Rational mNAV

Valuing the Wrapper in Bitcoin Terms Against the Self-Custody Hurdle

A Bitcoin-Denominated Monte Carlo Framework for Bitcoin-Treasury Equity: Strategy (MSTR) versus Self-Custody

@rationalmnav

August 4, 2026
First published June 26, 2026

The author welcomes feedback on ways to improve this document and revises it as new learnings emerge and strategies evolve.

Abstract

We develop a valuation framework for a Bitcoin-treasury company—one whose enterprise is the accumulation of Bitcoin (BTC) funded by issuing equity, preferred stock, and convertible debt—taking the perspective of an investor whose unit of account is Bitcoin itself. The central question is not whether the equity rises in dollars but whether owning it delivers *more Bitcoin per dollar committed* than simply holding self-custodied coins. We answer it with a Monte Carlo engine that (i) simulates the BTC price under a *selectable* process—a GBM whose median compounds at a flat rate, or one of three diminishing-returns trends whose CAGR decays from a high start toward a floor—each anchored to a fair price de-lagged from the 200-week moving average, with today's spot converging to that trend as a decaying dislocation, and optionally overlaid with an auto-calibrated stochastic-volatility process, (ii) simulates the short policy rate (SOFR) under a selectable process (Vasicek, Cox–Ingersoll–Ross, random walk, or jump-diffusion) with an optional fiscal-dominance cap on the policy rate and optional joint fiscal-stress shocks that re-rate Bitcoin in the same event, (iii) sets the firm's leverage target countercyclically off a *feasibly estimated* fair line—high near bear bottoms, low at manias, sized for constant stress amplification and capped by a fractional-Kelly test on realized carry, so the policy assumes neither cycle clairvoyance nor knowledge of the growth regime—and propagates an explicit capital-structure engine in which the floating preferred (STRC) dividend is set endogenously off the simulated short rate (SOFR), the firm's leverage, the USD reserve runway, and a migrating credit rating, the STRC price is then *derived* as the ratio of that coupon to the yield the market requires—so its discount to par, and hence the fraction of the time its at-the-market program can issue at or above par, are outputs of the Bitcoin and reserve path rather than assumptions; convertible notes resolve contingently at their put dates under an assumed-diluted share count, a proactive accretive equity program issues into strength at a pace that scales with the premium, and residual financing needs are met through a strict waterfall (free cash flow, cash reserve, at-the-market equity, at-the-market preferred, and—only as a last resort—BTC sales), and (iv) prices each share at exit as its gross Bitcoin backing per share multiplied by the premium the market grants—a premium built from a net-backing ratio (Bitcoin value less the senior claims ranking ahead of the common, per unit of gross coin value) plus the capitalized value of the firm's forward Bitcoin-per-share yield in excess of its structural cost of leverage, so that a wrapper compounding coins per share faster than its cost of capital trades above backing and one that cannot trades below it. The premium the equity trades at is computed *endogenously* along every path: the market inside the model grants $m_t = \rho_t + (\mathrm{BY}_t^{\mathrm{fwd}} - h_t)D$—net backing plus D years of the forward Bitcoin-per-share yield in excess of the structural cost of leverage—with the forward yield driven by a *momentum kernel* calibrated on the full 2020–2026 record, in which the premium tracks trailing twelve-month Bitcoin momentum far better than the price level, expanding convexly in uptrends and flooring near net backing in downtrends. Finite-supply brakes (execution slippage on the firm's own buys, a concentration brake that shrinks the sourceable purchase flow with the remaining float, and an adaptive preferred-carry test) cause the terminal Bitcoin holding to emerge as a derived arc rather than an input—calibrated so that accumulating more than a quarter of the network within three decades is unlikely. The investor's rule is to prefer the equity only when its expected *Bitcoin-denominated* return clears a counterparty/wrapper-risk hurdle relative to self-custody. From this we derive a closed expression for fair value—the share price at which the median Bitcoin-denominated return equals the hurdle, equivalently the price at which the probability of beating self-custody is one-half—and we translate it into a *rational mNAV*—and, because the brakes make long holds pay a real concentration cost, we derive the *optimal holding horizon* at which the marginal year stops clearing the hurdle and self-custody takes over—distinguishing the model's premium (what the trend-extrapolating market pays) from the investor's *personal* rational mNAV (the fair entry under the investor's own outlook and hurdle). The rational premium is thus not a single number but a function of beliefs. All figures are conditional outputs of the stated assumptions and are not forecasts.

1 Motivation and numeraire

A Bitcoin-treasury company presents a deceptively simple question. If one wishes to invest in Bitcoin, is it better to buy the coins and hold the keys, or to buy a share of a company that *owns* coins on its own balance sheet and finances additional purchases with leverage? A treasury company is not a custodian: it owns the Bitcoin outright, and the shareholder owns equity whose value derives from those coins but cannot be redeemed for them. There is no right of withdrawal—one cannot exchange shares back into the underlying coin—so the shareholder's claim is mediated entirely by the market price of the stock and by the corporate structure sitting between the coins and the equity. The two routes are therefore not equivalent: the equity wrapper layers on a premium or discount to net asset value, a senior preferred and convertible stack, dilution at management's discretion, and assorted counterparty, governance, and regulatory risks; in exchange it offers leveraged accumulation that, when it works, grows the Bitcoin backing *per share* faster than an individual could on their own.

The cleanest way to adjudicate is to change the unit of account. We denominate everything in Bitcoin. Holding one's own coin is then the risk-free benchmark: one Bitcoin remains one Bitcoin, a Bitcoin-denominated return of zero. The equity is worth buying only if its Bitcoin-denominated return is positive after a hurdle that compensates for the wrapper's incremental risk. This reframing removes the dollar price of Bitcoin as a confound (it cancels from both sides) and isolates the only thing that matters: whether the corporate structure is expected to deliver more coin per share than one started with.

2 The decision rule

Let B_t be the USD BTC price, H_t the coins held by the firm, N_t the diluted share count, and $b_t \equiv H_t/N_t$ the *Bitcoin backing per share.* Let m_t denote the market premium ("mNAV"), defined throughout this paper as the ratio of equity market capitalization to the value of the firm's Bitcoin,

$$m_t = \frac{N_t P_t}{H_t B_t}, \tag{1}$$

where P_t is the share price. An investor who buys one share at $t = 0$ pays, in Bitcoin terms, $P_0/B_0 = m_0\, b_0$ coins. At horizon τ the share is worth $b_T\, m_T$ coins. The Bitcoin-denominated gross multiple of the position is therefore

$$G \;=\; \frac{b_T\, m_T}{m_0\, b_0}, \tag{2}$$

against a benchmark multiple of 1 for self-custody. Writing r_p for the annual counterparty/wrapper-risk hurdle, the rule is to prefer the equity whenever

$$G^{1/\tau} - 1 \;>\; r_p. \tag{3}$$

Equivalently, the hurdle compounds over the full holding period: the equity is preferred whenever its Bitcoin-denominated exit value exceeds the entry cost grown at r_p for τ years,

$$b_T\, m_T \;>\; m_0\, b_0\, (1 + r_p)^{\tau}. \tag{4}$$

The hurdle is therefore an *annual* rate, not a one-time threshold: a 3.5% hurdle over a ten-year hold demands a Bitcoin-denominated gain of $(1.035)^{10} - 1 \approx 41\%$ across the horizon, not 3.5%.

Equation (2) cleanly separates the two engines of value: the growth in backing per share b_T/b_0 (the accretion "flywheel"), and the change in the premium m_T/m_0 (the market's willingness to keep paying above net asset value). As Section 5.3 makes precise, the premium term is computed *endogenously* along each path: the same Bitcoin outlook that drives the price beneath b_T also drives the realized momentum and demonstrated yield the market capitalizes into m_T, so both factors move together on every path; everything that follows is machinery for simulating their joint distribution.

The sign of the hurdle. A positive r_p is the usual case: the wrapper carries incremental counterparty, governance, dilution, and regulatory risk, and a rational holder demands compensation for bearing it. But r_p is a personal parameter, and for some investors it is legitimately *negative*. Consider someone who wants Bitcoin exposure yet cannot or will not self-custody—wary of their own key management, bound by institutional or fiduciary constraints, or simply unwilling to be their own bank—and who also wishes to avoid the recurring management fee of a spot ETF, a drag that compounds against the holder every year the position is held. For that investor the equity delivers Bitcoin exposure with the prospect of a *growing* Bitcoin-denominated balance (the accretion flywheel), no annual fee, and a custodian they trust more than themselves. The convenience, the avoided ETF fee, and the trust carry real value—the mirror image of the recurring fee a spot ETF charges—so such an investor demands no positive premium over self-custody to hold the wrapper, and may rationally assign it a *negative* hurdle. That lowers the bar in (3) and, through the fair-value relation (32), *raises* the rational mNAV—the wrapper can be worth a premium to its own Bitcoin backing for someone who values what it provides beyond the coins themselves.

2.1 Various measures of mNAV

The numerator and denominator of the premium are a modeling choice, and three conventions are in common use among analysts. They are not interchangeable: each measures a distinct quantity and answers a distinct question. We state the three explicitly and justify our selection on analytical rather than presentational grounds.

Let P denote the MSTR common share price, N the assumed-diluted common share count, N_{xc} the *ex-convert* share count (basic plus options and RSUs, before any convertible conversion),

H the quantity of Bitcoin held, B the Bitcoin spot price, D the outstanding convertible principal, and L the aggregate preferred liquidation preference. The three measures are

$$m^{\mathrm{eq}} = \frac{PN}{HB}, \qquad m^{\mathrm{EV}} = \frac{PN_{\mathrm{xc}} + D + L - R}{HB}, \qquad m^{\mathrm{net}} = \frac{PN_{\mathrm{xc}}}{HB - D - L + R}. \tag{5}$$

Here R is the firm's USD reserve. The two capital-structure measures net it against the senior stack—subtracting it from the aggregated claims in m^{EV} (the enterprise-value convention, in which cash reduces enterprise value) and adding it back to the asset base in m^{net} (the reserve is cash already set aside to service the senior claims being deducted, so a coherent floor must count it). The equity measure m^{eq} nets nothing and includes no non-Bitcoin asset, so the reserve does not enter it. The two capital-structure measures use the ex-convert count N_{xc} in their equity term rather than the assumed-diluted N: each already adds the convertible principal D as a senior claim, so counting the conversion shares in the equity term as well would charge the convertibles twice (developed below). The equity measure, which adds no senior claim, uses the assumed-diluted N. We adopt the first, m^{eq}—the market capitalization of the common equity over the gross value of the Bitcoin holdings.

A terminological caveat is in order. Strictly, "mNAV" is a misnomer for the quantity we adopt. A *net* asset value would subtract the firm's liabilities and add its non-Bitcoin assets—here, principally the USD reserve—and m^{eq} does neither. It is simply the ratio of the (assumed fully diluted) common-equity market capitalization to the gross market value of the Bitcoin holdings, with nothing netted out and no other assets included. A more accurate label would be a *multiple of asset value*—"mAV"—since the denominator is a gross asset value rather than a net one. We nonetheless retain "mNAV" throughout this paper, because the term is in common and well-understood use in discussions of Bitcoin-treasury equities; coining a private acronym would cost more in legibility than the added precision would return. The reader should simply keep in mind that the "N" is, in our usage, a convention rather than a literal description.

The justification is that the decision this framework addresses is posed from the standpoint of the residual (common) claimant, whose objective is the Bitcoin-denominated total return on the common over a multi-year horizon. The governing state variable is the Bitcoin backing per common share, H/N, and its evolution under the financing program. Of the three ratios, m^{eq} is the only one whose denominator is an observable, assumption-free quantity (HB) and whose numerator is exactly the claim being valued; it isolates the residual equity claim without imputing a valuation to the senior instruments, and it keeps the per-share backing legible. Accretive common issuance—sale of equity at a premium to $m^{\mathrm{eq}} = 1$ with proceeds deployed into coins—raises H/N, and it is the long-run trajectory of m^{eq}, rather than the accretion or dilution attributable to any individual issuance, that governs the realized return.

Enterprise value, m^{EV}, aggregates the senior instruments into the numerator and thereby values the entire capital structure relative to its assets. This is a coherent measure of the richness of the whole structure, but it does not isolate the residual claim: it superimposes the lenders' and preferred holders' positions on the common holders', and a single aggregate ratio cannot separate the return accruing to the residual claimant—the object of this analysis—from the cost of the senior financing. It answers a different question than the one posed.

Net asset value, m^{net}, deducts the senior claims from the asset base at face value—net of the USD reserve already held against them. This embeds an implicit wind-down assumption, because subtracting $D + L$ at par treats those claims as obligations due and payable in full at the valuation date. For this issuer they are not. The preferred is perpetual: it has no maturity, imposes no date on which coins must be liquidated to redeem principal, and burdens the common solely through its dividend stream—an ongoing carry rather than a principal repayment. The convertibles carry no covenant capable of compelling a sale of Bitcoin; each tranche resolves at its scheduled put either by conversion into equity or by refinancing from ordinary funding. Absent any claim that can force liquidation of the collateral, deducting the full face value of the senior stack prices the equity on wind-down terms even when continued operation is the modal outcome.

This is not to dismiss m^{net}, which is a legitimate downside reference—an estimate of the floor the common would approach were capital-market access to be withdrawn. The objection is to its use as the central estimate. If an investor assigns probability π to a forced wind-down, the liquidation outcome warrants weight π and the going concern weight $1-\pi$; adopting m^{net} as the headline valuation implicitly sets $\pi = 1$.

The decisive argument for excluding the senior stack from the denominator, and charging it through the cash flows instead, is that the alternative double-counts. Consider the convertibles under an assumed-diluted share count, which presupposes their eventual conversion to equity: that count has already recognized them as prospective dilution of the common. Deducting the convertible principal from the asset base in addition would charge the same claim twice— once as future shares and once as a cash repayment. A convertible resolves as exactly one of the two, conversion (in the diluted count N) or repayment/refinancing (in the ex-convert count N_{xc}), so a framework that dilutes for conversion must not also subtract the principal—which is precisely why m^{EV} and m^{net} use N_{xc} while m^{eq} uses N. The identical point applies to the perpetual preferred, whose economic burden is the dividend stream the simulation already charges period by period; subtracting its notional in addition would charge for it twice.

The issuer's own conventions (July 2026). Since this section was first written the issuer has adopted m^{net} as its *reported* headline metric (announcing that references to mNAV computed before July 23, 2026 are not comparable), alongside a companion per-share measure, *Net BPS*: net Bitcoin—coins less out-of-the-money debt and preferred notional net of the USD reserve, divided by the Bitcoin price—per *FDSO*, a share count comprising basic shares, options and RSUs, plus only the *in-the-money* convertible instruments. The convention is internally consistent: an instrument contributes either its conversion shares (when in the money) or its notional claim (when out), never both—the same single-counting principle this section develops, applied instrument by instrument at the prevailing share price rather than unconditionally. Our N_{xc} coincides with FDSO whenever every convertible and the convertible preferred (STRK, conversion price \$1,000) is out of the money, which is the case at the calibration snapshot; the decision tool implements the full price-dependent classification, including the exclusion of STRK's conversion shares whenever its liquidation preference is netted. The relationship between the two per-share measures is the wedge identity $\text{netBPS} = \text{BPS} - K/(B\,N)$ with K the net senior claims: a fixed-dollar term over a compounding denominator, so the two measures converge as the thesis plays out and diverge in drawdowns—which is precisely why this framework retains gross BPS as the objective and treats the net family as the credit constraint. At the snapshot the tool's Net BPS reproduces the issuer's reported figure to within 0.05%.

The senior claims are therefore not disregarded but charged in the form in which they actually fall due. Preferred dividends, convertible coupons, and any contingent redemptions are deducted as ongoing costs that reduce the asset backing per share, period by period; and the fair premium itself carries a net-backing term equal to Bitcoin value less the senior claims ranking ahead of the common plus the cash reserve, normalized by gross coin value, which marks the premium down as that coverage thins—so that a balance sheet on which the coins only narrowly cover the senior stack is valued more conservatively than one with ample headroom. The headline premium thus remains a direct comparison of the market valuation of the common equity to the Bitcoin it ultimately stands behind, with the senior stack borne where it economically falls: in the cash flows, over the holding period.

3 Exogenous state dynamics

Two exogenous drivers are simulated on a quarterly grid: the Bitcoin price and the short policy rate (SOFR). Each admits several processes so that the user can express a range of worldviews rather than a single house view.

3.1 Bitcoin price

Every price model compounds its trend forward from *today's fair price* F_0, derived from the 200-week moving average, while the observed price starts at spot and converges to the trend as a decaying dislocation. A single day's spot embeds whatever short-term dislocation the market is carrying; anchoring a multi-decade trend to it would bake that noise into the entire forward arc. The 200-week average summarizes four years of history instead—but it is a *lagged* statistic, and the model de-lags it exactly. For a trend growing at rate g, the trailing T-window average of that trend equals today's trend value divided by a known factor, so

$$F_0 \;=\; \mathrm{WMA}_{200} \times \frac{\bar{g}T}{1-e^{-\bar{g}T}}, \qquad T = 200 \text{ weeks} \approx 3.83 \text{ yr}, \tag{6}$$

where $\bar{g}$ is the *realized* growth over that same trailing window—an observable, not a belief. One shared F_0 serves all four models; letting each model imply its own fair price from its own growth assumption would let a bullish belief buy both a higher trend *and* a larger catch-up, double-charging the same view. The spot–fair gap $d_0 = \ln(B_0/F_0)$ then decays deterministically,

$$B_t \;=\; M_t\, e^{\,d_0 e^{-\kappa_d t}}, \tag{7}$$

where M_t is the unchanged model process started from F_0 and $\kappa_d = \ln 2/h_d$ with half-life $h_d = 1.25$ years, calibrated to the trough-to-trend reclaim times of the 2015, 2018, and 2022 bears. The factor is applied identically to every model, *outside* its stochastic machinery, and it cuts both ways: spot below fair produces an early upward convergence after which the selected model's trend governs; spot above fair (a mania) produces an early give-back before trend resumes. At the entry snapshot, $\bar{g} \approx 27\%$ gives a de-lag factor of 1.61, $F_0 \approx \$100{,}900$ against \$59,000 spot, and $d_0 \approx -0.54$: the model regards the entry price as $\approx 41\%$ below fair. The entry *balance sheet* remains marked at spot throughout—the coins are worth what they are worth today—and disabling the anchor recovers the from-spot regime exactly ($F_0 = B_0$, $d_0 = 0$). The process selected here is *your* outlook—it drives the Bitcoin price and hence both your Bitcoin-denominated return and, through the realized momentum and yield it produces, the premium the equity trades at along each path (Section 5.3).

Common structure. Each model specifies a deterministic trend G_t with $G_0 = F_0$, around which a mean-reverting residual $x_t \equiv \ln(B_t/G_t)$ fluctuates,

$$dx_t = -\kappa_B\, x_t\, dt + \nu(t)\, dW_t^B, \qquad \nu(t) = \sigma(t)\sqrt{2\kappa_B}, \qquad x_0 = 0, \tag{8}$$

so that the core process $G_t\, e^{x_t}$ starts at F_0 and the observed price $B_t = G_t\, e^{x_t} e^{\,d_0 e^{-\kappa_d t}}$ starts exactly at spot. The residual's relative volatility decays as the network matures, $\sigma(t) = \sigma_0\,\big(t_0/(t_0{+}t)\big)$ with $\sigma_0 = 0.50$ and Bitcoin's current age $t_0 \approx 17.5\,\mathrm{yr}$, shared identically by all four models (Table 1).

GBM (flat 10%). The neutral anchor, recast to mimic a mature equity index. Rather than fix a drift coefficient and let the realized return fall out, we *pin the median path* to compound at a fixed rate g (default 10%/yr, the S&P 500's long-run nominal total return) at *every* horizon:

$$d\ln B_t = \ln(1+g)\,dt + \sigma(t)\,dW_t^B, \qquad G_t = B_0\,(1+g)^t. \tag{9}$$

Because the drift is the constant $\ln(1+g)$ rather than $\mu - \frac{1}{2}\sigma^2$, volatility enters only the *dispersion*: the typical (median) path realizes exactly g regardless of horizon or of how σ evolves, while the mean still runs above it by the usual Itô margin. The corridor around the g-line narrows over the hold—the asset de-risks as it matures—without disturbing the central rate. This is the floor of the four: a steady, equity-like return from today's price, carrying no view that Bitcoin is currently cheap or rich.

Diminishing returns (three variants). The growth thesis: a trend that compounds from spot at an *arithmetic* CAGR starting high and decaying geometrically toward a floor,

$$c(t) = c_\infty + (c_0 - c_\infty)e^{-\lambda t}, \qquad G_t = B_0 \exp\Big[\int_0^t \ln\big(1 + c(s)\big)\,ds\Big], \tag{10}$$

so that at the floor the price compounds at exactly $(1 + c_\infty)$ per year—the same arithmetic convention as GBM, which is what makes a 10% DR floor and the 10% GBM line directly comparable. All three variants share the decay speed $\lambda = 0.15$, set so each lands within one percentage point of its floor by year 20, and differ only in the pair (c_0, c_∞): **DR** $20 \to 10$ (the conservative variant, which barely clears GBM in the long run since the two share a 10% floor), **DR** $30 \to 15$ (the moderate central case), and **DR** $40 \to 20$ (the aggressive variant, whose 20% floor compounds for decades). The shared OU residual fluctuates around each trend with $x_0 = 0$. Table 1 and Table 2 trace the resulting CAGR and median-price paths.

Stochastic volatility (auto-calibrated overlay, on by default). Each process is overlaid with a regime-switching volatility. The instantaneous volatility mean-reverts in log space to the volatility the model already uses—the age-decaying $\sigma(t)$, shared by all four models—so the overlay changes the *shape* of the distribution, not the average volatility. Writing $\sigma_t = \bar\sigma(t)\,e^{\,y_t - \frac{1}{2}s^2}$ with $dy_t = -\kappa_v y_t\,dt + \xi_v\,dW_t^v$ and $s^2 = \xi_v^2/(2\kappa_v)$ the stationary variance, the $-\frac{1}{2}s^2$ term makes the overlay mean-preserving ($\mathbb{E}[\sigma_t] = \bar\sigma(t)$). The two dynamics are set from stylized facts about Bitcoin rather than fit to market data: volatility regimes mean-revert with roughly a three-month half-life ($\kappa_v = 3$) and swing about ±45% in log terms ($\xi_v \approx 1.1$), so volatility occasionally runs to 2.5× or down to 0.4× of its anchor. Because a leveraged balance sheet responds to volatility asymmetrically—a spike can trigger an irreversible forced sale on the downside but, when the firm is strong, can balloon the premium and fund accretive issuance on the upside (Section 7)—turning this overlay on makes the tails both fatter, and its net effect on the rational mNAV is small; with the fair-value anchor off (the from-spot regime) it is modestly dilutive, since issuance into volatility spikes then rides no recovering price.

Bitcoin model	Years from today 0	5	10	15	20	25	30
GBM (flat 10%)	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%
DR 20 → 10	20.0%	14.7%	12.2%	11.1%	10.5%	10.2%	10.1%
DR 30 → 15	30.0%	22.1%	18.3%	16.6%	15.7%	15.4%	15.2%
DR 40 → 20	40.0%	29.4%	24.5%	22.1%	21.0%	20.5%	20.2%
Volatility $\sigma(t)$	50.0%	38.9%	31.8%	26.9%	23.3%	20.6%	18.4%

Table 1: Trend CAGR by year—the annualized arithmetic rate at which each model's median path grows—with the shared age-decaying volatility in the last row. GBM is pinned flat at 10%; each DR variant decays from c_0 toward its floor c_∞ at $\lambda = 0.15$, landing within a point of the floor by year 20. Because all four share the arithmetic convention, the DR 20 → 10 floor and the GBM line coincide at 10%. Volatility follows $\sigma(t) = \sigma_0\, t_0/(t_0 + t)$ with $\sigma_0 = 0.50$, $t_0 \approx 17.5\,\text{yr}$, identical for all four models.

Bitcoin model	Years from today 0	5	10	15	20	25	30
GBM (flat 10%)	$59,000	$157,117	$261,119	$421,362	$678,691	$1.09M	$1.76M
DR 20 → 10	$59,000	$214,121	$412,915	$715,214	$1.19M	$1.95M	$3.16M
DR 30 → 15	$59,000	$304,113	$780,242	$1.74M	$3.67M	$7.56M	$15.37M
DR 40 → 20	$59,000	$422,055	$1.42M	$4.02M	$10.65M	$27.28M	$68.82M
Volatility $\sigma(t)$	50.0%	38.9%	31.8%	26.9%	23.3%	20.6%	18.4%

Table 2: Median Bitcoin price by year under each model with the 200-WMA fair-value anchor: every trend compounds from $F_0 \approx \$100{,}900$ at the Table 1 CAGRs while the observed price starts at $59,000 spot and converges to trend (half-life 1.25 yr); the shared volatility is repeated in the last row. Blended 30-year CAGRs from *spot*—convergence plus trend—are 12.0% (GBM), 14.2%, 20.4% and 26.5%—the one-time convergence adds ≈ 1.8 points over three decades, front-loaded in the first two years, after which the DR floors do the long-run work.

3.2 Policy / short rate (SOFR)

The model carries a single rate state, and it is the *short* rate, not the ten-year. The reason is structural and is developed in Section 4.3: STRC is a *variable-rate* perpetual whose dividend the issuer resets frequently against a short benchmark, so the rate that sets its dividend is SOFR—the overnight financing rate the Federal Reserve controls—following the issuer's own "spread-over-SOFR" framing. This is also the rate channel a hawkish "higher-for-longer" stance actually moves. We therefore *forecast SOFR* with one of four processes, jump-diffusion by default to reflect a debt-saturated world in which policy moves in discrete steps on meetings and fiscal news with an upward skew:

$$
\begin{aligned}
\text{Vasicek:} \quad & dr_t = \kappa(\theta - r_t)dt + \sigma_r\, dW_t^r, && (11)\\
\text{CIR:} \quad & dr_t = \kappa(\theta - r_t)dt + \sigma_r\sqrt{r_t}\, dW_t^r, && (12)\\
\text{Random walk:} \quad & dr_t = \sigma_r\, dW_t^r, && (13)\\
\text{Jump-diffusion (default):} \quad & dr_t = \kappa(\theta - r_t)dt + \sigma_r\, dW_t^r + J_t\, dQ_t, && (14)
\end{aligned}
$$

with current SOFR $r_0 = 3.62\%$, long-run neutral $\theta = 3.5\%$, reversion $\kappa = 0.40$ (policy rates adjust faster than long-term yields would), volatility $\sigma_r = 0.9\%$, and a reflecting floor. In the jump variant Q_t is Poisson with intensity 0.15/yr and jumps $J \sim \mathcal{N}(\mu_J, \sigma_J^2)$ with $\mu_J = +50$ bps (fiscal and debt shocks skew rates upward) and $\sigma_J = 150$ bps. A hawkish "higher-for-longer" view is expressed by raising θ (and r_0), a normalization by lowering θ; the shock dW^r may be correlated with dW^B to encode an inflationary regime. There is deliberately no separate long-rate forecast— Section 4.3 explains why a second rate state would mis-model a floating-rate preferred rather than improve it.

Joint fiscal-stress shocks (off by default). Optionally, each Poisson rate jump is treated as a single fiscal-regime event that re-rates Bitcoin in the same instant: when the jump fires, $\ln B$ receives a permanent step $\sim \mathcal{N}(\mu_B^J, (\sigma_B^J)^2)$. With $\mu_B^J > 0$ this encodes the debasement thesis—the event that spikes yields also bids up Bitcoin—while $\mu_B^J < 0$ encodes a risk-off liquidation. It is left off by default because acute fiscal shocks are typically risk-off at the event horizon even if debasement dominates the long run.

A constrained policy-rate reaction function (on by default). The default policy overlay constrains the Fed's *reaction function*: a fiscally-dominant Fed caps how high it lets the policy rate go, and how long it holds it there, once debt service binds. We model this as a gentle, asymmetric pull. The rate evolves freely below a critical value r^*, but above it faces a downward force proportional to the excess, $dr_t \mathrel{-}= \lambda_{\text{fd}} \max(r_t - r^*, 0)\, dt$. The force is one-sided and *dormant* below r^*, so at today's sub-r^* rate it does not misstate the present; it engages only when policy would otherwise turn restrictive. It is therefore distinct from the symmetric reversion-to-neutral already in the SOFR process: it is centered *above* neutral and acts only on the upside, encoding an *inability to stay high* rather than a tendency to return to a mean. The default $r^* = 4.5\%$—roughly a percentage point above the 3.5% neutral—is the mildly restrictive stance a debt-constrained Fed is reluctant to exceed, and with $\lambda_{\text{fd}} = 1.0$/yr it leans on the upper tail of the rate distribution: with neutral SOFR near 3.5% it engages only when policy turns restrictive, so its effect on the base case is modest while it does real work in a hawkish or fiscal-dominance scenario. We deliberately do *not* set r^* at the prior cycle's $\approx 5.5\%$ peak, which would imply the Fed could re-scale that peak unencumbered before any constraint binds—the opposite of the fiscal-dominance premise. A hard cap (immediate, dated, or sticky-on-breach) remains available as the limiting case of infinite strength.

4 The capital-structure engine

The novel content of the framework is the propagation of Strategy's actual liability stack, so that b_T and m_T are outcomes of financing decisions rather than free parameters.

4.1 Leverage

Define amplification as senior claims over Bitcoin value,

$$a_t \;=\; \frac{S_t + D_t}{H_t B_t}, \tag{15}$$

where S_t is preferred notional (liquidation preference) and D_t is convertible principal. New issuance targets an amplification level, and a ceiling $a_{\max}$ throttles it; when a drawdown pushes a_t above the ceiling the firm's financing options narrow sharply (below).

A countercyclical target. A fixed target is the wrong policy for a ratio whose denominator is this volatile: a crash multiplies amplification by the drawdown factor, so a target that is safe at a bear bottom is reckless at a mania top—0.45 taken at a 3×-WMA peak becomes ≈ 1.9 at the historical floor price, insolvency of the common's backing. The model therefore sets the target *countercyclically*, keyed to the same fair line the price anchor defines. The organizing principle is *constant stress amplification*: size the senior stack so that even at the historical worst price ($\approx 0.70\times$ the 200-week average—the 2015 and COVID wicks; 2018 and 2022 stopped nearer $0.80\times$) amplification stays at or below a tolerable $a_{\text{tol}} \approx 0.75$, beneath which the net-backing ratio holds above ≈ 0.25 and no forced deleveraging spiral engages. That single tolerance fixes the whole schedule. Writing $x_t = \ln(B_t/F_t)$ for the log-deviation of price from the fair line,

$$a_{\text{tgt}}(x_t) \;=\; \operatorname{clip}\!\big(a_{\text{fair}}\, e^{-x_t},\; a_{\text{lo}},\, a_{\text{hi}}\big), \qquad a_{\text{fair}} = \frac{a_{\text{tol}} \times 0.70}{\lambda} \approx 0.33, \tag{16}$$

with $\lambda \approx 1.61$ the WMA de-lag factor of Section 3.1, floor $a_{\text{lo}} = 0.20$ and ceiling $a_{\text{hi}} = 0.60$. Equivalently, and more simply: *keep the senior stack below* $\approx 52\%$ *of the 200-week-average value of the coin holdings, regardless of spot.* The schedule this produces is high near bear bottoms (0.55–0.60 at or below the WMA, where residual downside is smallest and the carry widest), ≈ 0.33 at fair value, and 0.20 at manias ($\geq 2\times$ fair), where the crash multiplier is largest. At the entry snapshot the feasible estimator (below) reads $\hat{x}_0 \approx -0.53$ and the schedule licenses $a_{\text{tgt}} \approx 0.56$ against an entry level of ≈ 0.41—headroom, consistent with the real firm levering *into* this bear. The absolute never-issue ceiling is retained at $a_{\max} = 0.75$; a static target (0.45) remains available with the schedule disabled.

The issuer's amplification metric, and the unit conversion. In July 2026 the issuer began reporting *Amplification* as BTC Reserve over Net Reserve,

$$A_t \;=\; \frac{H_t B_t}{H_t B_t - (S_t + D_t) + R_t}, \tag{17}$$

a ratio of gross coin value to coin value net of claims (with the USD reserve offsetting the claims), reading $\approx 1.51\times$ at the snapshot. The decision tool now expresses the policy dials in these units—target $1.34\times$ at fair value, floor $1.15\times$ at manias, ceiling $2.10\times$ at bear bottoms, absolute ceiling $3.00\times$, static fallback $1.60\times$—and converts at the boundary to the internal claims-over-value ratio via $a = 1 - 1/A + R/(HB)$, evaluated at the entry reserve share, so that the calibrated behavior of the schedule in equation (16) is preserved exactly at the current balance sheet. One property of the new metric is worth noting: because the reserve enters the denominator's netting, an A-target is slightly more permissive in claims terms when the reserve is fat and tighter when it drains—a mildly pro-survival feedback the raw claims-over-value ratio lacks.

Regime ignorance and the carry cap. The schedule above answers "where in the cycle are we?"—but the firm cannot answer that question with an oracle, and it cannot answer a second question at all: *which growth regime is Bitcoin in?* A treasury that could distinguish a DR $40 \to 20$ world from a flat one in real time would not need a policy rule. The model therefore imposes two layers of epistemic honesty. *First, the cycle position is estimated, not observed*: by default the deviation $\hat{x}_t$ is computed from a rolling average of the path's own realized prices (an EMA matched to the 200-week window's lag, seeded at the entry WMA) de-lagged by a slow EMA of realized growth—the same believe-the-tape discipline as the carry gate. The true fair line is never consulted; the model deliberately contains no mode that would grant the firm knowledge of it. *Second, a fractional-Kelly carry cap makes the target robust to regime ignorance*:

$$a_{\text{tgt}} \;=\; \min\!\left(a_{\text{fair}}\, e^{-\hat{x}_t},\; \phi\,\frac{\max(\hat{g}_t - c_t,\, 0)}{\hat{\sigma}_t^2}\right), \qquad \phi = \tfrac{1}{2}, \tag{18}$$

where $\hat{g}_t$ is the slow realized-growth estimate, c_t the firm's own STRC coupon, and $\hat{\sigma}_t^2$ the realized variance. If realized growth stops clearing the cost of leverage, the cap chokes new issuance no matter how cheap the price looks against a fair line that may itself be wrong—which is precisely the failure mode of a firm that believes in a recovery that never comes. At the entry snapshot the two legs read ≈ 0.59 (cycle) and ≈ 0.32 (carry), so *the carry leg binds and the operative entry target is* ≈ 0.32: the cycle says lever up, regime uncertainty says moderately.

Why countercyclical wins. Sweeping static targets across belief worlds shows the policy problem directly: in fast-growth worlds the target barely matters (Bitcoin value outruns the issuance pace, so ambitious targets are aspirations, never exposures), while in flat worlds high targets are destructive—the firm levers into the early convergence, the world stalls, and the leverage is stranded against a coupon the asset cannot out-earn. The countercyclical schedule Pareto-dominates every static target in the slow and flat worlds—higher valuation, better tail outcomes, fewer dividend pauses, lower terminal leverage, *simultaneously*—while costing nothing in fast worlds. It wins by construction: it delevers paths that run above trend before they revert, and levers below trend where the pull back to fair makes the carry widest and the residual downside smallest. The carry cap is then priced insurance: it costs $\approx$ 2–3% of valuation in the fast world and improves every survival metric in every other—including the nightmare configuration in which the price is flat *and* the fair-value thesis itself was wrong, where the capped policy holds terminal leverage near 0.55 against 0.66 for a static target and cuts the pause and underwater probabilities. One honest limit: no *target* policy guarantees survival, because in a stagnant world the ratio rises passively as obligations accrue; the cap stops new leverage but can only shrink the numerator as fast as the retirement channels allow. There is no hard *low* boundary beyond the mania floor: in a sustained advance the denominator grows while the senior stack does not, so amplification drifts down on its own and the firm re-levers only as price falls back toward the fair line.

4.2 Endogenous preferred dividend

The preferred stack is split into a fixed-coupon tranche and a floating tranche (STRC). The *required* yield the market demands of STRC is SOFR plus a spread that is the security's inherent baseline plus three *fading stress adders*, each keyed off already-simulated state:

$$q_t^{\mathrm{STRC}} = r_t + \underbrace{s_0}_{\text{baseline}} + \underbrace{k_R\, e^{-w_t/\tau_R}}_{\text{reserve runway}} + \underbrace{k_B \max(a_t - \underline{a},\, 0)}_{\text{senior coverage}} + \underbrace{k_V \max(\sigma_t^B - \bar{\sigma},\, 0)}_{\text{volatility}} - \rho_{\mathrm{cr}}\,\theta_t - \beta_{\mathrm{coll}} + d_t, \tag{19}$$

floored at $\underline{s} = 3.0\%$. The baseline $s_0 = 3.5\%$ is the spread a junior, Bitcoin-backed, retail-held perpetual commands even when fully confident of the dividend; the August 2026 recalibration (below) shifted weight from this fixed baseline into the coverage-dependent runway term, which now carries most of any observed discount. The *reserve-runway* adder is the central new term: with $w_t = R_t/(\bar{q}_t S_t)$ the years of dividend coverage the USD reserve provides, the adder is large when the runway is short and decays as coverage lengthens ($k_R = 12\%$, $\tau_R = 2.0\,\mathrm{yr}$; calibrated below). The *coverage* adder uses the claims senior-through-STRC—convertibles plus the preferred stack—over Bitcoin value, a_t; as converts are repaid or convert to equity, a_t falls and STRC rises in the waterfall, so this adder is also a structural tailwind ($k_B = 16\%$, $\underline{a} = 0.30$). The *volatility* adder ($k_V = 9\%$, $\bar{\sigma} = 0.55$) prices the risk premium in a turbulent tape. The rating-migration term $\theta_t \in [-1, 1]$ seasons upward while the firm pays in full and keeps its reserve and takes a sharp, sticky downgrade on any pause (up to $\rho_{\mathrm{cr}} = 2.5\%$ tightening; $\alpha = 0.35/\mathrm{yr}$, $h = 0.30$); an optional Basel collateral regime adds β_{coll}. The distress term d_t fires *only* in genuine last-resort distress, which the model reaches when three things coincide: (i) the USD reserve is drawn down to its operational floor (a minimum the firm holds back rather than spend, so reserve cash can no longer cover the dividend), (ii) the equity mNAV (m_t) sits below the *enterprise-parity gate*—the level at which enterprise mNAV equals 1.0, computed each quarter from the live capital structure—so that meeting the dividend through the common ATM would no longer raise even net backing per share, and (iii) the STRC ATM is itself unavailable (below par). In that state the firm does not immediately pause: it sells just enough Bitcoin to cover the dividend, leaving the reserve at its floor, and pauses only if even an orderly per-quarter liquidation (capped at 5% of the stack) cannot meet the bill—a catastrophic-collapse tail. The two resolutions are not equally severe, so the term is tiered: a forced Bitcoin *sale*, which keeps the dividend current, adds 2.0%, while an outright *pause*, which defers the dividend into compounding arrears, adds the full 4.0%. Neither is triggered by the small, discretionary Bitcoin sales Strategy has used to demonstrate process.

The offered coupon ratchets, capped by affordability. Management does not pay the required yield instantly; it resets the dividend toward it in small monthly steps to defend par, but the achievable rate is capped by what the reserve can *afford*:

$$c_t = \mathrm{clip}\Big(c_{t-\Delta} + \mathrm{clip}\big(\min(q_t^{\mathrm{STRC}},\, c_t^{\max}) - c_{t-\Delta},\, -\Delta c,\, +\Delta c\big),\, \underline{c},\, \bar{c}_{\mathrm{will}}\Big), \qquad c_t^{\max} = \mathrm{clip}\Big(\frac{R_t}{\phi\, S^{\mathrm{pref}}},\, \underline{c},\, \bar{c}_{\mathrm{will}}\Big), \tag{20}$$

with a step limit $\Delta c = 0.75\%$ per quarter ($\approx 25\,\mathrm{bps/month}$), an affordability horizon $\phi = 0.6\,\mathrm{yr}$, and a *willingness ceiling* $\bar{c}_{\mathrm{will}} = 15\%$. The ceiling is adopted here as a reasonable modeling assumption rather than a figure attributable to any management statement. The mechanism is deliberately self-limiting: raising the coupon enlarges the dividend bill and *shortens* the runway, which widens the required yield—so beyond a point a higher coupon is counter-productive, and the security absorbs stress in *price* rather than in coupon. With a constant fixed-tranche coupon q^{fix} on notional S^{fix}, the blended preferred coupon is $\bar{q}_t = [S^{\mathrm{fix}} q^{\mathrm{fix}} + (S_t - S^{\mathrm{fix}})c_t]/S_t$.

4.3 STRC price, the reserve, and ATM availability

STRC is engineered as a low-volatility, near-par instrument, and management has stated it will issue new shares through the at-the-market (ATM) program *only at or above par*. Modeling the ATM therefore requires a model of the STRC *price*, not just its required yield—because there are extended periods (the snapshot is one) in which the price trades below par and the ATM is simply shut, even though the security remains over-collateralized.

Which rate drives STRC—and why one forecast suffices. It is worth being precise here, because STRC invites a genuine confusion. By *Macaulay* duration—computed on its cash flows *as if the dividend were fixed*—a perpetual has very long, effectively unbounded, duration, which would seem to make its price acutely sensitive to the *long* rate. But Macaulay duration measures the wrong thing for this security, because STRC is a *variable-rate* perpetual: the issuer resets the dividend frequently (it accrues and pays on a roughly twice-monthly cadence) against a short benchmark. For a frequently-resetting floater, a change in the general level of rates flows through into the *coupon* rather than into the price; its *effective* (rate) duration is short—on the order of the time to the next reset, not the maturity. The frequent dividend cadence is precisely what collapses the effective duration toward the front end. The right benchmark for the *dividend* is therefore the short rate, SOFR, and a one-for-one rise in SOFR raises the required dividend one-for-one (holding the spread fixed). Pricing the dividend off the ten-year instead—treating a floater as a fixed-rate long bond—would be the actual modeling error.

The "long-duration" character does not vanish; it *migrates from rate risk to spread risk*. A perpetual has no maturity to pull its price back to par, so what remains highly sensitive is the *credit/structure spread* the market demands over SOFR (its *spread* duration is long) and the *lag* between a move in the required yield and the issuer's next reset. A Bitcoin-bear, risk-off tape widens that spread regardless of where SOFR sits; a SOFR scare or a vol spike opens a temporary gap before the dividend catches up. Both are captured here—not by a second long-rate state, but structurally, by the pricing identity below: the spread enters through the reserve-runway, coverage, and credit terms, and the reset lag is the small-step ratchet of the coupon toward the required yield. So the two "rate forecasts" a reader might expect collapse into one explicit forecast (SOFR, Section 3.2) for the coupon, plus a long *spread*-duration price risk handled by par defense—rather than a short-rate coupon inconsistently discounted at a long rate.

Price is the dividend over the required yield. A perpetual's price is par scaled by the ratio of the dividend it pays to the yield the market demands,

$$P_t^{\mathrm{STRC}} \;=\; \Pi \cdot \min\Big(\bar{p},\; \frac{c_t}{q_t^{\mathrm{STRC}}}\Big), \qquad \Pi = \$100, \quad \bar{p} = 1.003. \tag{21}$$

The snapshot confirms the identity: a paid coupon of 12% against a required yield of $\approx 13.0\%$ implies $100 \times 12/13.0 \approx \$92$, matching the observed \$92.32 print. So we do not impose a price band; the price is the gap between what Strategy pays (c_t, the ratcheting coupon above) and what the market requires (q_t^{STRC}). The cap $\bar{p}$ encodes *par defense from above*: when STRC would trade at a premium, management trims the coupon (within the reset limit) so the fundamental hugs par rather than floating above it.

The reserve is what the market prices. The dominant term in the required yield is the reserve runway, and the recent history bears this out almost mechanically. Strategy established a USD reserve of \$2.25 B in early 2026—about 2.5 years of dividend coverage—and STRC held its stated amount. A subsequent convertible buyback funded from the reserve cut coverage to roughly six months and STRC slipped to $\approx$ \$97. By mid-2026 the reserve had been rebuilt to $\approx$ \$1.4 B, yet STRC traded near \$74: coverage had actually *improved* from the spring, but Bitcoin had fallen sharply and its volatility spiked, widening the coverage and volatility adders faster than the runway adder narrowed. Over the same window the dividend climbed from 9% toward 12% as Bitcoin weakened. By August 2026 the reserve stood at \$4.0 B—about 28 months of coverage—and STRC had recovered to $\approx$ \$92: closer to par, not yet at it. Reading the market's response along that build implies that par requires roughly 2.5–3 years of coverage at the current coupon and balance sheet, and the runway parameters ($k_R = 12\%$, $\tau_R = 2.0\,\text{yr}$, jointly with the reduced baseline $s_0 = 3.5\%$) are calibrated to exactly that pair of anchors: the model reproduces the $\approx$ \$92 print at 28 months and clears par near 2.95 years. The exponential form gives confidence-per-dollar the diminishing returns one would expect: an additional year of coverage relieves the required spread by $\approx$ 3.6 percentage points at one year of runway, $\approx$ 2.2 at two, and $\approx$ 1.3 at three. The model reproduces this non-monotonicity directly: the reserve sets the regime, but Bitcoin's level and volatility move the price within it. (The structural terms alone do not fully reproduce the spring's resilience—STRC held near par on only six months of coverage—which we read as the market extending confidence that was later withdrawn, rather than as a stable equilibrium.)

The ex-dividend sawtooth and a smooth ATM gate. STRC accretes its (roughly twice-monthly) dividend and drops on the ex-date, with documented selling pressure in the following days. A price defended exactly at par is therefore above par only about half of each cycle. Rather than a hard par gate, the ATM-open fraction is the smooth share of the cycle the price spends at or above par,

$$\text{open}_t \;=\; \text{clip}\!\left(\frac{P_t^{\text{STRC}} - \Pi}{A_{\text{saw}}} + \tfrac{1}{2},\; 0,\; 1\right) \cdot \left[q_t^{\text{STRC}} \le \bar{q}^{\,\text{ceil}}\right] \cdot \left[a_t < a_{\max}\right], \tag{22}$$

with sawtooth amplitude $A_{\text{saw}} = \$2.2$ (the dividend drop plus post-ex-div pressure). Issuance and reserve top-ups scale by this fraction. The behavioral upshot is sharp: even a STRC defended successfully at par funds the ATM only about half the time, so Bitcoin accumulation through it is materially slower than a continuously-issuable instrument would suggest.

The bands are outputs, not assumptions. Because the price is now *derived* from the simulated reserve, senior coverage, and volatility, the time STRC spends near par and the fraction of the horizon its ATM is open are emergent results of each Bitcoin path, not inputs. Under the anchored default the early spot-to-fair convergence lifts the reserve and coverage quickly, so the window opens meaningfully for every outlook—the ATM is open $\approx$ 45% of a ten-year horizon under the default DR $40 \to 20$, and somewhat less under the conservative outlooks— while remaining shut through the entry bear phase in all of them; the duty cycle and coverage tests, not the Bitcoin outlook, set the ceiling. With the anchor off the picture is far starker: the slower from-spot trajectory leaves STRC discounted for most of the hold and the ATM open only a sliver of the horizon. The reserve is rebuilt toward its *floating* target (Section 4.5) via *common* issuance gated on capital-market access (which fades as amplification rises), *not* gated on the STRC ATM. A below-par STRC could never fund the reserve rebuild that lifts it back to par. The engine reports both the fraction of the horizon the ATM is open and the Bitcoin accumulated through it (Section 10).

4.4 Convertible notes

Each convertible tranche is tracked to its put date under an *assumed-diluted* share count: every tranche's conversion shares are pre-counted in N from the outset, consistent with the entry mNAV. At an early call date, if the price is sufficiently above the strike the issuer calls and holders convert—the shares are already in the count, so only the debt is retired. At the put date, if the price is above the strike the notes convert (again no new shares); otherwise the principal is paid in cash through the waterfall *and the pre-counted shares are removed*, so the diluted count falls when dilution does not occur. This keeps entry and exit on the same share basis and correctly credits shareholders for the conversion option expiring worthless. A small blended coupon is carried as a cash drag, and the deep-out-of-the-money tranches dominate forced-cash events in the near-term stress window.

4.5 Financing waterfall

Any cash need—an out-of-the-money convertible principal or a dividend bill that free cash flow cannot cover—is met in strict order:

1. free cash flow;
2. at-the-market *common* equity: accretive when $m > 1$; when m is between the *enterprise-parity gate* m_t^* (the equity mNAV at which enterprise mNAV $= 1.0$, recomputed each quarter from the live senior stack, reserve, and coin value) and 1 the firm still issues to protect the senior stack—dilutive to coins-per-share but still accretive to *net backing* per share—capped at a fraction of Bitcoin value per quarter;
3. at-the-market *preferred* (STRC), only while it trades at or above par (Section 4.3);
4. the USD reserve, drawn down to a hard operating floor—the *stress* channel, reached only when the issuance channels above are closed (below the parity gate) or capped out;
5. only if the reserve is exhausted as well, a forced *sale of Bitcoin*.

The ordering reflects the issuer's stated 2026 commitments: dividends and interest are funded from ongoing sources first, the reserve is spent on dividends and interest *only* (buybacks, preferred retirement, and coin purchases are funded elsewhere), and a minimum of twelve months of dividend coverage is maintained. In consequence the reserve does not saw-tooth between a trigger and a ceiling; it *floats* at a target the model derives endogenously each quarter. The target is the coverage at which the runway term of the required-yield model (Section 4.3) leaves STRC exactly at its par-clearing yield given the current coupon—inverted in closed form from the spread equation—plus a three-month buffer, floored at the stated twelve-month minimum and clamped at forty-eight months as a sanity bound for stress regimes. The clamp rarely binds, because the ceiling is economically endogenous: beyond the par-clearing runway the marginal confidence purchased by another reserve dollar is approximately zero (the price is capped near par by the call) while its opportunity cost against coins remains positive, so the optimal reserve *is* the requirement. At the entry snapshot the requirement reads ≈ 38 months—implying a build from today's \$4.0 B toward $\approx$ \$5.5 B—and it declines as coverage, seasoning, and over-collateralization improve; the median simulated path settles near 26 months by the horizon. Deep-discount regimes, in which equity refills would be intolerably dilutive, rebuild only to the twelve-month minimum via the collateral-constrained coin-sale channel below. The base case starts the reserve at \$4.0 billion, with a hard operating floor of roughly three months. The preferreds carry no forced-liquidation trigger: rather than default, the firm funds the dividend through the waterfall above—including, as a last resort when the reserve is exhausted, common issuance is intolerably dilutive, and the STRC ATM is shut, an orderly sale of Bitcoin sized to cover the bill—and only *pauses* the dividend if even that capped sale falls short, with cumulative series then accruing arrears that compound until cleared. One further refill channel reflects the firm's observed 2026 playbook: when the common trades at a sub-1 mNAV—so that issuing equity to refill the reserve would be dilutive—the firm instead sells a collateral-constrained trickle of coins to restore the dividend reserve. The sale is non-dilutive but shrinks the stack, and its cost is charged where it belongs: in the realized Bitcoin-per-share yield that the premium capitalizes.

The defensive gate is *dynamic*: the firm keeps selling common to honor its senior commitments while the equity trades at or above the enterprise-parity level $m_t^* = (1 - s_t)/x_t$ (with s_t the net senior claims over gross coin value and x_t the ex-convert share fraction), and no further. Above that line, issuance at any equity discount still raises net backing per share; below it, dilution fails even that test and the waterfall prefers coins. The gate therefore *widens* as the firm de-levers (with a light senior stack, deep discounts still clear parity) and *tightens* as leverage builds—replacing the fixed floor of earlier versions (0.70) with a boundary that adapts to the balance sheet. It remains the single most consequential policy lever, trading dilution of the common against forced coin sales. It is worth being explicit about what this issuance does and does not do for the common holder. Below $m = 1$ it is *not* accretive: selling shares for less than the Bitcoin standing behind them lowers Bitcoin-per-share. What it buys is continuation of the strategy. By meeting the dividend and principal obligations with common equity rather than with coins, the firm stays current on its senior commitments and so defends its credit standing—which is precisely what allows the amplification machine to keep running into the next upcycle. Crucially, the senior claims carry no maintenance covenants and no forced-liquidation trigger, so below-par issuance is a discretionary cost the firm elects to bear in order to protect its standing, never an obligation that could compel asset sales at the bottom of a drawdown. Put plainly, the common holder accepts some Bitcoin-per-share dilution in the bad times as the regrettable but bounded price of preserving the amplification that delivers outsized Bitcoin-per-share accretion in the good times. The model treats the level of this floor as a policy choice rather than a fixed constant, since reasonable people will weigh that trade differently.

Proactive accretive issuance and the collateral-capacity channel. The waterfall above is defensive—it raises cash only to meet an obligation. The firm also issues *offensively*: when issuance is accretive and Bitcoin is in strength (above a trailing reference), it sells common equity and buys Bitcoin. The naive accretion condition is $m > 1$: issuing ΔN shares at premium m changes backing per share by a factor $1 + \Delta N(m - 1)/(N + \Delta N)$. But that condition ignores what the new coins do one level up the stack. Each dollar of Bitcoin added to the balance sheet raises the amplification-target STRC capacity by a_{tgt} dollars, and STRC issuance buys coins with *zero* common dilution. When that unlocked capacity can actually be deployed—the carry is open and the firm is capacity-bound rather than placement-bound—the *packaged* gross-yield effect of a dollar of common issuance is

$$\frac{E}{H_t B_t}\Big[(1 + a_{\text{tgt}}\, u_t) \; - \; \frac{1}{m}\Big], \tag{23}$$

where $u_t \in [0, 1]$ is the unlock fraction (one when the carry is open and remaining room is below the quarterly placement pace, zero otherwise). Issuance is gross-BY accretive whenever $(1 + a_{\text{tgt}} u_t)\, m > 1$, where a_{tgt} is the robust countercyclical target of equations (16) and (18): at the entry state ($a_{\text{tgt}} \approx 0.32$, carry-capped) fully unlocked, the effective floor of the accretive ATM slides from 1.01 down to $1/(1 + a_{\text{tgt}}) \approx 0.76$, clamped from below by the same enterprise-parity gate m_t^* (≈ 0.68 at the entry structure)—the firm can rationally issue at a *sub-1 mNAV*, paying measured dilution to expand the collateral base for non-dilutive preferred-funded accumulation, exactly the behavior Strategy exhibited through 2025. Two honest qualifications. In *net-backing* terms the bar is higher, nearer 0.8 under an open carry, because the unlocked coins arrive encumbered and their value to the common is the carry spread on them; the model prices this correctly through ρ_t and the realized yield. And quantitatively the channel proves second-order: the momentum kernel usually re-rates the premium through par before the unlock binds, so the mechanism matters for *which quarters* the ATM runs more than for terminal values. The parity clamp applies regardless: no issuance, however packaged, runs below the enterprise-parity line. The pace scales with the packaged margin, $\min\bigl(\eta\,[(1+a_{\mathrm{tgt}}u_t)m-1],\,\bar{a}\bigr)$ of Bitcoin value per quarter, with slope $\eta = 0.10$ and a market-depth cap $\bar{a} = 12\%$: a trickle near the floor, a truck through a 2–3× window. Because the proceeds buy coins, the program *de-levers* while it runs, and it is the mechanism through which a momentum-driven premium spike is converted into permanent Bitcoin-per-share.

Proactive discount buybacks. The mirror image runs on the discount side. Strategy's stated aim is to grow Bitcoin-per-share, not the coin stack as such, and when its own equity trades below net asset value it can advance that aim by selling Bitcoin to *retire* shares: a share bought back at a discount is backed by more Bitcoin than the sale raises, so the coins left behind are spread over fewer shares. Retiring ΔN shares funded by a Bitcoin sale changes backing per share by a factor $1 + \Delta N(1-m)/(N-\Delta N)$—the exact reflection of the issuance factor above, accretive precisely when $m < 1$ and more so the deeper the discount. The pace scales with the discount, $\min\bigl(\eta_b\,(1-m),\,\bar{b}\bigr)$ of Bitcoin value per quarter (slope $\eta_b = 0.10$, cap $\bar{b} = 5\%$). The trigger is not a fixed constant but the *same effective floor* that governs the accretive ATM: the firm buys back exactly when the mNAV sits below the packaged-accretion boundary of equation (23), and issues when it sits above. The two common-equity operations are therefore mutually exclusive by construction—the model can never simultaneously call for issuing MSTR to buy Bitcoin and selling Bitcoin to retire MSTR—and the boundary moves with the state: near 1.0 when the carry is closed (a wider buyback zone in bears, where harvesting even a shallow discount beats deploying into a closed carry), sliding toward $1/(1+a_{\mathrm{tgt}})$ when the carry is open and capacity-bound. Coin sales for buybacks are additionally collateral-constrained: a sale may not push amplification above its ceiling. It is *coverage-gated*: buybacks run only when the dividend was met from operating sources, no Bitcoin was force-sold, and the reserve sits above its low band, so the program never competes with a dividend defense; like the issuance side it carries no reflexive price impact.

Two consequences are worth stating. First, at a discounted entry the buyback is a live, first-order channel of the valuation itself: the model's firm spends the bear phase harvesting its own discount, converting coins into Bitcoin-per-share at better than par, and part of the personal rational mNAV of Section 7 is precisely the capitalized value of that harvesting. Second, the program deliberately *shrinks* the coin stack to raise Bitcoin-per-share, so the terminal-stack and share-of-supply figures of Section 6 fall even as backing per share rises: growing the stack and growing Bitcoin-per-share are no longer the same objective.

Preferred retirement. One rung above the buyback sits its mirror on the senior stack: when the variable-rate preferred trades at a deep discount to par ($\leq 95\%$), the firm can tender for it, retiring a perpetual coupon obligation for cents on the dollar and jumping the net backing of every remaining share. Two funding legs, each with its own rationality gate. The *coin-funded* leg sells Bitcoin within the collateral constraint, and fires only when the locked-in retirement yield beats holding coins—the prevailing coupon divided by the purchase price must exceed the firm's trailing growth belief g_t^{ema}; under a bullish tape the coins are worth more deployed than the coupon saved, and the gate correctly stays shut. The *common-funded* leg issues equity even at a sub-1 mNAV whenever the swap is accretive—the common's discount is shallower than the preferred's, $m > (P^{\text{STRC}}/\text{par})(1 + \text{margin})$—a capital-structure arbitrage that exchanges a cheap claim for a cheaper one. Pace scales with the discount depth, capped per quarter, and the channel is coverage-gated like the buyback. In the simulated bear phase the retirement channel is the firm's principal deleveraging tool, typically retiring 5–15% of the variable-rate tranche before the recovery re-rates it to par.

5 The premium (mNAV) function

5.1 The BY-hurdle identity

The premium the common trades at is modeled *endogenously*: at every step, on every path, the market inside the simulation grants the equity

$$m_t \;=\; \underbrace{\frac{H_t B_t - L_t + R_t}{H_t B_t}}_{\text{net-backing ratio } \rho_t} + \left(\mathrm{BY}_t^{\text{fwd}} - h_t\right) D, \tag{24}$$

floored at $m_{\min} = 0.30$: the net backing per unit of gross coin value, plus the capitalized value of D years of the *forward Bitcoin-per-share yield in excess of the structural hurdle rate.* A wrapper that converts capital-markets access into coins-per-share faster than its cost of leverage deserves to trade above its backing; one that cannot clear that bar deserves to trade below it. The three ingredients are the net-backing ratio ρ_t, the hurdle h_t, and the forward yield $\mathrm{BY}_t^{\text{fwd}}$; the capitalization duration D converts the annual excess into a multiple.

Net-backing ratio ρ_t. The first component anchors the premium to the Bitcoin that backs the common *after* the senior stack is satisfied: gross coin value $H_t B_t$, less senior liabilities L_t (preferred par, accrued dividend arrears, and convertible principal), plus the cash reserve R_t, all expressed per unit of gross coin value. It is the level to which the premium would collapse in the absence of any going-concern optionality, and it encodes the fact that the common is a *residual* claim: leverage levers its claim on the coin, so a balance sheet on which the coins only thinly cover the senior stack carries a low ρ_t and a correspondingly compressed premium. This same term is what keeps the *exit* valuation honest about the obligations a future buyer inherits. A buyer at the horizon purchases a junior claim on a company that must continue servicing the senior stack—most importantly the perpetual STRC dividend—and the model credits the common only with the coins net of those claims. For a perpetual preferred trading near par at a market-clearing coupon, subtracting its par from the asset base is precisely the capitalized value of the dividend stream the buyer inherits, since par equals that dividend divided by the required yield; the deduction is therefore the present value of an indefinite obligation, not a liquidation haircut. The burden enters through two non-overlapping channels: dividends paid *before* exit have already reduced the backing per share, charged period by period in the waterfall, while the dividends owed *after* exit are capitalized into L_T at the moment of sale. The buyer thus pays for the coin behind each share net of the senior claims, never the gross coin as though the seniors did not exist.

The hurdle h_t: the structural cost of the leverage machine. The rate the excess yield is measured against is the *structural* required yield on the marginal unit of leverage—the STRC required yield of the capital-structure engine, but restricted to its structural terms: the policy rate plus the inherent spread floor, the coverage adder in amplification, and the volatility adder, net of the seasoned-rating and regulatory credits, floored at the minimum spread. The *runway and distress adders are deliberately excluded.* This is a design principle, not an omission: if acute reserve stress fed the hurdle directly, a falling premium that closes the common ATM would drain the reserve, widen the spread, raise the hurdle, and depress the premium further—a self-referential spiral with no economic counterpart. In this architecture acute distress still reaches the common's premium, but through the *fundamental* channels: a stressed reserve throttles the issuance-capacity component of forward yield, and forced coin sales show up as negative realized yield. The coupling exists; its loop gain is bounded well below one. At the entry snapshot $h_0 \approx 11\%$.

The capitalization duration D. The market does not capitalize excess yield in perpetuity; empirically it pays for roughly the stretch of trend it is willing to extrapolate. D is backed out at 2.0 years from three independent anchors that agree: inverting the identity at the last observed premium-to-yield pair on the way into the 2026 compression (mNAV 1.2× against a realized BY of 22.8%); the spread between demonstrated yield and the STRC cost of leverage; and the breakeven against a 10%-BY skeptic case. D also fixes a sharp structural result. Under sustained leverage a the firm's yield engine produces BY $\approx a\,g$ (issuance keeps pace with Bitcoin growth g), so $m \approx (1-a) + (ag - h)D$ and

$$\frac{\partial m}{\partial a} \;=\; gD - 1\,: \tag{25}$$

leverage raises the multiple only when $gD > 1$. With $D = 2$ the flywheel ignites only above $g \approx 50\%/\text{yr}$ —which is precisely the difference between the 2024 euphoria (high momentum, premium 2.5–4×) and the 2026 compression, where no amount of leverage can lift the multiple and the wrapper trades at a discount. The model holds D fixed and lets *momentum* carry the regime, which reproduces both phases without a free switching parameter.

5.2 Forward Bitcoin Yield: demonstrated, momentum, capacity

The yield the market capitalizes is a blend of three measurements, all built strictly from *physical* quantities—coins, shares, balances, and the exogenous Bitcoin price—so that the equity price never feeds its own premium:

$$\mathrm{BY}^{\mathrm{fwd}}_t \;=\; w_r\,\mathrm{BY}^{\mathrm{ema}}_t \;+\; w_m\,\varphi(g^{\mathrm{ema}}_t) \;+\; w_c\,\mathrm{BY}^{\mathrm{cap}}_t, \qquad (w_r, w_m, w_c) = (0.35,\, 0.50,\, 0.15). \tag{26}$$

Demonstrated yield $\mathrm{BY}^{\mathrm{ema}}$. An exponential moving average (half-life 1.5 years) of the *realized* annualized growth in coins per diluted share. Because it is computed from the physical ratio, it captures every channel with no double counting: STRC-funded purchases and accretive issuance add, dilutive reserve refills and coin sales for dividends, buybacks retire shares, preferred tenders spend coins—each lands in the accountant's number with its true sign. It is seeded at the entry snapshot to the decelerating trailing realized yield (21.5%; FY2025 printed 22.8%).

The momentum kernel φ. The dominant term prices the market's habit of *front-running the trend.* Let g^{ema}_t be an EMA (half-life 1.0 year, matching the twelve-month window of the calibration below) of realized Bitcoin log returns along the path. The kernel is piecewise linear, convex, and asymmetric,

$$\varphi(g) \;=\; \beta_0 \;+\; \kappa^+ \max(g - g_0,\, 0) \;-\; \kappa^- \max(-g,\, 0), \qquad (\beta_0, \kappa^+, \kappa^-, g_0) = (0.15,\, 1.10,\, 0.10,\, 0), \tag{27}$$

The economic mechanism is the one a levered treasury actually experiences: a rising Bitcoin price deleverages the balance sheet and reopens issuance capacity, so trend translates into *deployable* future yield at a steep slope; a falling price closes the machine, but the premium cannot price yield far below zero—it compresses toward the net-backing floor and stops responding, hence the nearly flat downside. The kernel is evaluated on *realized* returns of the exogenous price process, so it introduces no self-referential feedback whatsoever.

Issuance capacity $\mathrm{BY}^{\mathbf{cap}}$**.** The forward-looking correction: what the machinery could add *this quarter*, annualized—the duty-cycled STRC room toward the amplification target when the adaptive carry is open, plus the accretive common-ATM contribution at its packaged margin (Section 4.5). Only capacity facts enter: room, pace caps, the preferred's duty cycle, the carry gate. Because the ATM leg depends on m_t itself, premium and capacity are solved together as a damped fixed point (six half-steps per quarter). The loop gain of that circle is bounded by $w_c/\Delta t \cdot D \cdot 2\bar{f}\,(1 + a_{\mathrm{tgt}}) \approx 0.17 \ll 1$ (asserted at initialization), so the fixed point is contractive by construction; the weights are therefore doing double duty, expressing a "show-me" market that mostly believes demonstrated yield while simultaneously keeping the only self-referential term small.

Empirical calibration. The kernel is fitted on the full two-cycle public record, roughly 23 quarters of paired (mNAV, Bitcoin price) observations from 2020Q4 through 2026Q2 assembled from public trackers (levels carry ± 0.15 of definitional noise across sources; the directional results are robust). Three findings shape the functional form. *First*, the premium tracks momentum, not price: correlation of mNAV with trailing twelve-month momentum is $+0.74$, against $+0.58$ for six-month momentum and $+0.55$ for the log price level—the price rose $\sim 30\times$ across the sample while the premium round-tripped twice. *Second*, the response is amplified: quarterly $\Delta \ln(\mathrm{mNAV})$ moves $\approx 0.5\times$ quarterly $\Delta \ln(B)$, which stacked on the mechanical one-for-one makes the equity a $\approx 1.5\times$ amplifier of Bitcoin moves with the premium contributing the excess. *Third*, the response is convex and asymmetric: inverting identity (24) at $D = 2$ across the sample, the implied forward yield reaches $\approx 105\%$ at the 2024Q4 peak (twelve-month momentum $\approx +80\%$), while in the 2022 and 2026 troughs it floors at 10–20% regardless of how negative momentum runs—hence $\kappa^+ \approx 1.1$ against $\kappa^- \approx 0.1$. The 2026Q2 inversion (≈ 0.15) independently matches the entry seeding, a consistency check the calibration did not impose.

5.3 A single endogenous pass

Earlier formulations of this framework separated *whose beliefs set the price* from *whose set the premium* by simulating a market premium path under a second, independently chosen Bitcoin outlook and injecting it into the investor's simulation. The momentum-calibrated premium makes that construction unnecessary, and the model now runs as *one* simulation. Your Bitcoin model drives the realized price path; the market inside the model prices each path's own realized momentum and demonstrated yield through equations (24)–(27). The premium is therefore *path-dependent*: a path that rips higher re-rates, a path that grinds sideways stays discounted, and the exit dispersion lives in both the backing per share b_T *and* the terminal premium m_T—as it does in reality.

Three properties of the construction are worth stating. *First, the entry premium is a prediction, not an input.* The model's justified entry m_0 is its own first-step premium evaluated at the entry balance sheet: $m_0 = 0.83\times$ against the observed quote of $0.74\times$. The two no longer coincide, and the wedge is informative rather than embarrassing: it is the premium's immediate response to the fair-value anchor's assertion that spot sits well below trend—the first quarter of convergence already registers in the kernel. One unobservable—the trailing-yield seed—is pinned by inverting the entry quote in the from-spot configuration (an anchor-independent calibration, exactly as an implied volatility is pinned by an option price); the resulting seed sits plausibly between the decelerating FY2025 realized print and the stalled current run-rate, and the sensitivity analysis of Section 10 shows the seed washes out of every conclusion beyond the first two years. *Second, the terminal mark uses the same identity.* At the horizon each path is marked at $\rho_T + (w_r \, \mathrm{BY}_T^{\mathrm{ema}} + w_m \, \varphi(g_T^{\mathrm{ema}}) - h_T) \, D$—the capacity term is dropped at exit, a mildly conservative choice. *Third, the model's market is deliberately myopic.* It extrapolates the trailing tape plus present capacity; it does not project the full future arc of the policy engine (the D parameter carries a crude two-year projection and no more). A patient investor simulating the whole arc can therefore rationally value the equity above what the model's market pays for years at a stretch. That wedge is not an inconsistency; it is the quantity this framework exists to measure.

6 Finite-supply accumulation brakes

The capital-structure engine, left unconstrained, would let the firm issue and accumulate Bitcoin without bound—incompatible with a fixed 21 million-coin supply. Rather than impose a terminal holding as an input, the model lets one *emerge* from three economically grounded negative feedbacks, each keyed to the firm's share of accessible supply,

$$\theta_t \;=\; \frac{H_t}{S_{\mathrm{acc}}}, \qquad S_{\mathrm{acc}} = 17.3\,\text{million BTC}, \tag{28}$$

where S_{acc} is the 21 million cap less coins presumed permanently lost. As θ_t rises the three brakes tighten, and the terminal stack asymptotes as a *derived* output whose level depends on beliefs.

(1) Float-depletion slippage. A large buyer is a price-*taker* on the global Bitcoin trend but a price-*mover* on its own executions: it must walk a thinning order book. The chosen Bitcoin process therefore remains the market price B_t—the numeraire, the mark, and the mNAV denominator—and is never overwritten; finite supply enters only as a slippage markup on Strategy's *own* marginal purchases,

$$c_t \;=\; B_t\Big(1 + \underbrace{k_{\mathrm{perm}}\,\tfrac{\theta_t}{1-\theta_t}}_{\text{permanent scarcity}} + \underbrace{k_{\mathrm{temp}}\,\tfrac{P_t/B_t}{S_{\mathrm{acc}}-H_t}}_{\text{temporary, buy-size}}\Big), \tag{29}$$

so a financing inflow P_t buys P_t/c_t coins. Both terms diverge as $H_t \to S_{\mathrm{acc}}$, so the brake binds on *quantity*: as the float thins each dollar buys fewer coins and the stack approaches the float asymptotically, while B_t continues to follow whatever model the user selected. Reflexive "buying lifts the price" is deliberately excluded—each Bitcoin model already embeds its own demand path. The permanent coefficient is calibrated to strong-hands scarcity: $k_{\mathrm{perm}} = 1.25$ implies a $\approx 6\%$ premium at today's $\approx 5\%$ ownership (consistent with the firm's observed near-spot execution) rising to $\approx 42\%$ at quarter-network ownership and 125% at half.

(2) The concentration brake: pace and yield. The second brake acts through two channels keyed to the same state variable. The *explicit* channel constrains quantity at the source: every purchase-side pace cap—the STRC issuance pace and the accretive-ATM ceiling—is scaled by

$$g_t^{\mathrm{conc}} \;=\; \left(\frac{1-\theta_t}{1-\theta_0}\right)^{p_{\mathrm{conc}}}, \qquad p_{\mathrm{conc}} = 5, \tag{30}$$

the statement that the flow of coins the market can source per quarter shrinks with the remaining free float, steeply: the multiplier is ≈ 1 at today's concentration (no present-day footprint), ≈ 0.3 at quarter-network ownership, and ≈ 0.1 at 40%. Because the premium's issuance-capacity term reads the same caps, the market inside the model prices the shrinking future capacity automatically. The *emergent* channel works through yield: slippage raises the all-in cost of every marginal coin, realized $\mathrm{BY}^{\mathrm{ema}}$ erodes, the forward yield the market capitalizes falls, the premium compresses through identity (24), and the accretive ATM self-halts. Together with the recalibrated slippage the arc is tamed: the median stack under the default beliefs reaches $\approx 25\%$ of accessible supply at thirty years, $\approx 35\%$ at fifty, and $\approx 45\%$ by eighty. An honest caveat applies to all deep-concentration parameters: no one has ever bought a quarter of Bitcoin, so p_{conc} and the deep-θ reach of k_{perm} have no empirical anchor; they are coherent statements about strong-hands scarcity, calibrated so that accumulating more than a quarter of the network within three decades is unlikely rather than routine.

(3) Adaptive STRC carry closure. Preferred-funded accumulation only makes sense while the asset out-earns the coupon. Proactive STRC issuance toward the amplification target is therefore gated on

$$g_t^{\mathrm{ema}} \;>\; c_t^{\mathrm{STRC}} + \mathrm{margin}, \tag{31}$$

where g_t^{ema} is the same trailing realized-momentum EMA that drives the premium kernel: the firm, like the market, believes the tape rather than an oracle. The margin is set at 2%—management does not run razor-thin carry. The gate is closed at the entry snapshot (trailing momentum is negative after the crash, matching the real firm's halted ATM) and reopens as the trailing EMA turns—inside the first year under the anchored default, whose convergence supplies the momentum, and roughly 1.2 years into an unassisted recovery, consistent with the $\sim$one-year re-issuance lag observed after the 2022 bear. As trend growth decays toward the prevailing coupon at long horizons, or a distressed coupon rises to meet it, the carry closes and the STRC channel stops adding coins. The same comparison, run in reverse, gates the preferred-retirement channel of Section 4.5: the firm tenders for its own discounted preferred exactly when the locked-in yield beats holding coins at the believed growth rate.

An emergent, belief-dependent terminal stack. Together the brakes replace an arbitrary terminal holding with a derived arc. Under the default beliefs (your DR 40 $\to$ 20 with the fair-value anchor and the robust leverage policy) the median path traces the full cycle endogenously: a bear-market entry in which the carry is closed and the firm's coin flow is *negative* (buybacks, preferred tenders, reserve refills funded by coin sales—the 2026 posture); the spot-to-fair convergence registering as momentum, reopening the carry gate inside the first year and re-rating the premium through par in roughly three quarters; and the accretive engine restarting against tightening brakes. The median stack runs ≈ 1.13 M coins at three years, 1.49 M at five, 2.29 M at ten, 3.44 M at twenty, and 4.39 M at thirty ($\theta \approx 25\%$), with the concentration brake governing the pace beyond. The asymptote remains belief-dependent and is properly an output rather than a dial.

7 Fair value and the rational mNAV

Each simulated path produces an exit value for a single share, expressed in Bitcoin: $v_T \equiv b_T\, m_T$, the product of the Bitcoin backing per share at the horizon, $b_T = H_T/N_T$ (coins held divided by diluted shares), and that path's own terminal premium m_T—the endogenous mark of Section 5.3, granted by the model's market for the momentum and yield the path actually realized, not a level inflated by the investor's optimism about the future beyond it. A share bought today for V_0 dollars costs V_0/B_0 coins at the current spot price B_0, and the buyer is compensated for the wrapper's incremental risk only if the median terminal coin value clears that entry cost grown at the hurdle. Equating the two and solving for the fair dollar price gives

$$V_0 \;=\; B_0\,\frac{\mathrm{med}(v_T)}{(1+r_p)^{\tau}}, \tag{32}$$

where V_0 is the fair MSTR price today, B_0 the current Bitcoin spot price, $\mathrm{med}(\cdot)$ the cross-path median, r_p the annual counterparty/wrapper hurdle, and τ the holding horizon in years. The factor $\mathrm{med}(v_T)/(1+r_p)^{\tau}$ is the present value, in coin, of the median exit share; multiplying by spot B_0 re-expresses that coin value as a dollar price. This choice of the median is not arbitrary: at a price set by (32), a path beats self-custody exactly when v_T exceeds its own median, which by construction happens with probability one-half. *The fair value is the price at which the probability of beating self-custodied Bitcoin equals* 50%—a self-consistent definition we verify numerically (re-simulating at V_0 returns a beat probability of 0.500). Translating to the premium of this paper,

$$m_0^{\star} \;=\; \frac{V_0\,N_0}{H_0 B_0} \;=\; \frac{\mathrm{med}(v_T)}{b_0\,(1+r_p)^{\tau}}, \tag{33}$$

the *rational mNAV*: the premium at which a marginal buyer with these beliefs is indifferent between the equity and the coin. A firm that merely held Bitcoin, never issued accretively, and exited at parity would have $m_0^{\star} = (1+r_p)^{-\tau} < 1$—a structural discount that compensates for wrapper risk. Every point of premium above one is, in this framework, a quantified bet that the accretion flywheel and the terminal premium will outrun that discount.

The model's premium versus your personal rational mNAV. Two distinct premia emerge, and they answer different questions. The *model's* premium is m_t of identity (24): what the equity trades at along each path, beginning at the justified entry $m_0 \approx 0.83\times$ (against the observed $0.74\times$; the wedge is the anchor's fair-value assertion, Section 5.3) and evolving with each path's realized momentum, demonstrated yield, and balance sheet. The investor's *personal* rational mNAV is $m_0^{\star}$ above: the entry premium at which an investor holding their own Bitcoin outlook and hurdle is indifferent, given that the position will be marked at the model's (median) terminal premium at sale. The wedge between them is the model's central output: a patient holder who simulates the whole arc—the recovery, the re-rating, the compounding of coins per share—can rationally value the equity well above what the myopic, trend-extrapolating market pays at entry. Because every input to the premium is a physical quantity, the equity price never enters its own valuation: re-feeding the fair price V_0 back as the input leaves $m_0^{\star}$ unchanged, so the rational mNAV is a function of beliefs and the balance sheet alone, not of the current market mispricing.

7.1 The optimal holding horizon

The finite-supply brakes make the fixed-horizon fair value a declining function of long horizons, and that shape carries a sharp implication: *the wrapper is a finite trade, not a forever hold.* Its alpha is front-loaded—the spot-to-fair convergence, the premium re-rating, and high Bitcoin-per-share accretion all occur while concentration is low and the carry is open—while the hurdle compounds at a constant rate indefinitely. The optimal pre-committed exit is therefore the horizon at which the *marginal* year's expected Bitcoin-denominated excess return stops clearing the hurdle,

$$\tau^{*} \;=\; \arg\max_{\tau}\; m_0^{\star}(\tau), \tag{34}$$

after which self-custody dominates and every additional year of holding subtracts value. Computing the full curve is nearly free, because nothing in the simulation dynamics depends on the horizon (every policy and estimator is trailing): a single long simulation contains every shorter horizon inside it, so annual *exit checkpoints* recorded along one 30-year run—each marked with the same capacity-less terminal formula—reproduce dedicated per-horizon runs exactly (bit-for-bit at matched step counts). Under the default profile Figure 1 shows the result: $\tau^* = 4$ years at a peak of 1.36×, with a broad plateau (3–6 years all within 1.33–1.36×)—the exact exit year matters little—before the concentration brakes and the compounding hurdle grind the curve through parity near 17 years. The headline should be read accordingly: the equity is worth its premium *held for roughly the plateau*, not held forever. Two qualifications. τ^* is the best *pre-committed* horizon, a lower bound on the value of the true path-dependent stopping rule (exit when the state, not the calendar, says the alpha is exhausted); and the real-world switch to self-custody bears tax friction the model does not price, so τ^* is a guide to the shape of the trade rather than a directive.

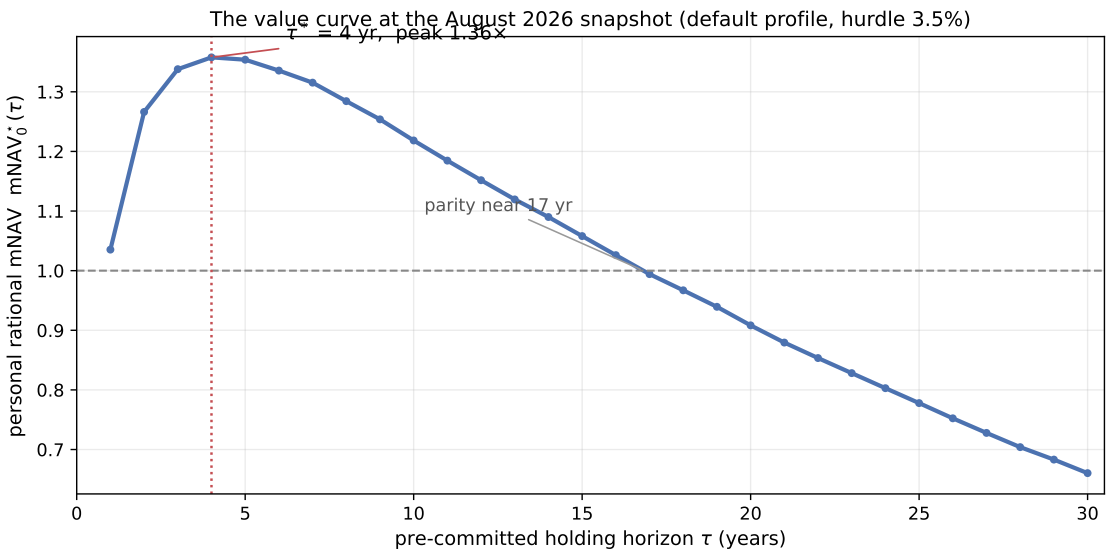

Figure 1: The value curve: personal rational mNAV against the pre-committed holding horizon under the default profile. The optimum τ^* marks where the marginal year's excess return stops clearing the 3.5% hurdle; the dashed line is parity with net asset value.

8 Calibration

Table 3 lists the snapshot used to anchor the base case. These are point-in-time inputs and are expected to move; the framework is designed to be re-calibrated rather than to assert fixed values.

8.1 Premium-model parameters

The endogenous premium of Section 5 adds one small parameter block, reproduced in Table 4. The empirically load-bearing entries are the capitalization duration D, the upside slope κ^+, and the momentum weight w_m—the three the sensitivity tornado of Section 10 identifies as dominant; the entry seeds are pinned by observables (the entry quote and the post-crash trailing return) and wash out beyond the first two years.

Input	Value (snapshot)
Bitcoin held	842,138 BTC
Bitcoin spot price	$\approx$ \$63,400
Bitcoin 200-week moving average	$\approx$ \$63,900
Trailing 200-week realized CAGR ($\bar{g}$)	$\approx 30\%$
Derived fair price (F_0, eq. 6; user-adjustable in the tool)	$\approx$ \$107,900
Implied dislocation (d_0)	-0.53 (spot $\approx 41\%$ below fair)
Entry amplification target (eqs. 16, 18)	cycle leg ≈ 0.56; entry level ≈ 0.41 ($A_0 \approx 1.51\times$)
Bitcoin value of treasury	$\approx$ \$53.4 B
MSTR share price	$\approx$ \$94.86
Entry mNAV (this paper's definition)	$\approx 0.74\times$
Diluted shares (BPS 201,822 sats basis)	417,267,691
FDSO / entry Net BPS (issuer convention)	$\approx$ 391.6 M / $\approx$ 142,300 sats
Preferred liquidation preference (5 series)	$\approx$ \$15.3 B
STRC share of preferred (floating)	$\approx 68\%$ ($\approx$ \$10.4 B)
STRC price (snapshot)	\$92.32 ($\approx 7.7\%$ below par)
STRC coupon / implied required yield	12% / $\approx 13.0\%$
Convertible principal (6 tranches)	$\approx$ \$6.71 B
USD reserve	\$4.0 B ($\approx$ 28 months of coverage)
SOFR / policy rate	$\approx 3.62\%$
Counterparty/wrapper hurdle r_p (base)	3.5%/yr

Table 3: Base-case calibration snapshot.

Parameter	Meaning	Value
D	capitalization duration (years of excess BY priced)	2.0
β_0	kernel level at zero momentum	0.15
κ^+	kernel upside slope	1.10
κ^-	kernel downside slope	0.10
(w_r, w_m, w_c)	weights: realized / momentum / capacity	$(0.35, 0.50, 0.15)$
$\mathrm{hl}(\mathrm{BY}^{\mathrm{ema}})$	realized-yield EMA half-life	1.5 yr
$\mathrm{hl}(g^{\mathrm{ema}})$	momentum EMA half-life	1.0 yr
$\mathrm{BY}_0^{\mathrm{ema}}$	entry trailing realized yield (implied by the entry quote)	21.5%
g_0^{ema}	entry trailing momentum (post-crash)	-20%
h_0	entry structural hurdle	$\approx 11\%$
$m_{\min}$	premium floor	0.30
loop gain	bound on the premium–capacity fixed point	≈ 0.17

Table 4: Endogenous-premium parameters (momentum kernel, blend, and seeds).

8.2 Per-instrument capital structure

The aggregate convertible and preferred figures in Table 3 resolve into the individual instruments the engine carries. Table 5 lists the six convertible tranches—each resolved contingently at its holder put date under the assumed-diluted share count—and Table 6 the five preferred series. Of these, only STRC is issued further and repriced over the simulation; the four fixed preferred series and all six convertibles are held at their snapshot terms.

Tranche	Principal ($M)	Coupon	Conversion px	Investor put (yrs)
2028	1,010	0.625%	$183.19	1.2
2029	1,500	0%	$672.40	2.0
2030 A	800	0.625%	$149.77	2.2
2030 B	2,000	0%	$433.43	1.7
2031	604	0.875%	$232.72	2.2
2032	800	2.25%	$204.33	3.0
Total	6,714			

Table 5: Convertible note tranches at the snapshot (assumed-diluted basis). "Investor put" is the approximate time to the next holder put date, the model's resolution trigger; conversion prices are per MSTR share.

Series	Liq. pref. ($M)	Coupon	Rate type
STRF	1,284	10%	fixed
STRK	1,402	8%	fixed, convertible
STRD	1,402	10%	fixed
STRE	883	10%	fixed
STRC	10,369	12%	variable (modeled dynamically)
Total	15,340		

Table 6: Preferred series at the snapshot. The four fixed series carry a par-weighted blended coupon of $\approx$ 9.44%; STRC, the largest tranche, carries the floating coupon the engine resets each period off SOFR, coverage, and the migrating credit rating.

9 Default assumptions

The snapshot above fixes the *market* inputs. The model also ships with a set of *behavioral* and *policy* defaults that govern how the capital structure responds to those inputs—how aggressively the firm issues, how it defends STRC near par, how it manages its reserve, and which thesis channels are active. Table 7 states them in one place. They are modeling choices rather than observables, and because several of them move the answer materially, they are listed explicitly so a reader can see exactly what the base case assumes and change any of it in the tool.

A handful of these carry most of the sensitivity, and it is worth being plain about which. The *Bitcoin model*, the *kernel parameters* (D, κ^+, w_m; Table 4), and the *hurdle* dominate the headline: the Bitcoin model now drives both the backing accumulated and, through realized momentum, the premium granted along the way, the kernel sets how generously that momentum is capitalized, and the premium falls monotonically as r_p rises. The *stochastic-volatility overlay* enters the premium only through the volatility adder in the structural hurdle and the realized paths themselves; its residual effect on the personal valuation is modest. The *reserve policy* and the *defensive mNAV*

Assumption (symbol)	Default
Decision frame	
Holding horizon (τ)	10 years
Counterparty / wrapper hurdle (r_p)	3.5%/yr
Bitcoin model (your outlook)	DR 40 $\rightarrow$ 20
200-WMA fair-value anchor	on (F_0 from eq. 6; $h_d = 1.25$ yr)
Monte Carlo paths × step (n_{p})	20,000× quarterly
Reserve and financing waterfall	
USD reserve, start (R_0)	$4.0 B
Reserve target	floating (par-clearing inversion +3 mo; floor 12 mo; clamp 48 mo)
Reserve hard operating floor	3 months
Defensive common-ATM gate	enterprise parity (m_t^*: enterprise mNAV ≥ 1.0, live)
Below-par common issuance cap	3%/qtr of BTC value
Last-resort BTC sale (fund dividend) / per-qtr cap	on / 5% of stack
Structural mNAV floor ($m_{\min}$)	0.30
Accretive common ATM	packaged accretion $(1 + a_{\mathrm{tgt}} u_t)m > 1$; clamped at m_t^*
Discount buyback — trigger / slope / cap ($\eta_b, \bar{b}$)	below the effective ATM floor / 0.10 / 5%/qtr
Preferred retirement — trigger / pace cap	STRC $\leq$ 95% of par / 10% of tranche per qtr
STRC pricing (required yield $\rightarrow$ coupon $\rightarrow$ price)	
Baseline spread / floor (s_0, $\underline{s}$)	3.5% / 3.0%
Reserve-runway adder (k_R, τ_R)	12%, 2.0 yr (calibrated Aug 2026)
Coverage adder (k_B, $\underline{a}$)	16%, 0.30
Volatility adder (k_V, $\bar{\sigma}$)	9%, 0.55
Distress add — pause / forced BTC sale (d)	4.0% / 2.0%
Coupon step / willingness ceiling (Δc, $\bar{c}_{\mathrm{will}}$)	0.75%/qtr / 15%
Affordability horizon (ϕ)	0.6 yr
Price cap / sawtooth amplitude ($\bar{p}$, A_{saw})	1.003 / $2.2
Rate ceiling ($\bar{q}^{\,\mathrm{ceil}}$)	16%
Amplification (leverage policy)	
Countercyclical target (eq. 16)	A: 1.34× fair, clip [1.15, 2.10] (a: 0.33, $[0.20, 0.60]$)
Fractional-Kelly carry cap (eq. 18)	$\phi = 0.50$ on estimated carry
Absolute issuance ceiling / static fallback	A 3.00× (a 0.75) / A 1.60× (a 0.45)
STRC issuance pace cap	5%/qtr of BTC value
Finite-supply accumulation brakes	
Accessible supply (S_{acc})	17.3 M BTC
Float-depletion impact (k_{perm}, k_{temp})	1.25 / 2.0
Concentration pace exponent (p_{conc})	5.0
STRC carry margin	2%
Short policy rate (SOFR)	
Neutral / long-run (θ)	3.50%
Mean-reversion speed / vol (κ, σ_r)	0.40/yr / 0.90%
Jump mean / intensity (μ_J, λ_J)	+50 bps / $\sim$0.15 yr^{-1}
Policy-rate cap threshold (r^*)	4.5%
Bitcoin dynamics	
DR start→floor CAGRs ($c_0 \rightarrow c_\infty$)	$20 \rightarrow 10,\ 30 \rightarrow 15,\ 40 \rightarrow 20$
Initial residual vol (σ_0, decaying)	0.50
Stochastic-vol overlay	on
DR decay speed (λ, shared)	0.15/yr
GBM median CAGR (g) / fair value	10% / spot
Premium (mNAV) function	
Endogenous premium (kernel, blend, seeds)	Table 4
Thesis toggles	
On by default	jump rates, vol flywheel, credit migration, policy-rate cap
Off by default	joint fiscal shock, hard rate cap, Basel eligibility

Table 7: Behavioral and policy defaults—modeling choices, not market data.

τ (yr)	m_0	med m_T	$[p_{10}, p_{90}]$	stack	$m_0^\star$	P(beat)	med. excess ret.
3	0.83×	1.30×	[0.92, 1.66]	1.22 M	1.34×	96%	+22.5%
5	0.83×	1.26×	[0.90, 1.61]	1.65 M	1.35×	98%	+13.2%
10	0.83×	1.16×	[0.87, 1.46]	2.65 M	1.22×	98%	+5.3%
20	0.83×	1.13×	[0.92, 1.33]	4.14 M	0.91×	89%	+1.1%
30	0.83×	1.14×	[0.99, 1.31]	5.25 M	0.66×	18%	−0.4%

Table 8: Headline outputs across horizons (default profile, DR 40 → 20, hurdle 3.5%): the model's justified entry premium, the median and 10–90th percentile terminal premium, the median terminal coin stack, the personal rational mNAV, the probability of beating self-custody after the hurdle, and the median annualized Bitcoin-denominated return in excess of the hurdle.

floor together govern behavior in stress: a higher floor trades more common dilution for fewer forced coin sales, and a thicker reserve postpones both. The *STRC band targets* and the *par gate* set how often the preferred ATM is open, and therefore how much Bitcoin the flywheel can accumulate. The three *thesis toggles* that ship off—a joint fiscal-shock channel, a hard rate cap, and Basel collateral eligibility—are deliberately conservative; switching them on mainly reshapes the tails. Every figure in the next section is conditional on the values above.

10 Results

All figures below use the default profile: the endogenous momentum-calibrated premium of Section 5, the 200-WMA fair-value anchor of Section 3.1, the robust countercyclical amplification policy of equations (16) and (18) (feasible estimator, half-Kelly carry cap), jump-diffusion rates, the stochastic-volatility overlay, credit-rating migration, and the policy-rate cap active, the finite-supply brakes engaged, and the full firm-policy set of Section 4.5—collateral-capacity ATM, unified buyback threshold, preferred retirement, and the adaptive carry—switched on. The default Bitcoin outlook is the aggressive DR $40 \to 20$; the hurdle is 3.5%; runs use 20,000 quarterly paths (seed dispersion in the headline outputs is $\approx 0.4\%$); the belief cube, one horizon-blind simulation per cell, uses 10,000.

The recovery arc. The model's central story at the entry snapshot is a full cycle traced endogenously, and under the fair-value anchor the cycle starts from a spot price the model regards as $\approx 41\%$ below fair. At entry the justified premium is 0.83× against the observed 0.74×—the model, unlike the from-spot regime, no longer matches the quote and does not intend to: the gap *is* its assertion that the quote underprices a dislocated spot. Trailing momentum is negative and the carry is closed at entry (buybacks, preferred tenders, coin-funded reserve refills—the 2026 posture), but the dislocation convergence registers as realized momentum almost immediately: the carry gate reopens within the first year, the kernel re-rates the premium through par within the first year, and the accretive engine restarts. At the ten-year default the fair value—the entry price at which the holder is exactly indifferent to self-custody at the hurdle—is $\approx$ \$156 against the \$94.86 quote, a $\approx 64\%$ premium to the market. Table 8 carries the headline outputs across horizons.

Three features of the table deserve comment. The near-horizon returns are dominated by the *recovery trade*, now with both legs live: the spot-to-fair convergence and the premium re-rating compound to a median +22.5% per year in Bitcoin terms over three years—an entry-timing result, not a steady state, and one that decays quickly with horizon as the one-time convergence amortizes. The terminal-premium distribution stays wide and honest: paths on which momentum fades after the convergence settle back toward a discount (the p_{10} column), while sustained re-ratings carry the p_{90} to ≈ 1.5–$1.7\times$ at near and mid horizons. And the personal valuation itself falls below parity past two decades as the maturing drift decays through the kernel and the hurdle compounds: the model does not hand the wrapper a permanent premium; it makes the premium earn its keep every quarter. Read jointly, the columns say the wrapper is a *finite trade*: the personal rational mNAV peaks near five years and declines thereafter, and Section 7.1 formalizes the optimal exit this implies.

The Bitcoin process still dominates. At the ten-year default the personal rational mNAV under the four Bitcoin models is $0.67\times$ (GBM), $0.84\times$ (DR $20 \to 10$), $1.03\times$ (DR $30 \to 15$), and $1.22\times$ (DR $40 \to 20$). The spread is wide for a structural reason: your Bitcoin model drives *both* the coins behind each share *and*, through realized momentum and yield, the premium the market grants along the way. A slow-growth world is punished twice (weak accretion, discounted exit); a fast one is rewarded twice. Note that the shared fair-value anchor lifts every model by the same convergence leg, yet GBM still cannot escape a discount: a one-time $\approx 60\%$ catch-up cannot rescue a flat 10% drift against a 3.5% hurdle once the wrapper's machinery stays parked.

The wrapper hurdle. Holding DR $40 \to 20$ at ten years, the personal rational mNAV runs $1.90\times$ at a -1% hurdle (the wrapper-premium investor of Section 2), $1.72\times$ at 0%, $1.22\times$ at the 3.5% default, and $0.96\times$ at 6%: a sufficiently skeptical view of the structure compresses the premium sharply, and under the current snapshot the 6% skeptic lands just below parity—the recovery trade absorbs most, not all, of the skepticism.

What moves the answer: sensitivity and validation. A one-at-a-time tornado on the ten-year personal rational mNAV (base $1.22\times$) remains dominated by exactly the three empirically calibrated kernel parameters: the capitalization duration D (± 0.21 across 1.5–2.5), the upside slope κ^+ (± 0.20 across 0.8–1.4), and the momentum weight w_m (± 0.19 across 0.40–0.60). The fair-value anchor's own dials, by contrast, barely register at this horizon: the dislocation half-life contributes ± 0.013 across 0.75–2.0 years (the convergence completes well inside a decade regardless), and the trailing-growth input ± 0.02 across 22–32%, and the Kelly fraction ± 0.015 across 0.35–0.75 (its value shows in the stressed worlds of Section 4.3's leverage discussion, which the default-world tornado does not probe). The anchor moves the *level* of every cell; the decision *uncertainty* still lives in how the market prices momentum, which is where the empirical work was spent. Validation held-outs bracket the kernel honestly: replaying it over the realized 2023 recovery implies a premium of $1.26\times$ against the observed $\approx 1.35\times$, and through the 2022 trough $0.53\times$ against $\approx 0.6\times$; the known bias is that the linear-in-momentum kernel *under-prices manias* (2024Q4 implied $2.0\times$ against an observed 2.8–$4\times$), so the model is conservative precisely in the states where a holder would be selling. Stability is verified rather than hoped for: the premium–capacity loop gain is bounded at ≈ 0.17, seed dispersion is 0.4%, and pushing the kernel to $D = 3$, $\kappa^+ = 2$ produces finite, sane premia throughout.

STRC availability, retirement, and the tails. Because the STRC price is derived from the simulated reserve, coverage, and volatility, both *when* and *how much* remain outputs. Over the ten-year default the STRC ATM is open 45% of the horizon—closed through the brief bear phase, reopening with the convergence—and the median path accumulates coins equal to 86% of the starting stack through the STRC channel (66–104% across the 10–90th percentiles), with STRC spending 86% of the horizon inside \$97–101. Discounted preferred tenders and sub-floor common buybacks continue to convert discount into Bitcoin-per-share along the way, on the same unified-threshold logic as before, with buyback-funded coin sales amounting to a mid-single-digit share of the starting stack on the median path. The tails are thinner still under the recalibrated reserve policy, because the floating target holds coverage at the confidence requirement before the convertible puts arrive: a 3.1% probability of at least one forced coin sale at a put, a 0.03% probability of a last-resort dividend-driven coin sale, and a 0.15% probability of a dividend pause. The reserve itself floats rather than saw-tooths: the median terminal balance sits at ≈ 26 months of coverage—between the twelve-month floor and the ≈ 38-month entry requirement, easing as seasoning and coverage improve—and terminal amplification settles near $a \approx 0.16$ ($A \approx 1.19\times$), the structure de-levering passively as the denominator compounds. (The near-disappearance of dividend-driven sales and the reduction in put-date sales are consequences of the dynamic parity gate: on healthy, de-levered paths the defensive window reaches far deeper than the old fixed floor, so equity defense replaces coin sales; the small rise in pauses is its mirror on highly levered paths, where the gate correctly refuses dilution that would no longer raise net backing.)

The belief cube. The valuation depends on the Bitcoin model and the hurdle; the holding horizon is *optimized out* using the machinery of Section 7.1, so every cell reports the value of the trade run to its own best exit. Figure 2 presents the 4×4 grid in the three capital-structure framings of Section 2.1—equity, enterprise, and net-assets—each cell showing *today's* rational mNAV under the best pre-committed strategy, alongside that strategy's horizon τ^*; Tables 9–11 tabulate the three panels. To be precise about what a cell is: it answers "if I buy *today* and run the hold-τ^*-then-switch-to-self-custody strategy, what entry price makes me exactly indifferent to self-custody at my hurdle?" Both the denominator (today's coin value per share) and the discounting (back to today, at the hurdle) anchor the number to the present. The optimization over τ merely selects *which* pre-committed strategy to price—each candidate horizon yields a today-valuation, and τ^* is the one where that today-valuation is highest. The cell is an entry valuation, not a forecast of the premium prevailing at the exit date (that object is the terminal-premium path of Section 7). Two patterns organize the cube. The growth assumption still dominates the *level*: the aggressive models sit far above parity, and even flat GBM ekes out $\approx 1.05\times$ at a zero hurdle by harvesting the spot-to-fair convergence over a short trade. But the hurdle sets the *clock*. At a 0% hurdle holding is free, so τ^* runs to the edge of the grid (25–30 years) in any world whose carry stays positive—the wrapper becomes a hold-forever machine. The moment the hurdle is positive, the finite trade reappears and the optimal exit collapses to 2–10 years, shortest where beliefs are weakest and the hurdle steepest: the marginal year must now *pay* for itself, and only the front-loaded convergence-and-re-rating years reliably do. The enterprise and net-assets framings shift every cell upward by construction (senior claims are $\approx 39\%$ of the coin value at entry), but the shape—growth sets the level, the hurdle sets the clock—is identical in all three.

	Wrapper-risk hurdle			
Your Bitcoin model	0%	2%	4%	6%
GBM (flat 10%)	1.05× / 3	0.99× / 3	0.93× / 3	0.89× / 2
DR $20 \to 10$	1.20× / 5	1.12× / 3	1.06× / 3	1.01× / 2
DR $30 \to 15$	1.54× / 30	1.28× / 4	1.19× / 3	1.13× / 3
DR $40 \to 20$	1.85× / 30	1.46× / 5	1.33× / 4	1.25× / 3

Table 9: The equity panel of the belief cube: today's rational equity mNAV under the optimal pre-committed hold / τ^* (years).

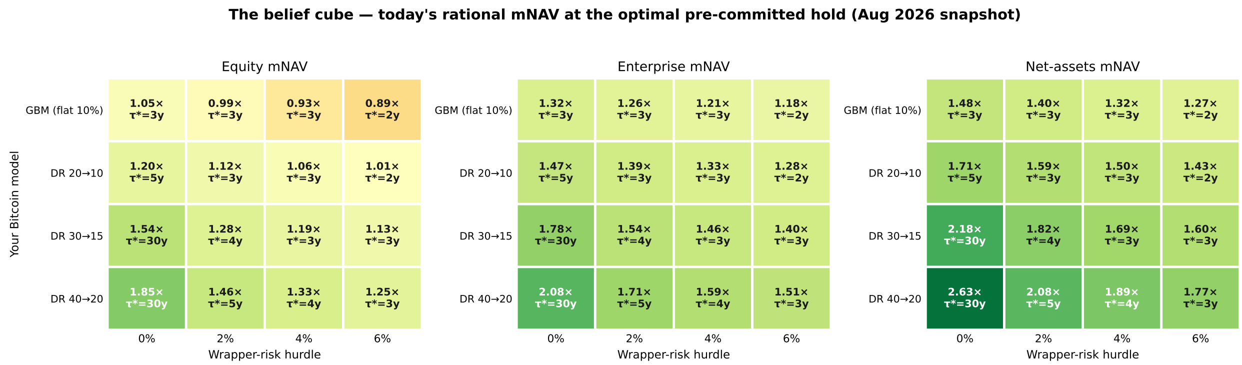

Figure 2: The belief cube. Panels are the three capital-structure framings of the personal rational mNAV (equity, enterprise, net-assets); within each panel, rows are *your* Bitcoin model and columns the wrapper-risk hurdle. Each cell shows *today's* rational mNAV—the entry premium at which buying MSTR now and holding for the optimal pre-committed horizon is exactly indifferent to self-custody at the stated hurdle—followed by that horizon τ^* in years. Green denotes a premium to net asset value, red a discount; the dividing line is 1.0×. A zero hurdle makes holding free and pushes τ^* to the grid edge; any positive hurdle restores the finite trade.

	Wrapper-risk hurdle			
Your Bitcoin model	0%	2%	4%	6%
GBM (flat 10%)	1.32× / 3	1.26× / 3	1.21× / 3	1.18× / 2
DR 20 → 10	1.47× / 5	1.39× / 3	1.33× / 3	1.28× / 2
DR 30 → 15	1.78× / 30	1.54× / 4	1.46× / 3	1.40× / 3
DR 40 → 20	2.08× / 30	1.71× / 5	1.59× / 4	1.51× / 3

Table 10: The enterprise panel of the belief cube: today's rational enterprise mNAV under the optimal pre-committed hold / τ^* (years).

	Wrapper-risk hurdle			
Your Bitcoin model	0%	2%	4%	6%
GBM (flat 10%)	1.48× / 3	1.40× / 3	1.32× / 3	1.27× / 2
DR 20 → 10	1.71× / 5	1.59× / 3	1.50× / 3	1.43× / 2
DR 30 → 15	2.18× / 30	1.82× / 4	1.69× / 3	1.60× / 3
DR 40 → 20	2.63× / 30	2.08× / 5	1.89× / 4	1.77× / 3

Table 11: The net-assets panel of the belief cube: today's rational net-assets mNAV under the optimal pre-committed hold / τ^* (years).

11 Limitations

The reserve target inverts a spread model with an unquantifiable residual. The floating reserve target is only as good as the required-yield decomposition it inverts. That model attributes STRC's discount to measurable state—coverage, runway, volatility, seasoning—but the calibration episode itself showed a residual the state variables do not carry: market sentiment extended confidence in early 2026 on six months of coverage and withdrew it later on more. The par-clearing requirement the model computes is therefore an estimate whose error is exactly that sentiment term, and the requirement should be expected to drift against any fixed calibration as confidence in the asset class moves. Relatedly, the amplification dials are entered in the issuer's reserve-embedding units and converted to the internal claims ratio *at the entry reserve share*; a reserve that departs far from its entry scale shifts the effective schedule slightly from what the labels state.

It is a capital-structure model, not a macroeconomic one. The engine has exactly two exogenous drivers—the short policy rate (SOFR) and the Bitcoin price—and everything else is the liability stack propagating off them. Broad macroeconomic conditions such as global liquidity, the growth of money supply, national debt and deficits, and reserve-currency status are *not* state variables and are never modeled directly. They enter the analysis only to the extent that the user believes they are already expressed in those two paths: a view that abundant liquidity and monetary debasement lie ahead is encoded by choosing a more bullish Bitcoin process and a more constrained (capped) policy-rate process, not by any separate macro module. The model inherits whatever macro thesis the inputs imply and no more.

The model does not assume fiscal dominance—the user selects it. This is worth stating plainly, because the default profile leans toward a fiscal-aware world and could be mistaken for a built-in assumption. It is not. Whether the simulated future resembles *fiscal* dominance (rates pinned below inflation while hard assets appreciate) or *monetary* dominance (a central bank free to hold real rates positive and anchor inflation) is determined entirely by the settings. Jump-diffusion rates with the policy-rate cap engaged and a aggressive DR Bitcoin trend describe a fiscal-dominance scenario; Vasicek rates reverting to a stable mean with the cap off, paired with a GBM or conservative-DR Bitcoin view, describe an orthodox monetary-dominance world in which the wrapper usually loses to self-custody. The joint-fiscal-shock and yield-curve-control toggles sharpen the fiscal case further, and all of them ship *off* except the policy-rate cap, which is dormant at today's rates. The framework therefore takes no stance on the debate; it prices whichever side the user assumes.

The fair-value anchor is an estimate, not an oracle. Anchoring the trends to a de-lagged 200-week average removes the dependence on one day's spot, but it replaces it with an estimate that carries its own assumptions: the de-lag identity treats the window's realized growth as the growth of a clean exponential trend, the WMA itself differs modestly across data sources, and the dislocation half-life is the one hand-set dial in the mechanism (though the sensitivity analysis shows both anchor dials wash out of decade-horizon conclusions). More consequentially, the anchor is a large *level* lever precisely because the momentum kernel prices the convergence it asserts: the near-horizon returns of Section 10 are entry-timing results, conditional on the fair-value thesis being right, and a reader who rejects that thesis should read the from-spot configuration instead. The discipline that keeps the mechanism honest is its symmetry—at a mania top the same formula marks the entire forward arc down—and it should only be trusted by a user prepared to accept that verdict in both directions.

The premium is a fitted behavioral rule, not an equilibrium. The momentum kernel was adopted precisely because market sentiment toward Bitcoin swings violently and largely *with* the price: the endogenous premium now re-rates and compresses along each simulated path, entering and leaving regimes as the realized trend dictates, which repairs the static-outlook limitation of earlier formulations. But the repair has its own honest edges. *First*, the kernel is a fitted description of how this market has priced momentum across two cycles, not a no-arbitrage equilibrium; nothing guarantees the mapping is stable out of sample, and the sensitivity analysis shows the verdict leans exactly on its three fitted parameters. *Second*, the fit is deliberately conservative at the extremes: the linear-in-momentum upside under-prices manias (the 2024Q4 peak implies $2.0\times$ against an observed 2.8–$4\times$), so blow-off states are marked below their historical prints. *Third*, the calibration data are coarse—quarterly premium observations assembled from public trackers with ± 0.15 of definitional noise across sources—adequate for the directional findings that shaped the functional form, but not for fine parameter precision. *Fourth*, the capitalization duration D is held fixed at its compression-era back-out; a regime-dependent D (the euphoria multiple) exists in the engine but is disabled to avoid double-counting momentum, so the model expresses bull re-ratings through the kernel alone.

The STRC price is derived; its calibration is assumed. Unlike an earlier version of this model, the price-stability bands of Section 4.3 are no longer an input: STRC's price, the fraction of time it trades near par, and its ATM-open duty cycle are *emergent outputs* of the pricing identity, driven by the simulated reserve runway, senior coverage, and volatility. What remains a *modeling assumption* is the calibration of the required-yield spread—the inherent baseline and the three adder coefficients—together with the coupon affordability logic, its 15% willingness ceiling (adopted as a reasonable figure, not one attributable to management), the small-step reset cadence, and the ex-dividend sawtooth amplitude. These were set against the instrument's short but informative history: at par on ≈ 2.5 years of reserve coverage in early 2026, near \$97 on roughly six months, and near \$74 on roughly ten months amid a Bitcoin drawdown—so the model reproduces both the level and the *non-monotonicity* (more coverage, lower price, when Bitcoin fell). One gap is acknowledged: the spring's resilience—STRC holding near par on only six months of coverage—is read as confidence the market later withdrew rather than a stable equilibrium, so the structural terms price it somewhat below the observed \$97–98. What still lies in the tail is an outright *breakdown* of the stabilization mechanism (demand evaporates, the reserve genuinely exhausts, and STRC re-rates to a distressed yield), which the distress term captures only on actual reserve exhaustion. These coefficients will need revision as more of the instrument's behavior is observed.

Two technical simplifications. Nominal and real rates are not separated, so a pure debasement that lifts Bitcoin and nominal rates together is captured only approximately, through the optional rate–Bitcoin correlation and joint shock rather than through an explicit inflation state. And because every input to the endogenous premium is a physical quantity, the exit value is exactly independent of the entry price: the fair value of (32) is a true one-pass solution, with no fixed-point residual.

It is conditional by design. The framework takes the user's beliefs as inputs and returns the internally consistent premium. Its purpose is to make a valuation follow rigorously from a worldview, not to tell the user which worldview is correct—the two largest uncertainties, which Bitcoin process governs the future and how much the equity wrapper truly deserves to be discounted, remain the user's to judge.

Long horizons warrant added caution. Results for holding periods beyond approximately thirty years should be read as directional rather than precise. Across such spans several of the model's dynamics—the Bitcoin price trajectory, the structural-premium path, the credit-spread evolution, and the accumulation asymptote—grow increasingly sensitive to parameter choices, precisely where the framework extrapolates furthest beyond the data that inform it. The same horizon encompasses Bitcoin's transition from a block-subsidy-dominated to a fee-dominated security budget, multiple plausible monetary and interest-rate regimes, an evolving regulatory landscape, and successive changes in Strategy's capital structure, management, and strategy, none of which the framework conditions upon. The asymptotic outcomes the model produces at such horizons—in which a single entity approaches a substantial share of accessible supply—would themselves likely provoke second-order responses in network governance, regulation, and market structure that lie beyond the model's scope. Outputs past this horizon are therefore best interpreted as illustrative of qualitative tendencies rather than as quantitative projections.

12 Conclusion

Denominating value in Bitcoin turns the question of a treasury company's worth into a clean comparison against self-custody and exposes the premium as a function of three beliefs rather than a number to be argued in the abstract. The rational mNAV follows from one's Bitcoin outlook, one's trust in the corporate wrapper, and one's holding horizon. Much of the public disagreement over Strategy's fair premium is, on this reading, a disagreement over these inputs conducted without making them explicit. The contribution here is a transparent engine that makes the premium the consequence of stated assumptions, so that debate can move to where it belongs: the assumptions themselves.

Disclaimer. This white paper and the accompanying tool are provided for educational and informational purposes only. They do not constitute financial, investment, tax, or legal advice, nor a recommendation, offer, or solicitation to buy or sell any security or asset, including MSTR shares, Strategy's preferred securities, or Bitcoin. Every figure herein is the conditional output of user-supplied assumptions and stylized parameters, is not a forecast, and may differ materially from realized outcomes; models are simplifications, calibrations can be wrong, and relationships fitted to history may not persist. All market data reflect a snapshot in time and will be stale by the time of reading; the document is versioned and dated for that reason, and earlier versions may circulate. **The author may hold positions in securities discussed herein, including MSTR common stock and Strategy preferred stocks, and such positions may change without notice or update to this document.** Nothing here has been reviewed or endorsed by Strategy or any regulator. Markets involve risk of loss, including total loss. Readers should conduct their own due diligence and consult a qualified, licensed professional before making any investment decision.

License. © 2026 @rationalmnav. This document is licensed under Creative Commons Attribution–NonCommercial–NoDerivatives 4.0 International (CC BY-NC-ND 4.0): you may share and repost it freely *with attribution to @rationalmnav and rationalmnav.com*; commercial use and modified or derivative versions are not permitted. Full license: `creativecommons.org/licenses/by-nc-nd/4.0`.

A Notation and symbols

This appendix collects the symbols used in the paper. Time-indexed quantities carry a subscript t, with 0 denoting the present (the calibration snapshot) and T the value at the holding horizon. Constants set as modeling defaults are listed with their base-case values in Table 7.

Symbol	Meaning
State: prices, holdings, shares	
B_t	USD Bitcoin spot price; B_0 today, B_T at the horizon.
H_t	Bitcoin held by the firm, in coins.
N_t	diluted common shares outstanding.
P_t	MSTR common share price.
b_t	Bitcoin backing per share, $b_t = H_t/N_t$.
Premium and valuation	
m_t	mNAV, the market premium: equity market capitalization P_tN_t over gross coin value H_tB_t.
$m_0^\star$	rational (fair) mNAV: the premium at which MSTR beats self-custody with probability one-half.
m^{eq}, m^{EV}, m^{net}	the three mNAV conventions of Section 2.1: equity-to-gross-coin (used here), enterprise-value, and net-asset.
m_t^*	enterprise-parity gate: the equity mNAV at which enterprise mNAV = 1.0; defensive and packaged issuance run only at or above it, buybacks below the packaged boundary (never in a defense quarter).
$m_{\min}$	structural floor on the premium function.
m_{acc}	premium above which the offensive (accretive) common ATM activates.
G	Bitcoin-denominated gross multiple of the position over the horizon.
v_T	exit value of one share in Bitcoin, $v_T = b_T\, m_T$.
V_0	fair MSTR price today, in dollars.
$\text{med}(\cdot)$	cross-path median.
r_p	annual counterparty/wrapper hurdle.
τ	holding horizon, in years.
mNAV-convention snapshot (Section 2.1)	
$P, N, N_{\text{xc}}, H, B$	snapshot share price, diluted shares, ex-convert shares, coins held, and spot price.
D	outstanding convertible principal.
L	aggregate preferred liquidation preference (snapshot); L_t below is its time-varying analogue.
π	(Section 2.1 only) the investor's assigned probability of a forced wind-down.
Short policy rate (SOFR) dynamics	

Symbol	Meaning
r_t	short policy rate (SOFR); r_0 today.
θ	long-run neutral rate.
κ	rate mean-reversion speed.
σ_r	rate volatility.
μ_J, σ_J	mean and volatility of rate jumps; λ_J the jump intensity (per year); Q_t the Poisson jump-counting process.
r^*	policy-rate cap threshold (the level above which the fiscal-dominance constraint pulls the rate down).
λ_{fd}	fiscal-dominance constraint strength (per-year downward pull on the excess above r^*).
W_t^r, W_t^B	Brownian drivers for the rate and for Bitcoin.
Bitcoin dynamics	
G_t	deterministic Bitcoin trend under the selected model (GBM line $B_0(1+g)^t$, or the DR integral); $G_0 = B_0$ (spot).
c_0, c_∞, λ	DR start and floor CAGR and shared decay ($\lambda = 0.15$): $c(t) = c_\infty + (c_0 - c_\infty)e^{-\lambda t}$.
μ, σ	Bitcoin drift and volatility governing reversion toward the fair-value line.
σ_t	Bitcoin (residual) volatility; σ_0 initial, $\bar{\sigma}_t$ the trailing anchor, σ_{ref} a reference level.
g	GBM median CAGR (10%); all four model trends compound from F_0.
F_0, d_0, κ_d	200-WMA fair price; entry dislocation $\ln(B_0/F_0)$; its decay rate ($\ln 2/h_d$).
g	GBM median CAGR: the rate at which the GBM median path is pinned to compound (10%/yr default).
μ_B^J, σ_B^J	mean and volatility of the Bitcoin jump under the optional joint fiscal-stress shock.
Capital structure and waterfall	
S_t	preferred notional (aggregate liquidation preference).
D_t	convertible principal outstanding.
R_t	USD reserve; R_0 today.
L_t	senior liabilities ranking ahead of the common: preferred par, accrued arrears, and convertible principal.
a_t	amplification, $a_t = (S_t + D_t)/(H_t B_t)$; $a_{\text{tgt}}, a_{\max}, a_{\min}$ its target, ceiling, and floor.
η	accretive-issuance slope; $\bar{a}$ the market-depth (issuance-pace) cap.
η_b	discount-buyback slope; $\bar{b}$ its per-quarter cap; buybacks fire below the effective ATM floor.
ΔN	common shares issued in a step.
STRC pricing (required yield $\to$ coupon $\to$ price)	

Symbol	Meaning
P_t^{STRC}	STRC price; $\Pi = \$100$ its par.
q_t^{STRC}	required (market-clearing) STRC yield = SOFR + spread.
c_t	offered STRC coupon (state); ratchets toward $\min(q_t^{\text{STRC}}, c_t^{\max})$ in steps $\pm\Delta c$.
s_0, $\underline{s}$	inherent baseline spread over SOFR and the spread floor.
k_R, τ_R	reserve-runway adder coefficient and decay scale; $w_t = R_t/(\bar{q}_t S_t)$ the runway (yrs).
k_B, $\underline{a}$	senior-through-STRC coverage adder coefficient and its threshold amplification.
k_V, $\bar{\sigma}$	volatility adder coefficient and reference vol; d_t the distress add (pause 4% / forced BTC sale 2%).
$c_t^{\max}$, $\bar{c}_{\text{will}}$, ϕ	affordability-derived coupon cap, willingness ceiling, affordability horizon.
$\bar{p}$, A_{saw}	above-par price cap and ex-dividend sawtooth amplitude.
$\bar{q}^{\,\text{ceil}}$	hard rate ceiling above which STRC cannot clear even at par.
ε_t	standard normal shock.
Premium (mNAV) function	
ρ_t	net-backing ratio, $\rho_t = (H_t B_t - L_t + R_t)/(H_t B_t)$; $\widehat{\rho}_t = \text{clip}(\rho_t, 0, 1)$.
h_t	structural hurdle: the structural required STRC yield (no runway/distress adders).
D	capitalization duration: years of excess Bitcoin Yield the market prices.
BY_t^{ema}, g_t^{ema}	trailing realized Bitcoin-per-share yield; trailing realized Bitcoin momentum (EMAs).
$\varphi(g)$	momentum kernel, $\beta_0 + \kappa^+ \max(g - g_0, 0) - \kappa^- \max(-g, 0)$.
u_t	collateral-capacity unlock fraction (carry open and capacity-bound).
a_{tgt}	robust amplification target: min of the cycle schedule $a_{\text{fair}} e^{-\hat{x}}$ (clipped) and the ϕ-Kelly carry cap.
$\hat{x}_t$, $\hat{g}_t$, $\hat{\sigma}_t$	feasible estimates (trailing observables) of the fair-line deviation, growth, and volatility.
Simulation	
n_{p}	number of Monte Carlo paths.
Δ	time step (one quarter).