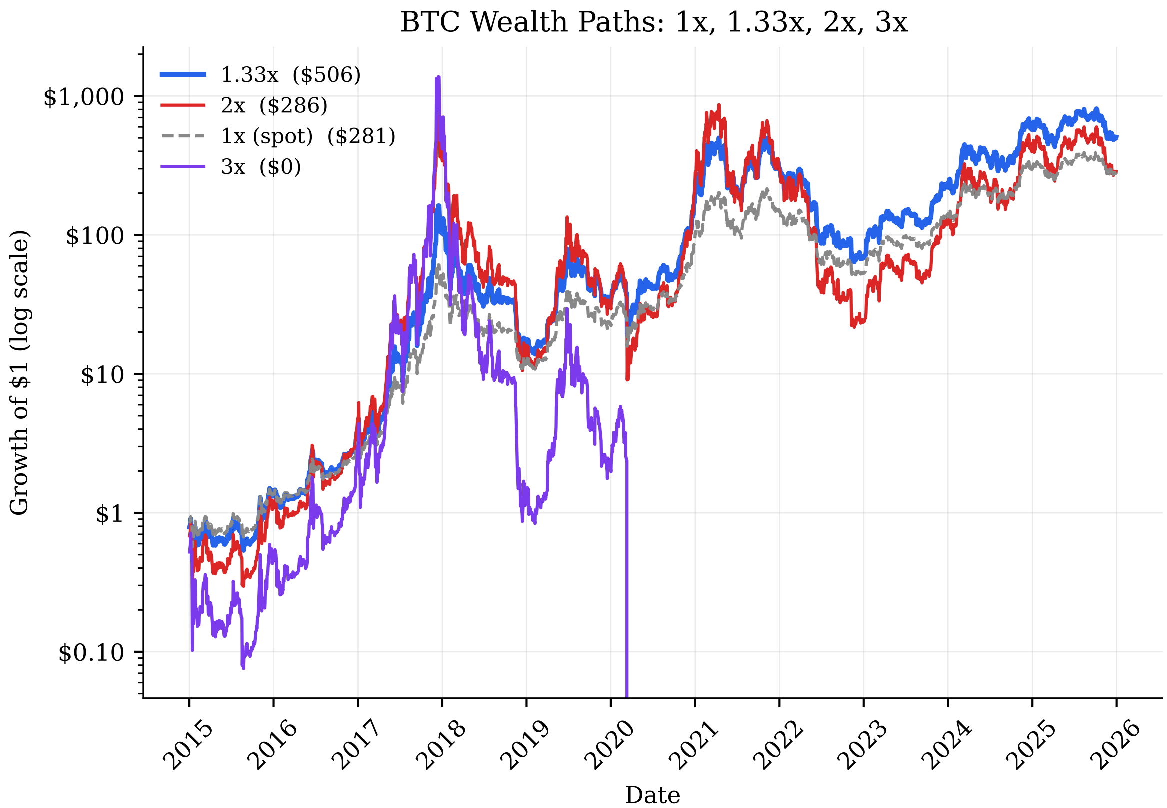

Figure 2: *Growth of $1 in BTC at four leverage ratios (2015–2025), all using BTC-Jr parameters. 1.33× dominates throughout; 2× suffers catastrophic drawdowns from which it never fully recovers; 3× is effectively annihilated. Terminal values shown in legend.*

The result is robust to assumptions and generalizes across assets and drift–volatility regimes.[6]

4.1 Leverage and Long-Run Wealth

Figure 2 plots the growth of $1 invested in BTC at 1×, 1.33×, 2×, and 3×, all using BTC-Jr parameters (0.50% AUM fee, financing premium $s = 1\%$).

The 1.33× path dominates not just in terminal value but throughout the sample: it overtakes spot early and never relinquishes the lead. The 2× position suffers catastrophic drawdowns during bear markets—the quadratic volatility penalty amplifies losses to the point where the position never fully recovers. At 3×, the penalty is so severe that the position is effectively annihilated within the first major drawdown. This is the Durability Condition playing out in real time: drift exceeds drag only at moderate leverage.

[6]Sensitivity analysis and stress tests are in the supplementary material, Section 1.

4.2 Regime Robustness

A long-duration leverage strategy should not be judged by whether it outperforms in every subperiod. In sustained drawdowns, leveraged exposure will typically underperform spot because it magnifies losses and drag accumulates precisely when capital is shrinking. The relevant question for strategic holders is therefore whether the *compounding surplus* earned in up and trending environments outweighs the *compounding deficit* incurred in down environments over a full market cycle.[7]

Put differently: the strategy is durable if it can compound faster than spot across complete cycles—even if it predictably lags during drawdowns.

Over the full 2015–2025 period, the surplus from up environments dominates the deficit from down environments, so the strategy outcompounds spot in aggregate.

To examine these dynamics more concretely, Figure 4 zooms into three representative periods: a bull-dominated stretch (January 2023 to December 2025), a sustained down market (November 2021 to December 2022), and a sideways/range-bound market (January to August 2024).

Bull-dominated. Over Jan 2023–Dec 2025, $10,000 in BTC grows to $53,200; the same capital in BTC-Jr reaches $74,200—40% outperformance despite the additional drag.

Down market. During the 2022 crypto winter, BTC-Jr draws down to $1,390 vs. BTC's $2,480—an extra $\sim$11 percentage points of additional drawdown, the premium paid for bull-market outperformance. Crucially the drawdown is *painful but not terminal*: modest leverage preserves enough capital for the subsequent recovery.

Sideways market. In the range-bound first half of 2024, BTC-Jr holds at $13,600, broadly tracking the underlying. By contrast, a 3× LETF decays to $10,031—and real-world products like the GraniteShares 3× Long MicroStrategy ETP [14] suffer even more severe drag—becoming effectively unholdable.

Taken together, the regime decomposition and the per-condition close-ups confirm the main empirical claim: 1.33× BTC-Jr is *durable*—it compounds in trending markets, survives drawdowns, and avoids decay in sideways conditions.

5. Economics

The preceding analysis establishes the empirical case. This section examines the structural economics: why the capital is cheaper, why moderate leverage is optimal, where the architecture applies, and where it does not.

[7] Regimes are classified by manual date ranges aligned to major market turning points (e.g. the November 2021 all-time high marks the start of the 2022 bear market). See Figure 3 for the exact boundaries.

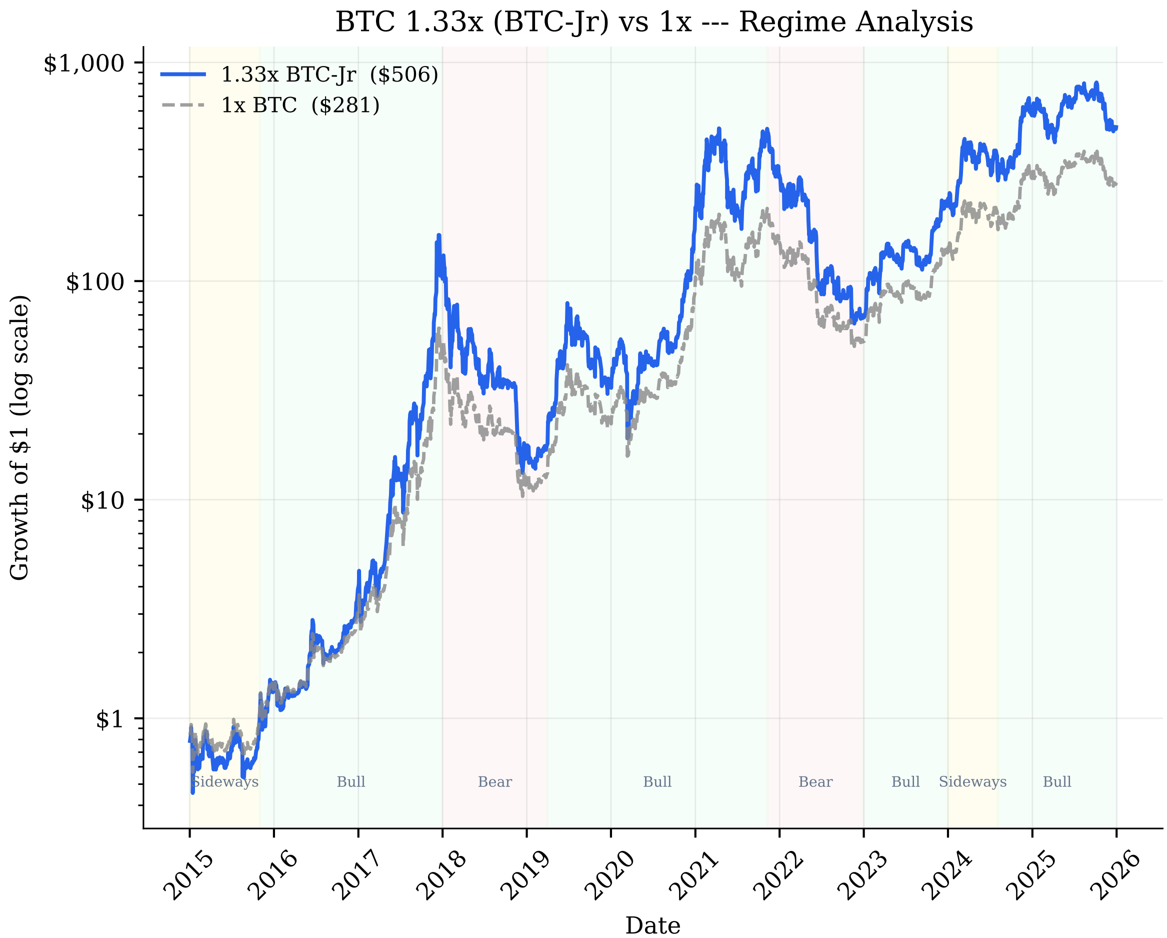

Figure 3: *BTC 1× vs 1.33× (BTC-Jr) with regime shading. Green: bull regimes. Yellow: sideways. Red: bear regimes.*

5.1 Capital Structure Advantage

The primary innovation is not the leverage ratio but the *source* of leverage. Every existing instrument couples its financing premium to the hedging costs of a specialized counterparty—and for Bitcoin, those costs are extreme. Volatility segmentation replaces these counterparties with yield-seeking depositors, lowering the financing premium by an order of magnitude and compounding into a material wealth gap over time.

The leverage providers behind existing products must actively hedge and consume balance sheet; those costs flow through as financing premia. A Senior depositor, by contrast, benchmarks yield against treasuries—not swap-dealing returns—so the required spread is small. The supply pool is also wider: yield-seeking capital is a broader base than regulated dealer balance sheets, and on-chain settlement extends it globally.

The Sr yield used throughout this paper ($r_f + \sim 1\%$, or $\approx 4.5\%$ at current rates) is our

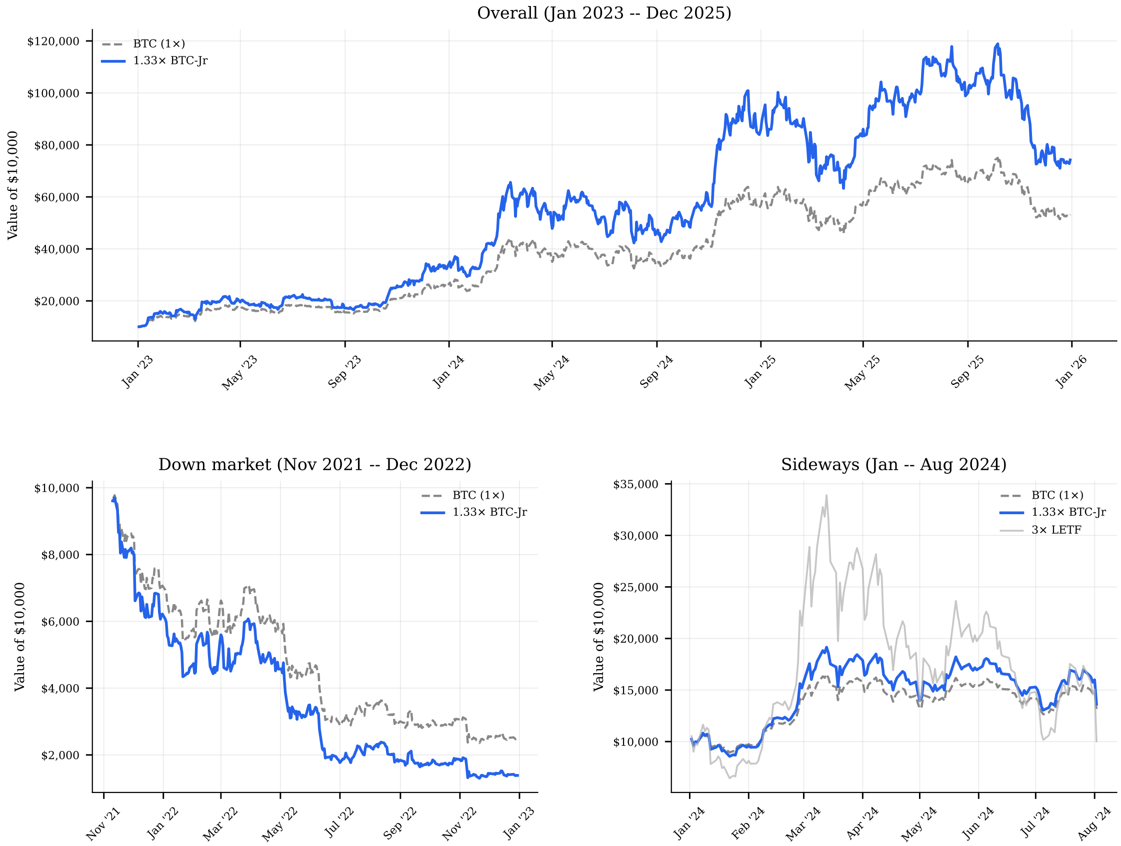

Figure 4: *Value of $10,000 invested in BTC (1×) vs 1.33× BTC-Jr across three market conditions.* ***Top:*** *Overall (Jan 2023 – Dec 2025).* ***Bottom left:*** *Down market (Nov 2021 – Dec 2022).* ***Bottom right:*** *Sideways (Jan – Aug 2024), with a 3× LETF shown for comparison. All simulations use financing premium $s = 1\%$ and 0.50% Jr AUM fee.*

best estimate, not an observed market rate; supply and demand will set the actual number. The tranche payoffs formalize why the spread should be thin. Junior's position is equivalent to a deep in-the-money call; Senior is the residual (Appendix A derives this). The financial risk to Senior is negligible—the underlying would need to drop more than 50% within a single 28-day bond period. Under log-normal scaling this is a $\sim 4\sigma$ event; the worst 28-calendar-day BTC drawdown on record is -54% (January 2018), which just breaches this threshold once in 11 years.[8] The spread above r_f is almost entirely protocol risk, which compresses as the codebase matures. Structural comparisons corroborate this: with $\alpha = 0.5$, Seniors sit above 50% Junior subordination on short-dated collateral, a conservative secured-credit profile. On-chain, large lending venues clear near r_f for overcollateralized lending with *no* subordination; Seniors here receive strictly stronger protection. Early yields will likely be higher—the market will price

[8] Senior impairment at $\alpha = 0.50$ requires a $>50\%$ drop in a single 28-day bond window. The January 2018 episode (-54% over 28 calendar days) is the only instance in the dataset. Shorter bond periods or a lower α would eliminate it entirely.

novelty risk—then compress as capital scales in. (Ethena's sUSDe, a tokenized yield product, converged toward the risk-free rate after supply grew from $85M to $12B [15, 16].)

Existing leverage products fail the Durability Condition for the reasons identified in Sections 3.2 and 2.2: leverage ratios imported from equity markets fall outside the productive band for Bitcoin-class volatility; LETFs embed hidden swap costs that dwarf their stated expense ratios; and perpetual futures charge funding on the full notional—not just the borrowed portion. In each case, total drag exceeds the drift from leverage.

5.2 Moderate by Design

The 1.33× target is not an optimization residual—it is determined by the capital structure: a 3:1 system ratio at $\alpha = 0.50$ yields $L = 1.33$ (Section 3), independent of any backtest. At Bitcoin-class volatility, the penalty for overshooting the productive band is orders of magnitude larger than the penalty for undershooting it (Section 3.2), which means the regime-robust choice sits at or below the point estimate of L^* across the range of plausible regimes. For high-volatility assets, moderate leverage is not a compromise but the structural optimum—the ratio that compounds reliably across the widest range of market conditions.

The structure is also self-dampening under stress. If Senior depositors withdraw during a drawdown, the system ratio shifts toward Junior capital, which mechanically *reduces* leverage—Junior exposure moves closer to 1× spot. The risk is asymmetric in the safe direction: stress compresses leverage rather than amplifying it. Junior holders may temporarily receive less leverage than they pay for, but the capital structure prevents the dangerous scenario—leverage spiking precisely when drawdowns are deepest.

5.3 Beyond Bitcoin

The mechanism generalizes beyond Bitcoin. Volatility segmentation does not rely on Bitcoin-specific market structure; it relies on a generic property of leveraged compounding: the value of leverage depends on the underlying drift–volatility regime and on the all-in financing premium. Where the underlying exhibits positive long-run drift, and where the financing premium can be compressed, moderate leverage can sit inside the productive band and remain compatible with long-horizon allocation.

The more interesting question is why this regime has been so difficult to serve. Historically, "durable leverage" was largely out of reach because the standard way to industrialize leverage is through deep derivatives infrastructure. In traditional markets, efficient, continuously rebalanced leverage depends on the ability of intermediaries to warehouse and hedge risk at scale—which, in turn, requires extremely deep futures, options, and swap markets. Those markets take years to develop and often never develop at meaningful depth; when they do, it is typically for major indices and a small set of benchmark underlyings. Outside those venues, the financing premium remains structurally high and unstable, and continuously rebalanced leverage tends to sit outside the productive band.

Given that constraint, the market optimized for a different customer and a different objective: tactical short-horizon leverage. At trading horizons, users are not especially sensitive to compounding efficiency; directional P&L dominates, and the product can tolerate wide premia and implementation friction. The result is an ecosystem of leverage tools designed for speculation rather than long-duration capital allocation.

Volatility segmentation changes the feasibility frontier. Because the leverage is embedded in a capital structure that can be implemented on-chain, it does not require pre-existing, institution-scale derivatives markets to approach the productive band. A reliable price feed and a depositor base are sufficient. This opens a new product category: moderate, automated leverage engineered for buy-and-hold investors, not just short-term traders.

6. Limitations

1. **Outperformance depends on realized drift.** All simulations are backward-looking; historical returns are not representative of future performance. Leverage adds value only when incremental drift exceeds incremental drag. If the underlying's realized drift is persistently negative, leveraged positions underperform 1× by construction. The framework handles any μ, but this paper scopes to assets with positive secular drift. The protocol allows adjusting L over time as the asset evolves.

2. **Constant leverage, constant premium assumption.** The Durability Condition is a mathematical result; the protocol is an architecture designed to satisfy it. All backtests reset the leverage ratio to L after each day's return and hold the financing premium s constant, the standard approach in the LETF literature [9, 10]. In practice, Jr's exposure can drift from the target. When BTC moves, leverage deviates; the protocol restores it through a feedback loop: the funding rate adjusts, arbitrageurs deposit or withdraw Sr capital, and leverage returns to target.

   If the feedback is slow, the two quantities drift together. During a rally, leverage gradually falls below target while s rises; during a drawdown, leverage rises above target while s drops. Fee curves bound the deviation by making the rebalancing trade cheap near the target, and the sensitivity analysis (Supplementary Section 1) shows durability survives s rising far above 1%. The residual risk is rate-insensitive Sr capital: if depositors do not respond to the funding signal, leverage drifts and Jr's realized path differs from the simulation. But data shows on-chain yield-seeking depositors are rate-sensitive [15].

3. **Smart-contract and oracle risk add a yield premium.** Appendix A shows that the

$\sim$1% spread above r_f is almost entirely protocol risk—smart-contract bugs, oracle failures, governance attacks—not financial risk. For a new protocol this premium is real and may initially widen s above the 1% estimate. Historical DeFi exploit rates have fallen sharply [17]; as the protocol matures, the premium compresses.

4. **Intraday exposure drift.** The simulations assume end-of-day rebalancing to the target leverage ratio and do not model intraday entry/exit timing effects (exposure drift / tracking error). These effects can matter over short horizons, but for long-horizon holders they are typically second-order relative to volatility drag and financing drag, because the impact of any single day's entry/exit timing shrinks as the holding period grows.

7. Conclusion

Leverage is durable when incremental drift exceeds drag. Every leveraged product pays a cost to stay leveraged. You cannot avoid volatility drag—that is physics—but you can choose who you borrow from and how much you pay. Today's options are expensive. LETFs hide double-digit swap rates inside an expense ratio. Perpetual futures charge funding on the full notional, not just what you borrowed. Margin loans price in the lender's worst case. At Bitcoin volatilities, these costs eat most or all of the extra return leverage is supposed to deliver.

Volatility segmentation changes the economics. Depositors, not dealers, supply the capital, and they accept a fixed yield instead of demanding a risk premium. The financing cost drops by an order of magnitude. That difference—a few percent a year—compounds quietly in the background, and over a decade it separates a position that builds wealth from one that merely matches spot.

The mechanism is simple: a price feed and a depositor base. No derivatives market required. Any asset with positive long-run drift clears the bar, which means the same structure works beyond crypto. The point of this paper is that durable leverage is now a design choice, not a market constraint.

The companion Supplementary Material [18] provides multi-asset simulations, LETF cost decomposition, fee capacity analysis, and a leverage instrument survey.
https://drive.google.com/file/d/1VPu2SOqdjIAz22j7rs_k-DOupghCBoWm/view

---

Table 5 pools the full 2015–2025 history. Figure 5 and Tables 1–3 repeat the sweep within each of the three BTC market regimes studied in Section 4.2. In the bull (2023–2025) and sideways (Jan–Aug 2024) regimes, durability holds across the entire tested range: ΔR remains positive even at $r_f + 20\%$. In the bear regime (Nov 2021–Dec 2022), leverage amplifies losses regardless of the Sr yield—the lines in Figure 5 are nearly flat and uniformly below zero—confirming that the bear-market outcome is driven by the direction of spot returns, not by financing costs.

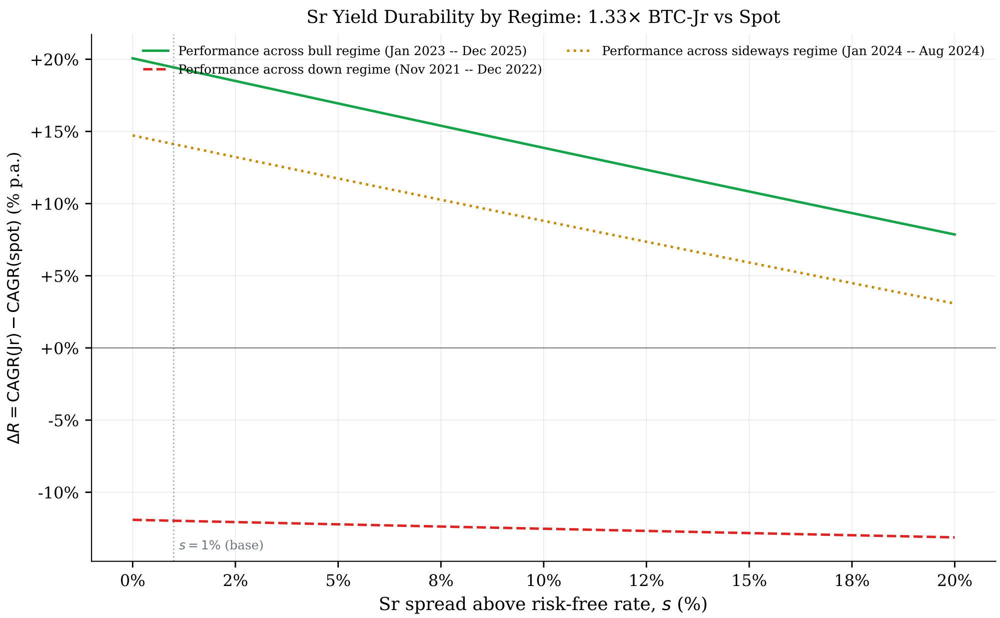

Figure 5: *Excess compound growth $\Delta R = R(1.33\times) - R(1\times)$ as a function of Sr yield spread s for each BTC market regime. Bull and sideways regimes remain durable across the full $[r_f,\, r_f + 20\%]$ range. The bear regime shows near-flat negative ΔR at all s: leverage amplifies losses independent of Sr yield.*

y_{Sr}	Jr \$1 $\rightarrow$	ΔR	Durable?
r_f	\$7.50	+20.1%	Yes
$\boldsymbol{r_f + 1\%}$ **(baseline)**	**\$7.42**	**+19.4%**	**Yes**
$r_f + 5\%$	\$7.13	+16.9%	Yes
$r_f + 10\%$	\$6.77	+13.9%	Yes
$r_f + 15\%$	\$6.43	+10.8%	Yes
$r_f + 17.5\%$	\$6.27	+9.3%	Yes
$r_f + 20\%$	\$6.11	+7.9%	Yes

Table 1: *Performance across bull regime (Jan 2023 – Dec 2025). Daily-rebalanced simulation on BTC (2023-01-01 to 2025-12-31) with time-varying $r_f(t)$. Durable across the full tested range (s = 0–20%). Spot: \$1 $\rightarrow$ \$5.32.*

y_{Sr}	Jr \$1 $\to$	ΔR	Durable?
r_f	\$0.14	-11.9%	No
$r_f + 1\%$ (baseline)	**\$0.14**	**-12.0%**	**No**
$r_f + 5\%$	\$0.14	-12.2%	No
$r_f + 10\%$	\$0.13	-12.5%	No
$r_f + 15\%$	\$0.13	-12.8%	No
$r_f + 17.5\%$	\$0.13	-13.0%	No
$r_f + 20\%$	\$0.13	-13.1%	No

Table 2: *Performance across down regime (Nov 2021 – Dec 2022). Daily-rebalanced simulation on BTC (2021-11-10 to 2022-12-31) with time-varying $r_f(t)$. Not durable across the tested range; leverage amplifies losses in a down market regardless of Sr yield. Spot: \$1 $\to$ \$0.25.*

y_{Sr}	Jr \$1 $\to$	ΔR	Durable?
r_f	\$1.45	$+14.7\%$	Yes
$r_f + 1\%$ (baseline)	**\$1.45**	**$+14.1\%$**	**Yes**
$r_f + 5\%$	\$1.43	$+11.7\%$	Yes
$r_f + 10\%$	\$1.42	$+8.8\%$	Yes
$r_f + 15\%$	\$1.40	$+5.9\%$	Yes
$r_f + 17.5\%$	\$1.40	$+4.5\%$	Yes
$r_f + 20\%$	\$1.39	$+3.1\%$	Yes

Table 3: *Performance across sideways regime (Jan 2024 – Aug 2024). Daily-rebalanced simulation on BTC (2024-01-01 to 2024-08-04) with time-varying $r_f(t)$. Durable across the full tested range ($s = 0$–20%). Spot: \$1 $\to$ \$1.38.*

References

[1] Kelly, J. L. (1956). A new interpretation of information rate. *Bell System Technical Journal*, 35(4), 917–926.

[2] Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3), 247–257.

[3] Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51(3), 239–246.

[4] Ayres, I. and Nalebuff, B. (2013). Diversification across time. *Journal of Portfolio Management*, 39(2), 73–86.

[5] Gayed, M. A. and Bilello, C. (2016). Leverage for the long run—A systematic approach to managing risk and magnifying returns in stocks. Winner of the 2016 Charles H. Dow Award, *CMT Association*.

[6] Peters, O. (2011). Optimal leverage from non-ergodicity. *Quantitative Finance*, 11(11), 1593–1602.

[7] BestETF.net (2026). Leveraged ETF list.

[8] BitMEX Research (2025). Q2 2025 derivatives report: 9 years of funding rate data.

[9] Avellaneda, M. and Zhang, S. (2010). Path-dependence of leveraged ETF returns. *SIAM Journal on Financial Mathematics*, 1, 586–603.

[10] Cheng, M. and Madhavan, A. (2009). The dynamics of leveraged and inverse exchange-traded funds. *Journal of Investment Management*, 7(4), 43–62.

[11] Thorp, E. O. (2006). The Kelly criterion in blackjack, sports betting, and the stock market. In *Handbook of Asset and Liability Management*, North Holland.

[12] Fragments Research (2022). SPOT — An Inflation Resistant Store of Value.

[13] Buttonwood Foundation (2021). ButtonTranche: Zero-Coupon Bond Protocol.

[14] GraniteShares (2025). 3× Long MicroStrategy Daily ETP (3LMI). Product page and KID. ISIN: XS2617255760.

[15] DefiLlama (2026). Ethena protocol dashboard.

[16] Ethena Labs (2025). Ethena 2025: Convergence.

[17] Lantzsch, C. / Cicada Partners (2025). The state of DeFi exploit risk. *CoinDesk Indices*, October 8. coindesk.com.

[18] Fragments Research (2026). Durable Leverage: Supplementary Material.

[19] Roy, S. (2025). Leveraged ETFs: The hidden costs that eat your returns. *etf.com*.

[20] Garleanu, N., Pedersen, L. H., and Poteshman, A. M. (2009). Demand-based option pricing. *The Review of Financial Studies*, 22(10), 4259–4299.

[21] Morningstar (2025). These ETFs had one job. They didn't do it, costing traders billions.

[22] Aramonte, S., Schrimpf, A., and Shin, H. S. (2024). The prime broker–hedge fund nexus: Recent evolution and implications for bank risks. *BIS Quarterly Review*, March. bis.org.

[23] Charles Schwab (2025). Margin requirements and interest rates.

[24] Interactive Brokers (2026). Margin rates and financing.

[25] Ledn (2026). Best Bitcoin-backed loan rates: Compare terms and LTV.

[26] FINRA (2025). Margin statistics.

[27] ProShares (2025). Fund documentation for SSO, UPRO, and TQQQ.

[28] BitMEX (2024). Perpetual contracts guide: Funding.

[29] CoinGlass (2025). Bitcoin futures open interest.

[30] Muravyev, D. and Pearson, N. D. (2020). Options trading costs are lower than you think. *The Review of Financial Studies*, 33(11), 4973–5014.

Appendix

A. Senior Yield Assumption

A key input to the BTC-Jr model is the equilibrium Senior yield y_{Sr}. In the main results we use $y_{\mathrm{Sr}} \approx 4.5\%$ (roughly $r_f \approx 3.5\%$ plus a $\sim$1% premium). This appendix derives that level from the tranche payoffs, which are equivalent to standard options. At the baseline configuration the *financial* risk to Senior over a single bond window is negligible; the spread above r_f is therefore dominated by protocol risk, the one component that compresses with time.

A.1 Tranche payoffs and option equivalence

A bond holds underlying worth S_0 at inception, with tranche ratio α and maturity τ. At settlement, the Senior tranche is paid first up to par; Junior receives the residual:

$$V_{\mathrm{Sr}} = \min\big(S_T,\ \alpha\, S_0\big) \tag{6}$$

$$V_{\mathrm{Jr}} = \max\big(0,\ S_T - \alpha\, S_0\big) \tag{7}$$

with $V_{\mathrm{Sr}}+V_{\mathrm{Jr}} = S_T$. Junior's payoff is equivalent to a call on the underlying with strike $K = \alpha\, S_0$; Senior is the residual.

Junior pays Senior a running yield for this leveraged exposure. The protocol implements this by rolling maturing bonds into fresh ones every τ days. Each period, Jr's cost is the time value of the embedded call, paid to Sr as yield.

A call's value has two parts: intrinsic value and time value (the premium for optionality). If α is sufficiently small, the call is deep in the money: it moves nearly dollar-for-dollar with the underlying and the chance of expiring below the strike is negligible. In that regime, the time value reduces to (approximately) the interest on the strike:[9]

$$\mathrm{TV} \;\approx\; K\, r_f\, \tau.$$

Whether this approximation holds depends on α, σ, and τ.

[9] The Black–Scholes call price is $C = S_0\, N(d_1) - K\, e^{-r_f \tau}\, N(d_2)$. When the call is deep in the money, both $N(d_1)$ and $N(d_2)$ approach 1, giving $C \approx S_0 - K\, e^{-r_f \tau}$. Subtracting intrinsic value $(S_0 - K)$ leaves time value: $K(1 - e^{-r_f \tau}) \approx K\, r_f\, \tau$.

A.2 Senior yield decomposition

Senior depositors provide capital and receive a yield funded by Junior. The required yield compensates for three components:

$$y_{\mathrm{Sr}} = r_f \ + \underbrace{\pi_{\mathrm{put}}(\alpha, \sigma, \tau)}_{\text{collateral risk}} \ + \underbrace{\pi_{\mathrm{lock}}(\tau)}_{\text{exit friction}} \ + \underbrace{\pi_{\mathrm{contract}}}_{\text{protocol risk}} . \tag{8}$$

π_{put} captures any compensation beyond r_f for the risk that the call moves out of the money within one bond period (i.e. the underlying drops below the strike). It depends on tranche ratio α (which sets the strike), underlying volatility σ, and maturity τ. π_{lock} captures the cost of exiting (fees and any lockup during bond maturity). π_{contract} covers smart contract, oracle, and governance risk.

A.3 Evaluating the terms at the baseline configuration

Collateral risk (π_{put}). Senior loses value only if the underlying drops below the strike within one bond window. With short-dated bonds ($\tau = 28$ days), this requires an extreme move in a narrow window. Table 4 evaluates several candidate deployments. BTC and NVDA are the tightest cases: a >50% or >45% 28-day drop corresponds to an event on the order of ∼4 standard deviations under diffusion scaling, with $P(\mathrm{OTM}) < 0.002\%$ under the Black–Scholes benchmark.[10] The worst 28-calendar-day BTC drawdown on record is −54% (January 2018), which just breaches the $\alpha = 0.50$ threshold—the only such instance in 11 years of data. Shorter bond periods or a lower α would eliminate it. For lower-volatility assets (S&P, Gold), strikes can sit closer to spot and still remain deep in the money. At all configurations shown, $\pi_{\mathrm{put}} \approx 0$.

Asset	$\boldsymbol{\sigma}$	$\boldsymbol{\alpha}$	**Sys. ratio**	**Target L**	**Drop to Sr loss**	$\boldsymbol{P}$**(OTM)**
BTC	60%	0.50	3.0	1.33×	>50%	< 0.002%
NVDA	50%	0.55	1.0	2.0×	>45%	< 0.002%
S&P 500	19%	0.75	1.3	1.75×	>25%	≈ 0
Gold	18%	0.80	0.7	2.35×	>20%	< 0.001%

Table 4: *Candidate configurations at $\tau = 28$ days. Target leverage $L = 1 + 1/\text{system ratio}$. $P(OTM) = N(-d_2)$ is the Black–Scholes probability the call expires out of the money. When $P(OTM) \approx 0$, $\pi_{put} \approx 0$.*

Exit friction (π_{lock}). Senior depositors are not structurally locked in. Fast redemption can convert Sr tokens to underlying in a single transaction; regular redemption returns bond tranches that settle at maturity (at most 28 days). During that wait, Sr's par claim is fixed in dollar terms, so the opportunity cost is roughly $r_f \times 28/365 \approx 27$ bps at current rates. The remaining

[10] Bitcoin returns exhibit fat tails, so this should be interpreted as a conservative calibration aid rather than a literal tail model.

friction comes from deviation-based fees tied to the system's distance from target; those fees are highest precisely when Sr is least likely to exit (when Sr is earning well and the system has excess Jr), and near-zero when the system wants Sr to leave to rebalance. Overall, π_{lock} is small.

Protocol risk ($\pi_{contract}$). The only component that cannot be eliminated by configuration. Smart contract bugs, oracle failures, governance attacks—engineering risk rather than financial risk. We calibrate π_{contract} from observable proxies: DeFi insurance markets price coverage at ∼1–3% for newer protocols, declining toward ∼0.1–0.3% for mature ones. Historical exploit rates have declined from ∼30% of TVL in 2020 to $< 0.5\%$ in 2024 [17] as audit practices and operational track record accumulate. The exact number is uncertain, but the direction is clear: this component compresses with time.

Result. Collateral risk drops out at the baseline configuration and exit friction is small. The spread above r_f is therefore dominated by protocol risk:

$$y_{\text{Sr}} \approx r_f + \pi_{\text{contract}} \approx r_f + 1\%.$$

Whatever Senior earns above r_f is not primarily compensation for expected financial loss; it is compensation for trusting the settlement mechanism. The Sr yield is market-determined and will fluctuate with the Jr–Sr capital balance, but the equilibrium level is set by r_f plus protocol risk. As the codebase accumulates operating history, π_{contract} compresses and y_{Sr} converges toward r_f.

A.4 Sensitivity to Senior yield

The preceding subsections derive $y_{\text{Sr}} \approx r_f + 1\%$ as a mature-deployment estimate. To show that the durability result is not fragile to this point estimate, Table 5 sweeps the Sr yield from r_f to $r_f + 20\%$, using time-varying $r_f(t)$ from 13-week T-bill yields, and simulates 1.33× BTC-Jr on daily returns (2015–2025). We report excess compound growth (Equation 4):

$$\Delta R \;=\; R(1.33) - R(1).$$

y_{Sr}	Jr \$1 $\rightarrow$	ΔR	Durable?
r_f	\$525	+9.3%	Yes
$r_f + 1\%$ (baseline)	**\$506**	**+8.7%**	**Yes**
$r_f + 5\%$	\$435	+6.4%	Yes
$r_f + 10\%$	\$361	+3.6%	Yes
$r_f + 15\%$	\$299	+0.9%	Yes
$r_f + 17.5\%$	\$272	−0.5%	No
$r_f + 20\%$	\$248	−1.8%	No

Table 5: *Sensitivity of BTC-Jr (*$1.33\times$*, 0.50% AUM) to the Sr yield. Daily-rebalanced simulation on BTC 2015–2025 with time-varying* $r_f(t)$*. Breakeven at* $y_{Sr} \approx r_f + 16.6\%$*. Spot: \$1* $\rightarrow$ *\$281.*

The baseline estimate sits at the bottom of a wide band: even at $r_f + 10\%$, Jr still delivers $\Delta R > +3.6\%$. Durability breaks only at $y_{\mathrm{Sr}} \approx r_f + 16.6\%$. Since the preceding subsections show that the premium is dominated by protocol risk—calibrated at ∼1–3% for new protocols and compressing toward 0.1–0.3% with maturity—realistic Sr yields sit far from the breakeven. The Supplementary Material (Section 1.3) extends this to a two-dimensional sweep over both the Sr yield and AUM fee, confirming the same conclusion.

Window	r_f	$r_f + 1\%$	$r_f + 5\%$	$r_f + 10\%$
Bull: 2023–25 (BTC +430%)	+20.1%	+19.4%	+16.9%	+13.9%
Sideways: 2024 H1 (BTC +32%)	+14.7%	+14.1%	+11.7%	+8.8%
Bear: 2021–22 (BTC −74%)	−11.9%	−12.0%	−12.2%	−12.5%

Table 6: ΔR *at varying Sr yields for the three windows examined in Section 4.2—up, down, and sideways markets—plus the full period. Each row simulates* $1.33\times$ *BTC-Jr with time-varying* $r_f(t)$*.*

Table 6 shows that in the bull window, ΔR stays above +13% even at $r_f + 10\%$—higher s shaves the surplus but does not eliminate it. The sideways window tells a similar story: moderate drift is enough for Jr to outperform across the entire spread range. In the bear window, ΔR is nearly insensitive to s: in a sharp decline, the leveraged price loss dwarfs the financing cost, so varying the spread barely moves the outcome.