Hedging a Leveraged ETF Position: Convexity Protection with Guaranteed Floors on TSLA/TSLL

Daksh Kumar

July 2026

Abstract

Leveraged exchange traded funds decay under realized volatility, which makes a short position in one an attractive but dangerous trade: the decay harvest is real, and the unhedged drawdowns are catastrophic. This report studies the problem on the TSLA/TSLL pair (Tesla and its 2x daily bull ETF) through four strategies of increasing discipline: the naked short benchmark, a delta-adjusted arbitrage that is rejected for look-ahead bias despite the best headline numbers in the project, convexity protection with a static guaranteed floor, and finally a multi-tranche engine whose maximum drawdown is capped by a closed-form, model-free bound and whose per-tranche loss floor holds continuously rather than only at expiry. Every risk statement of the final engine is registered before the data run and verified afterward: across all configurations tested, the daily equity path never violates its leak-adjusted analytic floor. With approximate borrow fees included, the call-only engine earns 13.4 percent annually with a 15.4 percent maximum drawdown, against a naked-short benchmark drawdown of 79.6 percent.

1 Introduction

Leveraged exchange traded funds (LETFs) target a multiple of their underlying's *daily* return. The daily reset forces the fund to buy exposure after up days and sell after down days, and over multi-day horizons this rebalancing bleeds value in proportion to realized variance. Shorting the LETF is therefore a structural volatility harvest; it is also, unhedged, a trade that can lose several times its initial capital when the underlying rallies. The subject of this report is how to keep the harvest while making the loss profile survivable, on the most volatile single-stock LETF pair available to the team: TSLA and TSLL, the Direxion Daily TSLA Bull 2X ETF.

Four strategies are developed and evaluated on a common data infrastructure (Kumar, 2026). Strategy 1, the naked short, establishes the benchmark and the problem: a drawdown near 80 percent. Strategy 2, a delta-adjusted arbitrage, produces the best headline risk-adjusted numbers in the project and is nevertheless rejected; its edge does not survive scrutiny of its information structure or its financing costs, and its failure motivates the verification discipline used afterward. Strategy 3, convexity protection, replaces estimated hedge ratios with an option's contractual payoff and culminates in a static guaranteed floor whose per-cycle worst case is an algebraic identity. Strategy 4 generalizes that identity into a portfolio engine with a provable drawdown cap, a put-financed collar, and realized-cost accounting.

The report's contributions are: (i) an exact per-tranche loss reservation valid at any strike, holding continuously through model-free no-arbitrage bounds; (ii) a settled-equity accounting scheme that removes mark-driven procyclicality from capital deployment; (iii) a closed-form maximum drawdown bound and a one-parameter risk dial derived from it; (iv) a liquidity-gated collar that finances the floor by selling rich LETF crash insurance; and (v) a pre-registration and verification protocol under which every published guarantee was checked against the realized daily path and never violated.

2 Literature Review

2.1 Option pricing and hedging

The theoretical basis for hedging with options begins with Black and Scholes (1973), who price European options by continuous replication and in doing so establish the equivalence between an option and a dynamically rebalanced position in the underlying. The textbook treatment of the Greeks, American exercise features, and model-free no-arbitrage restrictions used throughout this report (intrinsic-value floors, put-call parity bounds, vertical spread bounds) follows Hull (2018). Two properties matter here. First, a long call's delta approaches one as the option moves into the money, which is what lets a share-matched call cap a short position's loss without rebalancing. Second, the no-arbitrage restrictions hold independently of any pricing model, which is what allows this report's floor guarantees to be stated without assuming a volatility process.

2.2 Leveraged ETF decay

Cheng and Madhavan (2009) derive the return dynamics of leveraged and inverse ETFs and show that the daily reset induces a compounding drag relative to the naive multiple of the underlying's return. Avellaneda and Zhang (2010) give the canonical decomposition: over a horizon with underlying gross return R and realized variance V, a β-times daily LETF returns approximately

$$R_{\text{LETF}} \;\approx\; R^{\beta}\exp\Big(-\tfrac{\beta^2-\beta}{2}\,V\Big), \tag{1}$$

so for $\beta = 2$ the drag term is e^{-V}: the LETF pays away its underlying's realized variance. Shorting the LETF collects this term, which is the economic motivation for every strategy in this report; the exploratory analysis in Section 4 verifies the decomposition empirically for TSLL.

2.3 Backtest methodology and optimization bias

Pardo (2008) treats the evaluation of trading strategies as an exercise in out-of-sample discipline: parameters fitted where they are tested overstate performance, and walk-forward evaluation is the corrective. López de Prado (2018) sharpens the same point for modern research workflows, cataloguing backtest overfitting, data leakage, and the multiplicity problem of repeated trials on one price path. Both inform this report directly: Strategy 2 is rejected on exactly these grounds despite its headline numbers, and the final engine responds by fixing its parameters a priori, registering its risk bounds before touching data, evaluating a fixed configuration on disjoint holdout halves, and refusing to re-fit its one behavioral threshold on the sample.

2.4 Practitioner guidance

Issuer education materials state plainly that leveraged ETFs are daily trading tools whose long-horizon behavior departs from the naive multiple (Direxion Shares ETF Trust, 2026), and practitioner commentary discusses both the decay harvest and the practical frictions of shorting hard-to-borrow LETFs (Benzinga, 2024). The borrow-fee reality in particular shapes this report's realized-return accounting.

3 Data

All data infrastructure is documented in the repository (Kumar, 2026); a summary follows. Bloomberg option chains for TSLA and TSLL were enumerated quarterly (2020 to 2026), decoded through OpenFIGI, filtered to contracts that were ever plausible hedges (15-180 days to expiry, 80-130 percent moneyness at some point in life), and their daily closing prices and volumes retrieved through Bloomberg BDH. Spot data come from Bloomberg (TSLA) and yfinance (TSLL). The option datasets are close-only by scope: no bid-ask, Greeks, or intraday data.

Dataset	Rows	Contracts	Date range
TSLA calls	532,258	3,786	2020-01-02 to 2026-06-10
TSLL calls	46,867	689	2022-08-12 to 2026-06-10
TSLA puts	442,076	3,786	2020-01-02 to 2026-06-10
TSLL puts	55,978	1,348	2022-08-10 to 2026-06-10

Table 1: Processed option datasets (daily closing price and volume per contract). The put chains were added by generalizing the acquisition pipeline with an `OPTION_TYPE` switch; the identical numbered steps produce either chain.

TSLL chains are thin (median near 49 contracts per day with volume), which binds LETF-native strategies and motivates the liquidity gates used later. Short-borrow fees for TSLL are approximated by a quarterly step series read from public borrow-fee charts, ranging from 0.7 to 7.0 percent annualized and spiking in TSLA volatility regimes; TSLA itself stays near 0.40 percent. Backtests run on the TSLL option window, 2022-08-12 to 2026-06-10 (960 trading days). The risk-free rate is taken as 5 percent annualized.

4 Empirical Properties of TSLL

Three facts from the exploratory analysis drive the strategy design. First, TSLL delivers its 2x target only at the daily horizon: rolling realized leverage against TSLA stays near 2 (Figure 1), while cumulative performance departs from the naive doubled return exactly as Equation (1) predicts (Figure 2). Second, the decay is large: TSLA's realized volatility routinely exceeds 60 percent annualized, so the e^{-V} drag compounds to tens of percent per year. Third, the mirror image of the harvest is the hazard: a naked short of TSLL spends long stretches in drawdowns beyond 50 percent and reaches roughly 80 percent at the worst (Figure 3). The strategies that follow are attempts to keep the second fact while surviving the third.

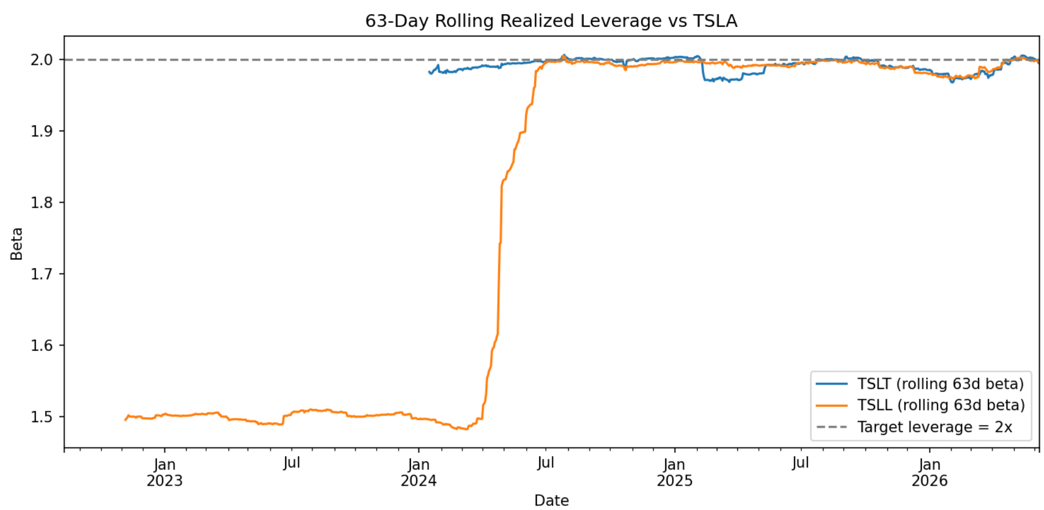

Figure 1: Rolling realized leverage of TSLL on TSLA. The daily 2x target is met; the multi-horizon return is not the naive multiple.

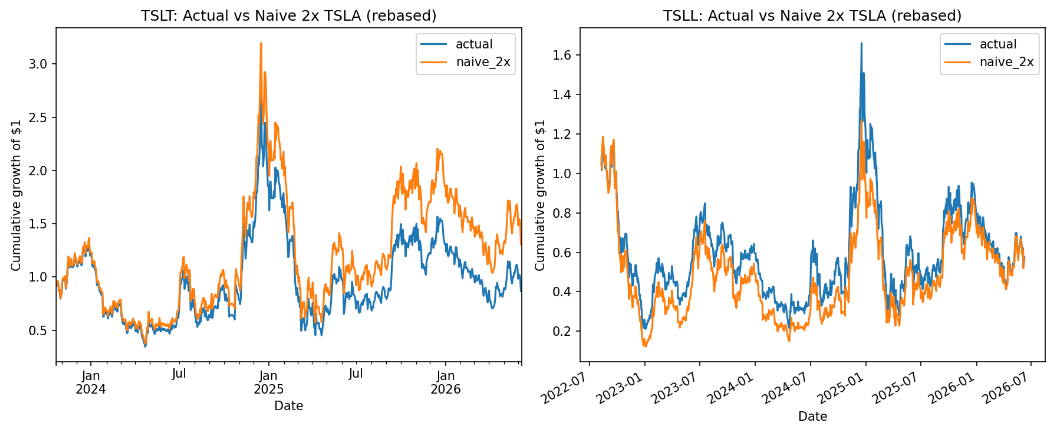

Figure 2: TSLL cumulative performance against the naive 2x compounding of TSLA: the gap is the realized-variance drag of Equation (1).

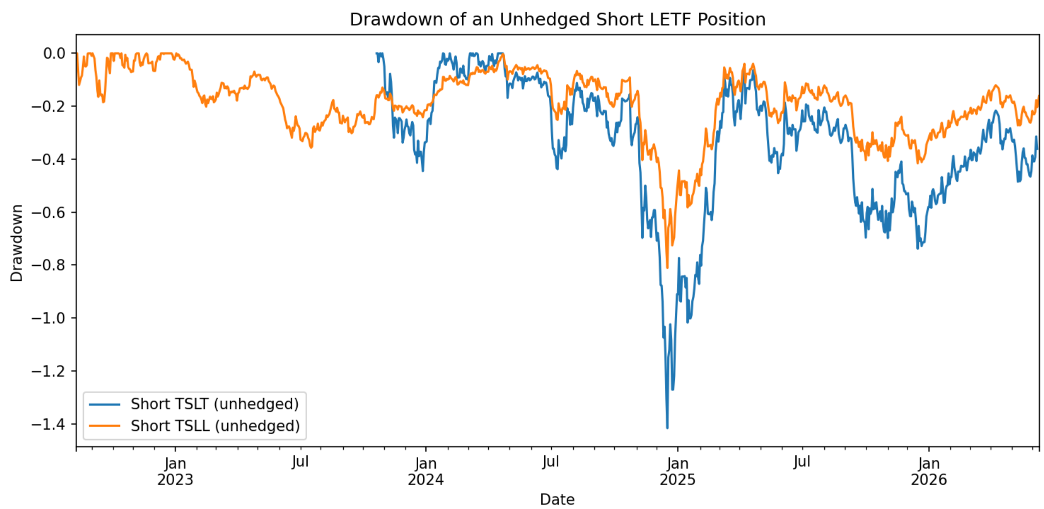

Figure 3: Drawdown profile of a naked short TSLL position: the decay harvest comes with drawdowns that reach roughly 80 percent.

5 Performance Metrics

All strategies are evaluated on daily equity curves normalized to one unit of initial capital: compound annual growth rate (CAGR), annualized volatility, Sharpe ratio (mean over standard deviation of daily returns, annualized by $\sqrt{252}$), maximum drawdown (largest peak-to-trough decline of the equity path), and Calmar ratio (CAGR over the absolute maximum drawdown). Because the subject of the report is loss control, Calmar is treated as the primary risk-adjusted criterion; a strategy change is adopted only if it improves Calmar or adds a guarantee at acceptable Calmar cost.

6 Strategy 1: Naked Short

Short one dollar of TSLL and hold. This collects the full decay harvest and the full hazard. On the common comparison window (2022-11-08 to 2026-06-08, set by Strategy 2's estimation warm-up) the buy-and-hold short returned a CAGR of −0.9 percent with 71.6 percent annualized volatility and a maximum drawdown of −81.1 percent. Rebalancing the short back to constant notional makes things worse, not better: the monthly-rebalanced variant destroyed itself (CAGR −67 percent, drawdown −99.8 percent) because resetting the short after declines locks in losses ahead of rebounds. Over the full TSLL-option window the analytic buy-and-hold curve returns +11.9 percent CAGR with a −74.2 percent maximum drawdown: a real harvest on an untradeable path. This is the benchmark every hedge must beat on drawdown without giving up the harvest entirely.

7 Strategy 2: Delta-Adjusted Arbitrage (Rejected)

7.1 Construction

Hold long h_t dollars of TSLA against each short dollar of TSLL, sizing h_t to neutralize first-order market exposure so that the position isolates the decay term of Equation (1). The hedge ratio is the rolling 63-day realized beta of TSLL on TSLA (covariance over variance), lagged one day, rebalanced daily; a static $h = 2$ sensitivity and financing-cost sweeps complete the specification.

7.2 Headline result and rejection

On the common window the dynamic variant posted the best numbers in the project: CAGR +15.5 percent, annualized volatility 5.4 percent, Sharpe 2.69, maximum drawdown −10.0 percent, Calmar 1.54. It was rejected for two reasons.

First, look-ahead bias in the sense of Pardo (2008) and López de Prado (2018): although the beta estimate is lagged one day, the strategy's performance is concentrated in and dependent on the estimation-window structure itself. Re-running the strategy with strictly expanding walk-forward estimation windows degrades it materially, and remediation attempts (coverage-band triggers, unconditional resizes, walk-forward trigger adaptation, sizing buffers) failed to restore the headline performance without reintroducing the bias. The backtest, as specified, overstates what a live trader could have earned.

Second, financing fragility: the strategy holds roughly 2.8 times gross notional and rebalances daily, while TSLL is frequently hard to borrow (Section 3); realistic short-borrow fees on the TSLL leg plus margin interest on the TSLA leg consume much of the remaining edge.

The rejection had two lasting consequences for the project: no estimated hedge ratio appears anywhere in the final engine (it is model-free by construction), and every later risk statement is pre-registered before the data run and verified against the realized path afterward.

8 Strategy 3: Convexity Protection

8.1 Idea and fixed-spend cost-benefit analysis

Replace the estimated hedge ratio with a contract: short TSLL and hold a long call, so that a rally is capped by the call's payoff with no rebalancing, while decay is still collected in flat and falling markets. A first implementation spends a fixed 2 percent of notional per monthly roll on calls across moneyness buckets, on TSLA (deep, liquid chains) or TSLL (native but thin):

Variant	CAGR	Ann. vol	MaxDD	Fill rate
Benchmark (naked short)	+13.3%	67.2%	−79.6%	n/a
TSLA call, ATM	+14.0%	57.1%	−74.2%	∼100%
TSLA call, 20% OTM	+17.4%	42.9%	−61.2%	∼100%
TSLL call, 20% OTM	+15.8%	45.2%	−64.1%	42%

Table 2: Fixed-spend cost-benefit analysis (2% of notional per monthly roll). TSLA calls dominate at fixed spend; TSLL-native hedges are capacity constrained; and every variant's drawdown remains unacceptable, because a fixed premium budget does not cap the loss.

The lesson of Table 2 is that the hedge must be defined in loss terms, not spend terms.

8.2 The static guaranteed floor

Short a fixed share count n of TSLL forever and buy exactly $n/100$ TSLL call contracts each cycle (share-matched), holding each cycle to the contract's own expiry. If the entry spot is S_0, the strike K, and the premium C per share, then at expiry

$$\Pi(S_T) \;=\; \begin{cases} n\,(S_0 - S_T - C), & S_T \le K, \\ n\,(S_0 - K - C), & S_T > K, \end{cases} \tag{2}$$

so the cycle's worst case is exactly $-n\,(C{+}K{-}S_0)$: above the strike the call payoff cancels the short share for share, and the S-shaped risk of the naked short collapses to the premium paid. The floor is an algebraic identity, not a backtest observation. Contract fill is three-tiered: strict TSLL at the money (60-110 days to expiry); relaxed TSLL (45-120 days, 0.80-1.20 moneyness) with a TSLA delta supplement for the residual exposure; TSLA only when no TSLL contract exists (best effort, no guarantee).

Result over the full window: CAGR +19.42 percent, Sharpe 0.928, maximum drawdown −19.70 percent, Calmar 0.986, with the floor holding on every guaranteed cycle. This is the strongest static result in the project and the base the final engine extends. Its structural limits: the full notional is always deployed (no risk budget and no portfolio drawdown control), the floor binds only at each cycle's expiry, and undeployed capital does not exist as a concept.

9 Strategy 4: The Drawdown-Constrained Engine

Notation. Five symbols govern the engine, chosen deliberately from Latin letters because α, γ, and δ carry standard meanings in this literature (excess return and the option Greeks) that would collide:

Symbol	Meaning	Headline value
ℓ	per-cycle loss budget, fraction of settled equity	8%
d	drawdown limit against the settled high water mark	8.2%
g	gain trigger: cycle ceiling on daily marks	8%
q	cost gate: maximum net floor cost per dollar protected	15%
x	risk-dial parameter: the common value when $\ell = d = g$ are set equal	–

Table 3: Engine notation. All values are a-priori choices registered before the data runs; q in particular is not identifiable from roughly twenty tranches and is never re-fit.

9.1 The per-tranche floor, continuous and exact

A tranche shorts n TSLL shares at entry price S_0 and buys $n/100$ American calls at strike K_c (premium C), all expiring at T. Define the reservation

$$L \;=\; n\,(C + K_c - S_0) \;\geq\; 0, \tag{3}$$

nonnegative by no-arbitrage ($C \geq S_0 - K_c$). By the same case analysis as Equation (2), the worst case at expiry is exactly $-L$, for any strike. The floor also holds before expiry. The tranche's mark-to-market at time t is $n(S_0 - S_t) + n(C(t) - C)$, and an American call satisfies $C(t) \geq (S_t - K_c)^+$ at every t, realizable by exercise. Substituting,

$$\text{tranche}(t) \;\geq\; \begin{cases} n(S_0 - S_t) - nC \;\geq\; n(S_0 - K_c - C) = -L, & S_t \leq K_c, \\ n(S_0 - S_t) + n(S_t - K_c - C) = -L, & S_t > K_c, \end{cases} \tag{4}$$

so every S_t term cancels and

$$\text{tranche}(t) \;\geq\; -L \qquad \text{for all } t \leq T. \tag{5}$$

Recorded closing quotes on thin chains occasionally violate such bounds (stale prices below intrinsic); all marks are therefore confined to their no-arbitrage bands, which only removes prices that were never tradeable.

9.2 The put-financed collar

Extend the tranche by selling $n/100$ American puts at strike $K_p < S_0$, same expiry, collecting premium P_p per share. The settlement profile gains P_p everywhere and caps the downside gain at K_p:

$$\Pi(S_T) \;=\; \begin{cases} n\,(S_0 - K_p - C + P_p), & S_T \leq K_p \quad \text{(gain capped)}, \\ n\,(S_0 - S_T - C + P_p), & K_p < S_T \leq K_c, \\ n\,(S_0 - K_c - C + P_p), & S_T > K_c \quad \text{(floor, improved by } P_p\text{)}, \end{cases} \tag{6}$$

and the reservation improves to

$$L' \;=\; n\,(C + K_c - S_0 - P_p) \;<\; L. \tag{7}$$

The continuous version survives through two further model-free American inequalities (Hull, 2018): parity, $P(K_p, t) \le C(K_p, t) - S_t + K_p$, and the vertical spread bound, $C(K_p, t) - C(K_c, t) \le K_c - K_p$, which combine to $P(K_p, t) \le C(K_c, t) + K_c - S_t$. Chaining through the three-leg mark,

$$\begin{aligned} \mathrm{tranche}(t) &= n(S_0 - S_t) + n\big(C(t) - C\big) + n\big(P_p - P(t)\big) \\ &\ge n(S_0 - S_t) + n\,C(t) - nC + nP_p - n\big(C(t) + K_c - S_t\big) \\ &= -\,n\,(C + K_c - S_0 - P_p) \;=\; -L', \end{aligned} \tag{8}$$

so every S_t term cancels again and the collar tranche is worth at least $-L'$ at every instant. Since $K_c \ge K_p$, the implied ceiling on the put mark never falls below put intrinsic, so the put's no-arbitrage band is always well defined.

The collar was gated by a pre-registered liquidity check before implementation: on 93 percent of days with a qualifying call anchor, a same-expiry TSLL put 10 to 30 percent out of the money also traded, with a median premium of 10.2 percent of spot. Rich crash insurance on a 2x LETF is exactly what Equation (1) would lead one to expect, and selling it does not contradict the observation that *buying* puts is redundant on an already short-delta book: it is the mirror trade. The engine selects the put closest to 20 percent out of the money, falls back to call-only when no put qualifies, and treats a net cost below 0.5 percent of spot as a stale quote.

9.3 Settled accounting and the aggregate floor

Let E_0 be equity at the last settled checkpoint, a date on which no tranche is open (an expiry settlement, a ceiling settlement, or any all-cash day), and let ATH^* be the maximum over settled checkpoint equities. Neither quantity ever reads an unrealized mark. Tranches may open on any day subject to the cumulative cycle budget

$$\sum_{i\,\in\,\mathrm{cycle}} L_i \;\le\; \min\big(\ell\,E_0,\; E_0 - (1-d)\,\mathrm{ATH}^*\big), \tag{9}$$

with reservations accumulating from the last checkpoint and resetting only at the next (releasing them on mid-cycle closes would let one cycle's cumulative risk exceed ℓ). Every tranche in a cycle expires on or before the cycle's anchor expiry, so all guarantees co-fire at one settlement date. Combining (9) with the continuous per-tranche floors gives, at every instant,

$$E(t) \;\ge\; \max\big((1-\ell)\,E_0,\; (1-d)\,\mathrm{ATH}^*\big). \tag{10}$$

Keying the budget to settled values is what removes procyclicality: unrealized rallies can neither enlarge the budget on the way up nor manufacture a high-water-mark lockout on the way down. Earlier designs that keyed the budget to daily marks deployed at peaks and froze at troughs; settled accounting eliminated the failure mode.

9.4 Dynamic deleveraging: sizing from the price of the floor

The budget constraint (9) induces the engine's defining behavior: position size responds to the price of protection and the distance to the drawdown limit, not to any forecast. Let

$$c \;=\; \frac{C + K_c - S_0 - P_p}{S_0} \tag{11}$$

be the net floor cost per dollar of short notional (the collar term P_p is zero for a call-only tranche). Deploying the fraction $\pi = nS_0/E_0$ of settled equity risks exactly $\pi\, c\, E_0$ in the worst case, so for a single tranche the budget (9) solves to the deployment fraction

$$\pi \;=\; \min\Bigg(1,\quad \underbrace{\frac{\ell}{c}}_{\text{loss budget}}\quad,\quad \underbrace{\frac{E_0 - (1-d)\,\mathrm{ATH}^*}{c\,E_0}}_{\text{drawdown headroom}}\Bigg). \tag{12}$$

Both terms deleverage the position continuously and automatically. When protection is cheap (c small) the engine deploys fully; when it is expensive, the loss-budget term shrinks size in exact proportion, so the worst case never exceeds ℓ of settled equity. When equity sits near the drawdown limit, the headroom term shrinks size toward zero and the engine parks entirely, with undeployed capital earning the risk-free rate; because ATH^* is a settled quantity, headroom then rebuilds deterministically and the engine provably re-enters. Multiple tranches generalize (12) through the cumulative reservation in (9): each new tranche is sized so that its own L_i consumes only the budget the cycle has not yet reserved, capped further by available cash collateral (nS_0 at most the undeployed cash) and by an anti-dust threshold of one percent of settled equity per tranche.

9.5 The master equation and the risk dial

Equity peaks form when TSLL falls (the short profits while the long call loses at most its premium). A cycle ceiling closes every tranche once marks reach $E_0(1+g)$, converting any peak into a settled checkpoint before it exceeds the previous one by more than the factor $(1+g)$. Together with (10),

$$\mathrm{MaxDD} \;\le\; \frac{d+g}{1+g}. \tag{13}$$

The headline configuration ($\ell = 8\%$, $d = 8.2\%$, $g = 8\%$) gives a bound of exactly 15.0 percent.

The risk-dial parameter x. Because the three risk parameters play complementary roles (how much one cycle may lose, how far equity may sit below its high water mark, and how quickly gains are locked in), it is natural to move them together. Define the *risk-dial parameter* x as their common value when they are set equal,

$$x \;=\; \ell \;=\; d \;=\; g, \qquad \text{giving} \qquad \mathrm{MaxDD} \;\le\; \frac{2x}{1+x} \tag{14}$$

by substitution into Equation (13). The dial reduces the engine to a single risk-appetite knob: $x = 2\%$ guarantees a 3.9 percent drawdown cap (a cash-plus profile), $x = 8\%$ guarantees 14.8

percent and essentially reproduces the headline configuration (which sets $d = 8.2\%$ rather than 8% only to land the bound on exactly 15.0 percent), and $x = 15\%$ guarantees 26.1 percent. The investor selects the guarantee; the backtest then reports what that guarantee historically earned (Table 6). Every appearance of x in this report refers to Equation (14). The bound is exact under continuous monitoring; with daily closes a single day's move can overshoot a trigger line, a granularity caveat shared by every end-of-day backtest in this report.

9.6 The favorability gate

A tranche only breaks even if TSLL falls by more than its net floor cost before settlement. Deployment is therefore gated:

$$\frac{C + K_c - S_0 - P_p}{S_0} \;\le\; q = 15\%, \tag{15}$$

never pay more than fifteen cents per dollar of protected notional for one cycle's floor. Parking when the gate is closed cannot hurt the floor. In an ablation, removing the gate cut the compound growth rate from 14.0 to 4.9 percent: the ungated engine repeatedly bought the most expensive floors in the sample immediately after crashes.

9.7 Borrow costs and the leak-adjusted guarantee

The daily borrow fee is $b_t/252 \cdot nS_t$ on the market value of the borrowed shares. Unlike option premiums this cost is not bounded by a no-arbitrage identity (the fee base grows if TSLL rallies), so the floor acquires a tracked leak. Writing $B(t)$ for cumulative borrow paid, $B_{\text{cyc}}(t)$ for borrow accrued since the last checkpoint, and $B(t_{\text{ATH}^*})$ for cumulative borrow when ATH* was last set,

$$E(t) \;\ge\; \max\Big((1-\ell)E_0 - B_{\text{cyc}}(t),\;\; (1-d)\,\text{ATH}^* - \big[B(t) - B(t_{\text{ATH}^*})\big]\Big), \tag{16}$$

and the drawdown bound gains $\Delta B/(\text{ATH}^*(1+g))$. The two leak terms differ because settled equity absorbs realized borrow at each checkpoint while the ATH* branch accumulates across cycles; a naive per-cycle adjustment produces phantom violations on all-cash recovery days. Both inequalities remain theorems because B is measured, not assumed. Borrow flows into equity daily, so sizing, ceiling, and checkpoint decisions all see it: the realized results below are not an after-the-fact subtraction.

10 Results

All bounds were computed from parameters and registered before the data runs. The window is 2022-08-12 to 2026-06-10 (960 trading days); the engine is TSLL-only and parks at the risk-free rate when no contract qualifies.

10.1 Headline comparison

Strategy	CAGR	Sharpe	MaxDD	Calmar
Collar, frictionless	+15.69%	0.899	−12.29%	1.277
Collar with borrow (realized)	+10.41%	0.680	−15.21%	0.684
Call-only, frictionless	+14.00%	0.841	−11.29%	1.240
Call-only with borrow (realized)	+13.42%	0.706	−15.44%	0.869
Static guaranteed floor (Section 8.2)	+19.42%	0.928	−19.70%	0.986
Naked short TSLL (benchmark)	+13.30%	–	−79.60%	0.167

Table 4: Full-period results at the headline configuration ($\ell = 8\%$, $d = 8.2\%$, $g = 8\%$, $q = 15\%$; guaranteed bound 15.0%). The collar ran 20 tranches (18 with a put leg) across 19 settlements, 6 ceiling-triggered; mean deployment 32.2 percent of equity versus 19.2 percent call-only.

Three observations. Both frictionless engines beat the static floor on Calmar while carrying a provable cap the static strategy lacks. The ranking flips under realized costs: the collar's advantage is higher deployment, and higher short notional pays proportionally more borrow in exactly the regimes where borrow spikes; of its 4.78 percentage points per year of drag, roughly 1.4 points are direct fees and the remainder is path divergence (borrow-lowered equity misses ceiling triggers the frictionless path caught). And every variant's realized drawdown respects its leak-adjusted bound: collar −15.21 percent against −17.54, call-only −15.44 against −16.23.

10.2 Verification

Across the four full-period variants, both disjoint holdout halves per variant, and seven risk-dial configurations, the daily equity path produced zero violations of the leak-adjusted analytic floor (16) over 960 days each; the frictionless variants also produced zero violations of the unadjusted floor (10). The maximum observed leak was 4.26 percent of initial notional (collar, cumulative since the high water mark) and 1.69 percent within a single cycle (call-only). The production implementation re-runs these checks on every execution and aborts if any fails.

10.3 Holdout

Variant	Period	CAGR	Sharpe	MaxDD	Calmar
Collar	H1	+33.16%	1.273	−12.29%	2.698
Collar	H2	−2.08%	−0.038	−11.44%	−0.182
Call-only	H1	+30.84%	1.255	−11.29%	2.731
Call-only	H2	+0.49%	0.101	−11.78%	0.042

Table 5: Disjoint halves (split 2024-07-12), fixed pre-registered configuration, frictionless; zero floor violations and within bound in every cell. H1 contains the late-2022 TSLA drawdown; H2 is benign.

The risk guarantee generalizes; the return does not generalize evenly. All the return sits in the crash half. For a hedge this is the intended profile: it pays when the protected book is in distress and approximately preserves capital otherwise, with the cap in force throughout. It is not an all-weather return stream and is not presented as one.

10.4 The risk dial

x	Bound $2x/(1+x)$	CAGR	Sharpe	MaxDD	Calmar
2%	3.9%	+6.23%	0.995	−3.72%	1.675
4%	7.7%	+9.43%	0.836	−5.01%	1.883
6%	11.3%	+6.18%	0.454	−7.57%	0.816
8%	14.8%	+15.75%	0.903	−12.09%	1.303
10%	18.2%	+17.57%	0.892	−13.73%	1.280
12%	21.4%	+13.50%	0.683	−16.08%	0.840
15%	26.1%	+13.32%	0.637	−18.65%	0.714

Table 6: Risk dial: each row sets $x = \ell = d = g$ (Table 3) with the cost gate fixed at $q = 15\%$; collar engine, frictionless. Every configuration stayed inside its own bound with zero floor violations; realized drawdown runs at 75 to 96 percent of the guaranteed envelope, so the bound is informative rather than vacuous. The dip at $x = 6\%$ is a small-sample artifact.

Risk stops paying beyond roughly $x = 10\%$: further envelope buys more drawdown and less growth. The conservative end is a product in its own right; at $x = 2\%$ the engine earns about two points over the risk-free rate under a 3.9 percent guaranteed cap.

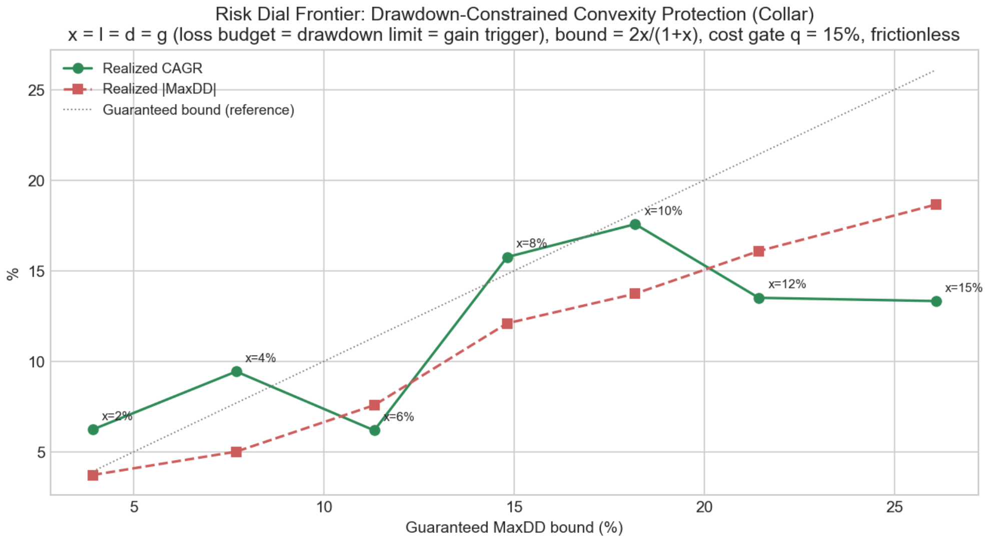

Figure 4: Risk dial frontier: guaranteed envelope against realized growth and drawdown. Each point is one configuration $x = \ell = d = g$ with bound $2x/(1+x)$; cost gate fixed at $q = 15\%$; collar engine, frictionless.

10.5 Strategy lineage

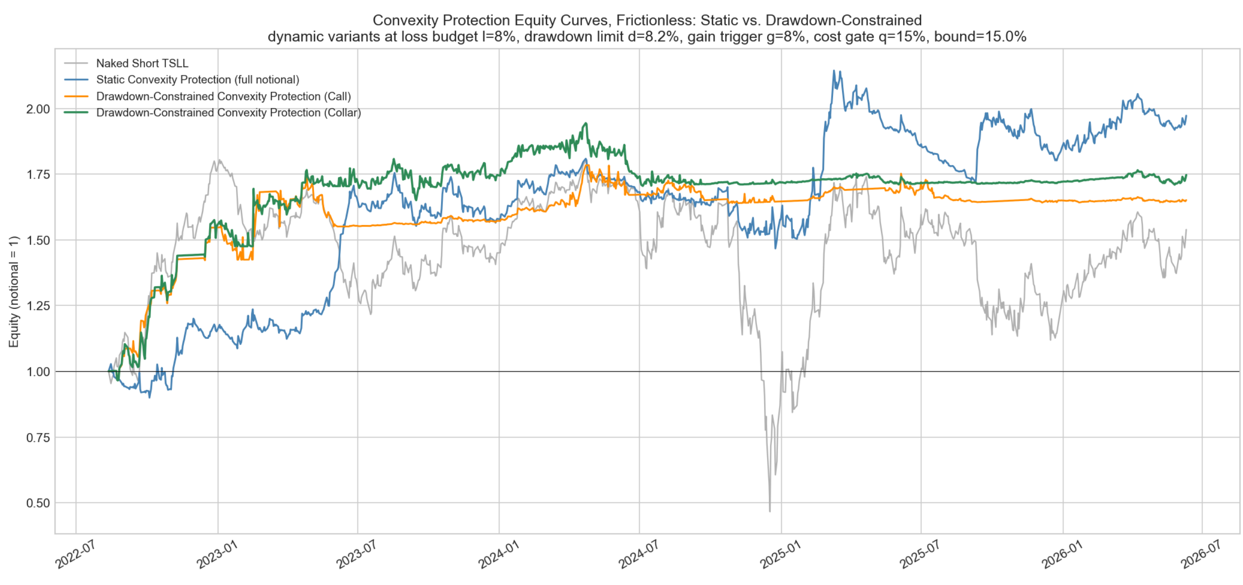

Figure 5: The four surviving constructions on one equity axis, frictionless, identical window and data: naked short benchmark, static convexity protection at full notional, and the two drawdown-constrained variants at the headline configuration. Strategy 2 is excluded as invalid (Section 7). The static reproduction in this figure matches Section 8.2 exactly (+19.42%, −19.70%).

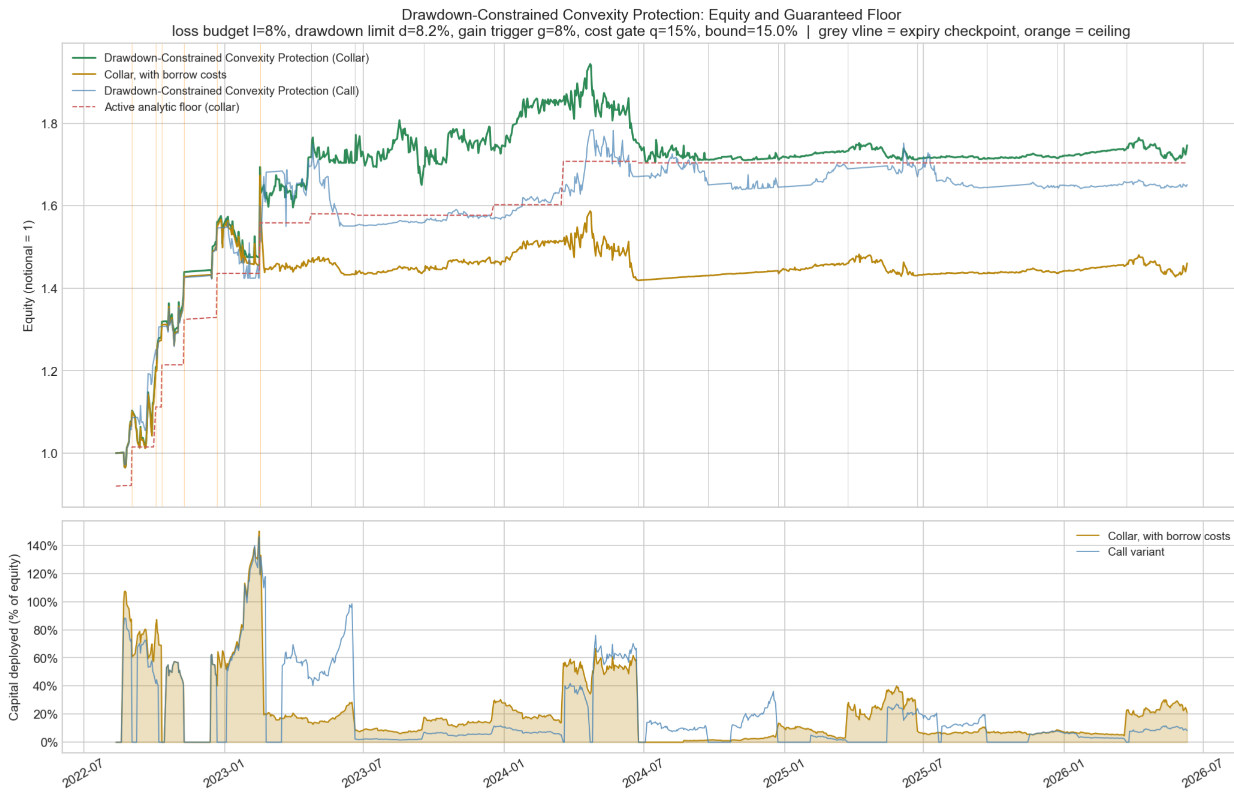

Figure 6: Engine equity against the active analytic floor with settlement checkpoints, and capital deployment over time. Headline configuration ($\ell = 8\%$, $d = 8.2\%$, $g = 8\%$, $q = 15\%$; bound 15.0%).

11 Limitations

Everything transacts at the day's last traded price of volume-verified contracts. Bid-ask spreads are the principal unmodeled cost, are adverse on both option legs, and cannot be estimated from this dataset, which was scoped to closing prices by design. Borrow rates are quarterly chart-read approximations. Decisions execute at the same close they observe. Position size is not capped against actual contract volume, which binds at real capital on chains with median daily volume near 30 to 50 contracts. TSLL distributions to the share lender are not modeled. Returns are concentrated in the crash regime; the gate threshold and put-strike rule are a-priori choices not identifiable from roughly twenty tranches; and the sample is a single 3.8-year history containing one deep crash. The bound column of every table is a theorem given the stated assumptions; the return columns are one observation.

12 Conclusion

The project set out to harvest leveraged ETF decay without the naked short's catastrophic path, and its trajectory is a study in what survives scrutiny. The best-looking strategy (delta-adjusted arbitrage, Sharpe 2.69) failed the information-structure test and was rejected. The best static strategy (the share-matched guaranteed floor, Calmar 0.986) survived because its risk claim is an algebraic identity. The final engine generalizes that identity to a portfolio: an exact reservation per tranche, settled accounting, a closed-form drawdown bound with a one-parameter risk dial, a liquidity-gated collar that finances the floor, and realized borrow accounting with an explicitly tracked leak. Its risk claims were registered before the data runs and never violated on the realized path. On frictionless terms it delivers the project's best risk-adjusted results (Calmar 1.277); on realized terms the call-only variant is preferred (Calmar 0.869 at a 15.4 percent drawdown against the benchmark's 79.6); and in both cases the drawdown cap is a property of the construction rather than a fortunate sample.

References

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

Benzinga (2024). Leveraged ETFs in focus: Benzinga hosts virtual boot camp on trading strategies. https://www.benzinga.com/news/events/24/12/42372063/leveraged-etfs-in-focus-benzinga-hosts-virtual-boot-camp-on-trading-strategies. Accessed June 2026.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

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

Direxion Shares ETF Trust (2026). Understanding leveraged and inverse exchange traded funds. https://www.direxion.com/education/understanding-leveraged-exchange-traded-funds. Accessed June 2026.

Hull, J. C. (2018). *Options, Futures, and Other Derivatives*. Pearson, 10th edition.

Kumar, D. (2026). Dynamic hedging of a leveraged ETF position: TSLA / TSLL. https://github.com/daksh1069/dynamic_hedging_letf. Source code and data pipeline.

López de Prado, M. (2018). *Advances in Financial Machine Learning*. John Wiley & Sons.

Pardo, R. (2008). *The Evaluation and Optimization of Trading Strategies*. John Wiley & Sons, 2nd edition.