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

June 2026

This is the first issuance of this white paper. The author welcomes feedback on ways to improve it and reserves the right to revise it in the future 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 power-law trend with mean-reverting residual, a geometric Brownian motion, or a diminishing-returns trend), each anchored to its own estimate of *today's fair value*—the power law to its formula, the other two to a lag-corrected 200-week moving average—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 soft financial-repression ceiling and optional joint fiscal-stress shocks that re-rate Bitcoin in the same event, (iii) 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, and a migrating credit rating, the STRC price is held within a band by an active par-defense process—at a persistent modest discount to par in the base case rather than pinned to it—so that its at-the-market program issues only at or above par, and thus only part of the time; 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 a justified premium (mNAV) anchored to the net-backing ratio—Bitcoin value less the senior claims ranking ahead of the common—and carrying a convexity component, a convex upside "flywheel" that lets a volatility spike balloon the premium when the firm is strong, and a mean-reverting dislocation. 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*. We show that the rational premium is not a single number but a function of beliefs: across plausible combinations of Bitcoin outlook, wrapper-risk hurdle, and holding horizon it ranges from a discount near 0.2× to a premium above 2×. 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 own 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

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

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). Everything that follows is machinery for simulating the joint distribution of b_T and m_T.

2.1 Why mNAV should be common-equity market cap over total Bitcoin holdings

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, 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

$$m^{\mathrm{eq}} = \frac{PN}{HB}, \qquad m^{\mathrm{EV}} = \frac{PN + D + L}{HB}, \qquad m^{\mathrm{net}} = \frac{PN}{HB - D - L}. \tag{4}$$

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. 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 (already in the share count) or repayment/refinancing (not in the share count), so a framework that dilutes for conversion must not also subtract the principal. 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 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 the excess of Bitcoin value over the senior claims ranking ahead of the common, 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

Each price model is anchored to its own estimate of *today's fair value* F_0, and the gap between spot and F_0 drives a catch-up. This removes a hidden inconsistency: the power law treats today's price as a discount to a trend line, so it must not be compared against models that silently assume spot *is* fair value. We give each model a fair value computed the way that model would compute one.

Power law (default). Following the long-observed log-log linearity of Bitcoin's price against time, the trend is $F_t = A\, d_t^n$ with d_t the number of days since the genesis block, calibrated to $A = 1.6 \times 10^{-17}$ and $n = 5.77$; today's fair value F_0 is read straight off this Santostasi line. The log-deviation $x_t \equiv \ln(B_t/F_t)$ follows a mean-reverting Ornstein–Uhlenbeck process around the line,

$$dx_t = -\kappa_B\, x_t\, dt + \nu(t)\, dW_t^B, \qquad \nu(t) = \sigma(t)\sqrt{2\kappa_B}, \tag{5}$$

with a relative volatility that decays as the network matures, $\sigma(t) = \sigma_0\,(t/t_0)^{-\gamma}$ ($\sigma_0 = 0.50$, $\gamma = 1$). The residual is initialized at today's deviation $x_0 = \ln(B_0/F_0)$, so a price below the line is a discount expected to be recovered—a structurally bullish, low-ruin prior. $B_t = F_t\, e^{x_t}$.

Fair value from the 200-week moving average. The diminishing-returns and GBM models take their fair value from an empirical anchor: the 200-week moving average, WMA_{200}. But the moving average is a *floor*, not a central value—a trailing average of a rising series sits below the series' current level for a purely mechanical reason. We therefore advance it by a model-specific *lag factor*, $F_0 = \phi \cdot \mathrm{WMA}_{200}$, where ϕ equals the current trend value divided by the trend averaged over the trailing 200-week window. A faster-growing trend leaves its trailing average further behind, so the factor is larger for more bullish models: $\phi \approx 1.75$ under the diminishing-returns growth rate and $\phi \approx 1.43$ under the GBM mean. With $\mathrm{WMA}_{200} \approx \$61\text{k}$ these give fair values near \$107k and \$87k respectively—both well above spot, and both below the power-law line.

Geometric Brownian motion. The skeptic's model: a random walk with constant log-drift μ and volatility σ (defaults 0.20 and 0.55),

$$d\ln B_t = \left(\mu - \tfrac{1}{2}\sigma^2\right)dt + \sigma\, dW_t^B + \delta_t\, dt, \tag{6}$$

to which we add a purely deterministic catch-up drift $\delta_t = -\kappa_B\, x_0\, e^{-\kappa_B t}$ that carries the median from spot up to the fair-value line $F_0\, e^{(\mu - \frac{1}{2}\sigma^2)t}$ over the reversion window while leaving the random-walk dispersion (and its fat left tail) intact. We use the *mean* lag factor rather than the median: an investor skeptical enough to consult this tool is unlikely to also believe Bitcoin's central tendency is its high-volatility median, so the mean keeps the skeptic honest without tipping into pessimism.

Diminishing returns. A conservative middle ground in which the trend grows from F_0 at a compound rate decaying toward a floor, $c(t) = c_\infty + (c_0 - c_\infty)e^{-\lambda t}$, giving $G_t = F_0 \exp\!\left[c_\infty t + \frac{c_0 - c_\infty}{\lambda}(1 - e^{-\lambda t})\right]$ around which the same OU residual fluctuates ($c_0 = 0.30$, $c_\infty = 0.08$, $\lambda = 0.10$), with $x_0 = \ln(B_0/F_0)$ so spot catches up to the line exactly as in the power law.

Stochastic volatility (auto-calibrated overlay, on by default). Each process is overlaid with a regime-switching volatility, switched on with no parameter for the user to choose. The instantaneous volatility mean-reverts in log space to the volatility the model already uses—the age-decaying $\sigma(t)$ of the power law, or the constant of the GBM—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 $\pm 45\%$ in log terms ($\xi_v \approx 1.1$), so volatility occasionally runs to $2.5\times$ or down to $0.4\times$ 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, under bullish Bitcoin views, net favorable.

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:

$$\text{Vasicek:} \quad dr_t = \kappa(\theta - r_t)dt + \sigma_r\, dW_t^r, \tag{7}$$

$$\text{CIR:} \quad dr_t = \kappa(\theta - r_t)dt + \sigma_r \sqrt{r_t}\, dW_t^r, \tag{8}$$

$$\text{Random walk:} \quad dr_t = \sigma_r\, dW_t^r, \tag{9}$$

$$\text{Jump-diffusion (default):} \quad dr_t = \kappa(\theta - r_t)dt + \sigma_r\, dW_t^r + J_t\, dN_t, \tag{10}$$

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 N_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.

Soft financial repression (on by default). Rather than an explicit peg, the default policy overlay is a gentle, asymmetric restoring force: rates float freely below a critical value r^*, but above it face a downward pull proportional to the excess, $dr_t \mathrel{-}= \lambda_{\text{rep}} \max(r_t - r^*, 0)\, dt$. This captures the aggregate of regulatory demand, central-bank purchases, issuance management, and tolerated inflation without committing to a mechanism, and—crucially—it is *dormant* below r^*, so at today's sub-r^* rate it does not misstate the present. The default $r^* = 4.5\%$—roughly a percentage point above the 3.5% neutral, a mildly restrictive policy stance a fiscally-dominant Fed is reluctant to exceed once debt service binds—together with $\lambda_{\text{rep}} = 1.0$/yr 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,

$$a_t = \frac{S_t + D_t}{H_t B_t}, \tag{11}$$

where S_t is preferred notional (liquidation preference) and D_t is convertible principal. New issuance targets a chosen amplification a_{tgt}, and a ceiling $a_{\max}$ throttles it; when a drawdown pushes a_t above the ceiling the firm's financing options narrow sharply (below). The base-case assumptions are a target of $a_{\text{tgt}} = 0.45$ and a high boundary of $a_{\max} = 0.60$, against an entry level of ≈ 0.42 implied by the calibration snapshot. There is no hard *low* boundary: in a sustained Bitcoin advance the denominator grows while the senior stack does not, so amplification drifts down toward zero on its own, and the firm re-levers (issuing preferred toward the target) only as the ratio falls back below a_{tgt}. The practical operating band is therefore roughly 0 (deep bull) to 0.60 (stress ceiling), centered on the 0.45 target.

4.2 Endogenous preferred dividend

The preferred stack is split into a fixed-coupon tranche and a floating tranche (STRC). The *required* (par-clearing) STRC rate is SOFR plus a credit spread,

$$q_t^{\text{STRC}} = r_t + \max\big(s_0 + s_1\, a_t - \rho_{\text{cr}}\, \theta_t - \beta_{\text{coll}},\ \underline{s}\big), \tag{12}$$

with a base spread $s_0 = 4.25\%$ over SOFR (calibrated so the par-clearing dividend is $\approx 11.5\%$ at the snapshot), a leverage component $s_1 = 8.0\%$, and a spread floor $\underline{s} = 2.0\%$. The remaining terms model *credit-quality migration* (on by default): a per-path rating state $\theta_t \in [-1, 1]$ seasons upward while the firm pays in full and keeps its reserve, $\dot{\theta} = \alpha(1 - \theta)$ when healthy, and takes a sharp, sticky *downgrade* $\theta \leftarrow \theta - h$ on any dividend pause—agencies' slow-to-upgrade, quick-to-downgrade asymmetry. A fully seasoned rating tightens the spread by up to $\rho_{\text{cr}} = 2.5\%$; $\alpha = 0.35/\text{yr}$, $h = 0.30$. An optional collateral-eligibility regime (a Basel recognition of Bitcoin, off by default) adds a further permanent tightening β_{coll} from a chosen year. This makes financing a *credit cycle*: good behavior compounds into cheaper capital, while a stress-driven pause raises the floating cost precisely when the firm can least afford it. With a constant fixed-tranche coupon q^{fix} on notional S^{fix}, the blended preferred coupon is

$$\bar{q}_t = \frac{S^{\text{fix}}\, q^{\text{fix}} + (S_t - S^{\text{fix}})\, q_t^{\text{paid}}}{S_t}, \tag{13}$$

where q_t^{paid} is the dividend *actually paid* (defined below), endogenous and drifting upward over time as the floating share of the stack grows (entry $\approx 11\%$, rising toward $\approx 13\%$ at ten years in the base case).

4.3 STRC price stability 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 par-defense process below: the spread enters through the leverage and credit terms (and the stress pushes), and the reset lag is the mean-reversion of the price toward par. 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 pinned by the dividend/required-yield ratio. A perpetual's price is par scaled by the ratio of the dividend it pays to the yield the market demands,

$$P_t^{\mathrm{STRC}} \;=\; \Pi \cdot \frac{q_t^{\mathrm{paid}}}{q_t^{\mathrm{STRC}}}, \qquad \Pi = \$100. \tag{14}$$

The snapshot confirms the identity exactly: a paid dividend of 11.50% against a required (effective) yield of 13.17% implies $100 \times 11.50/13.17 = \$87.3$, matching the observed \$87.31. So we need not simulate a separate price process; the price is the gap between what Strategy pays and what the market requires.

An active par defense. We model the price deviation from par, $\delta_t \equiv P_t^{\mathrm{STRC}}/\Pi - 1$, directly. Management adjusts the dividend (and, as needed, other tools) to keep the price close to par, but the market prices a persistent modest discount, so δ_t mean-reverts not to zero but to a small negative equilibrium δ_{cal}, pushed further below par in stress. We use an exact Ornstein–Uhlenbeck step (stable at any reversion speed),

$$\delta_{t+\Delta} \;=\; \bar{\delta}_t + (\delta_t - \bar{\delta}_t)e^{-\kappa_\delta \Delta} + \sigma_\delta \sqrt{1 - e^{-2\kappa_\delta \Delta}}\, \varepsilon_t, \qquad \bar{\delta}_t = \delta_{\mathrm{cal}} - \frac{u_t}{\kappa_\delta}, \tag{15}$$

with par-defense speed $\kappa_\delta = 12/\mathrm{yr}$ (a ≈3-week half-life, matching the roughly twice-monthly dividend resets), a calm equilibrium $\delta_{\mathrm{cal}} = -2.4\%$ (an ≈ \$97.6 center), and a calm-regime noise $\sigma_\delta = 2.95\%$. The stress term $u_t \geq 0$ is the source of transient *sub-par* excursions and has two channels,

$$u_t \;=\; \gamma_v \left(\max(\sigma_t^B / \bar{\sigma}_t^B - 1, 0) \right)^2 \;+\; \gamma_r \max(\Delta r_t, 0), \tag{16}$$

a *convex* response to Bitcoin volatility running above its anchor (risk-off demand) and a transient response to a positive rate shock (a SOFR scare makes the fixed-ish payout less attractive). The convexity is deliberate: ordinary volatility wiggles barely move the price off par, so it hugs par on a normal day, but a genuine volatility *spike* punches it below the band. Both pushes fade as management restores par.

Calibration to the intended stability. The free constants are set so the long-run price distribution matches management's stated intentions while reflecting that the security has in practice traded at a discount: STRC sits within \$97–101 about 40% of the time and within \$90–102 about 90% of the time, with the $\approx$ 10% of time outside the wider band concentrated in genuine stress and skewed well to the downside ($\gamma_v = 0.35$, $\gamma_r = 5.0$). Today's \$87 print sits below even the equilibrium discount and is treated as an *abnormal starting excursion* that reverts toward the $\approx$ \$97.6 center, not as normal performance; the engine initializes δ_0 at the observed price and lets the par defense pull it back. These bands are an assumption adopted for modeling purposes, not a fitted or guaranteed outcome; they will very likely need to be revised and refined for accuracy as the instrument's behavior is observed over more time and becomes better understood.

The par gate and the paid dividend. Two consequences follow. First, the ATM is available only when the price is at or above par,

$$\text{ATM open}_t \;=\; \left[P_t^{\text{STRC}} \geq \Pi\right] \wedge \left[q_t^{\text{STRC}} \leq \bar{q}^{\,\text{ceil}}\right] \wedge \left[a_t < a_{\max}\right]. \tag{17}$$

Second, below par the dividend actually paid lags the required rate—that lag is *why* the price is below par—so $q_t^{\text{paid}} = q_t^{\text{STRC}}\,(1 + \min(\delta_t, 0))$, floored, and the blended coupon and reserve band accrue on this paid rate. The important behavioral upshot is that a low-volatility product is *not* the same as a continuously-issuable one: management can hold STRC within \$90–102 yet, because the price sits *below* par most of the time—and is pinned further below in a hawkish-rate or Bitcoin-bear regime—the ATM is open only a fraction of the time and Bitcoin accumulation through it stalls, exactly the dynamic observed at the snapshot. The engine reports both the fraction of the horizon the ATM is open and the Bitcoin accumulated through it (Section 9).

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, then the cash reserve down to an operating floor;

2. at-the-market *common* equity: accretive when $m > 1$; when m is between a defensive floor m_{floor} (default 0.70, a 30% discount) and 1 the firm still issues to protect the senior stack, diluting the common, 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. only if m is below the defensive floor or the quarterly cap is exhausted, a forced *sale of Bitcoin*.

The cash (USD) reserve is targeted to a band running from a low boundary of twelve months to a high boundary of twenty-four months of forward preferred dividends—refilled toward the upper bound in calm periods and drawn down in stress—with a hard operating floor of roughly three months below which the waterfall must resort to selling Bitcoin. The base case starts the reserve at $\approx \$1.4$ billion. The preferreds carry no forced-liquidation trigger: their dividends may be paused, and cumulative series accrue arrears that compound until cleared.

The defensive floor m_{floor}—default 0.70, meaning the firm keeps selling common to honor its senior commitments down to a thirty-percent discount but no further—is 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. The waterfall above is defensive—it raises cash only to meet an obligation. The firm also issues *offensively*: when the premium is present ($m \geq 1.05$, so issuance is accretive) and Bitcoin is in strength (above a trailing reference), it sells common equity and buys Bitcoin, raising Bitcoin-per-share. Issuing ΔN shares at premium m changes backing per share by a factor $1 + \Delta N(m-1)/(N + \Delta N)$, accretive precisely when $m > 1$. Crucially the pace *scales with the premium*: the quarterly issuance is $\min\bigl(\eta\,(m-1),\,\bar{a}\bigr)$ as a fraction of Bitcoin value, with slope $\eta = 0.10$ and a market-depth cap $\bar{a} = 12\%$. At a 1.1× premium this is a trickle; at a 2–3× premium the firm drives a truck through the window. Because the proceeds buy coins, this program *de-levers* (it raises Bitcoin value without adding senior claims), and it is the mechanism through which a volatility-driven premium spike is converted into permanent Bitcoin-per-share.

5 The premium (mNAV) function

At any horizon the justified premium comprises four components—a net-backing ratio, an option/convexity premium, a convex upside flywheel, and a decaying entry dislocation,

$$m_t = \underbrace{\frac{H_t B_t - L_t + R_t}{H_t B_t}}_{\text{net-backing ratio } \rho_t} + \underbrace{\left(\pi_{\text{floor}} + \pi_{\text{cvx}} \frac{\sigma_t}{\sigma_{\text{ref}}}\right)}_{\text{option/convexity premium } \pi_t} + \underbrace{\pi_{\text{kick}} \left[\max\left(\frac{\sigma_t}{\bar{\sigma}_t} - 1, 0\right)\right]^2 \widehat{\rho}_t}_{\text{convex upside flywheel}} + \underbrace{g_0 \, e^{-\kappa_\pi t}}_{\text{entry dislocation}}, \tag{18}$$

where L_t is senior liabilities including arrears, R_t the reserve, and $\widehat{\rho}_t = \text{clip}(\rho_t, 0, 1)$. Each component is explained in turn.

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.

Option/convexity premium π_t. The second component is the premium the common commands *above* its net backing for the option-like character of a levered Bitcoin balance sheet. A levered, long position in a volatile, trending asset has a convex payoff: the upside is effectively unbounded while the downside is capped at the residual claim, and that asymmetry is worth a premium that rises with volatility. We represent it as a floor $\pi_{\text{floor}} = 0.15$ plus a volatility-scaled component governed by $\pi_{\text{cvx}} = 0.30$ acting through $\sigma_t/\sigma_{\text{ref}}$, the prevailing Bitcoin volatility relative to a reference level. The term is suppressed when the preferred dividend is paused: a structure visibly in distress does not command an optionality premium.

Convex upside flywheel. The third component corrects an asymmetry the smooth convexity premium misses. The largest historical premium expansions occur in *volatility spikes on otherwise healthy balance sheets*, when management can issue aggressively into strength and convert that volatility into permanent Bitcoin-per-share. The term activates only when realized volatility exceeds its own trailing anchor $\bar{\sigma}_t$, scales with the *square* of that excess—so it is dormant in calm markets and explosive in genuine spikes—and is multiplied by the clipped strength $\widehat{\rho}_t$ so that it fires only on well-collateralized paths. A volatility spike on a strong path therefore balloons the premium toward the 2–3× levels the security has historically reached, whereas a high-volatility *crash*—where

ρ_t is low—leaves the term near zero and lets the premium compress through ρ_t as it should. With kick coefficient $\pi_{\text{kick}} = 1.0$, the component is convex (squared) in the volatility excess and vanishes identically when the stochastic-volatility overlay is off.

Entry dislocation. The final component carries the gap $g_0 \equiv m_0 - (\rho_0 + \pi_0)$ between today's *observed* premium and its model-justified level, decaying exponentially at mean-reversion speed $\kappa_\pi = 0.25$/yr. It lets the simulation begin from the market's actual quote—a discount or premium that need not equal the model's fundamental value on day one—without imposing that dislocation in perpetuity. Because the gap has largely decayed by a multi-year exit, the terminal premium that drives the valuation is governed by fundamentals (ρ_T together with the optionality terms) rather than by the entry mispricing.

6 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 the terminal premium m_T delivered by the premium function of Section 5. 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

$$V_0 \;=\; B_0\,\frac{\operatorname{med}(v_T)}{(1+r_p)^{\tau}}, \tag{19}$$

where V_0 is the fair MSTR price today, B_0 the current Bitcoin spot price, $\operatorname{med}(\cdot)$ the cross-path median, r_p the annual counterparty/wrapper hurdle, and τ the holding horizon in years. The factor $\operatorname{med}(v_T)/(1+r_p)^{\tau}$ 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 (19), 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,

$$m_0^\star \;=\; \frac{V_0\,N_0}{H_0 B_0} \;=\; \frac{\operatorname{med}(v_T)}{b_0\,(1+r_p)^{\tau}}, \tag{20}$$

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.

7 Calibration

Table 1 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.

Input	Value (snapshot)
Bitcoin held	847,363 BTC
Bitcoin spot price	≈ $62,650
Bitcoin 200-week moving average	≈ $62,457
Bitcoin value of treasury	≈ $53.1 B
MSTR share price	≈ $98
Entry mNAV (this paper's definition)	≈ 0.71×
Diluted shares	386,052,000
Preferred liquidation preference (5 series)	≈ $15.5 B
STRC share of preferred (floating)	≈ 68% (≈ $10.5 B)
STRC price (snapshot)	$87.31 (below par)
STRC paid dividend / effective yield	11.5% / 13.2%
Convertible principal (6 tranches)	≈ $6.71 B
USD reserve	≈ $1.4 B
SOFR / policy rate	≈ 3.62%
Counterparty/wrapper hurdle r_p (base)	4.5%/yr

Table 1: Base-case calibration snapshot.

8 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 2 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* and the *hurdle* dominate the headline: the entire 0.29-to-1.50× spread reported below is the choice among the three price processes, and the premium falls monotonically as r_p rises. The *stochastic-volatility flywheel* is the largest single *structural* contributor, worth on the order of four-tenths of a turn. The *reserve band* and the *defensive mNAV 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.

9 Results

All figures below use the default profile: jump-diffusion rates, the stochastic-volatility flywheel, credit-rating migration, and soft financial repression all active, with each Bitcoin model anchored to its own fair value. The optional thesis levers (joint fiscal shocks, hard yield-curve control, Basel collateral eligibility) remain off.

The Bitcoin process dominates. At a ten-year horizon and the base hurdle, the rational mNAV is about 1.50× under the power law, 1.36× under diminishing returns, and 0.29× under

Assumption (symbol)	Default
Decision frame	
Holding horizon (τ)	10 years
Counterparty / wrapper hurdle (r_p)	4.5%/yr
Bitcoin model	power law
Monte Carlo paths × step (n_p)	20,000× quarterly
Reserve and financing waterfall	
USD reserve, start (R_0)	≈ $1.4 B
Reserve refill band (fwd preferred divs)	12 mo / 24 mo
Reserve hard operating floor	3 months
Defensive common-ATM floor (m_floor)	0.70
Below-par common issuance cap	3%/qtr of BTC value
Structural mNAV floor (m_min)	0.30
Accretive common-ATM trigger (m_acc)	$m \geq 1.05$
STRC par defense	
Target occupancy $97–101 / $90–102	≈ 40% / ≈ 90%
Calm equilibrium discount (δ_cal)	−2.4% (≈ $97.6)
Par-revert speed (κ_δ)	12/yr (∼3-wk half-life)
Calm price-deviation vol (σ_δ)	0.0295
Convex vol-stress sensitivity (γ_v)	0.35 (squared)
ATM opens only at/above par (Π)	$100
Base spread / rate ceiling (s_0, $\bar{c}$)	4.5% / 16%
Amplification (leverage policy)	
Target / ceiling / floor ($a_\mathrm{tgt}, a_\mathrm{max}, a_\mathrm{min}$)	0.45 / 0.60 / 0.30
STRC issuance pace cap	5%/qtr of BTC value
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 / ∼0.15 yr^{-1}
Soft-repression threshold (r^*)	4.5%
Bitcoin dynamics	
Power-law exponent (α)	5.77
Initial residual vol (σ_0, decaying)	0.50
Stochastic-vol overlay	on
Diminishing / GBM lag factors ($\phi_\mathrm{dim}, \phi_\mathrm{gbm}$)	1.75 / 1.43
Premium (mNAV) function	
Durable / convex premium ($\pi_\mathrm{floor}, \pi_\mathrm{cvx}$)	0.15 / 0.30
Flywheel kick / dislocation decay ($\pi_\mathrm{kick}, \kappa_\pi$)	1.0 / 0.25 yr^{-1}
Thesis toggles	
On by default	jump rates, vol flywheel, credit migration, soft repression
Off by default	joint fiscal shock, hard rate cap, Basel eligibility

Table 2: Behavioral and policy defaults—modeling choices, not market data.

GBM. The spread is the entire mNAV debate compressed into one assumption: whether Bitcoin reliably appreciates and recovers from a discount. Because each model now carries its own fair-value catch-up, even the conservative diminishing view sits well above net asset value—it, too, reads today's price as a discount to a lag-corrected trend. Under GBM the same leverage, dividend obligations, forced sales, and pauses *destroy* backing per share, and the rational price is a steep discount: holding the coin wins decisively.

The wrapper hurdle is underweighted in practice. Even on the bullish power law at ten years, the rational mNAV falls monotonically with the hurdle: 2.32× at 0%, 1.90× at 2%, 1.50× at 4.5%, 1.25× at 8%, and 0.87× at 12%. A sufficiently skeptical view of the *structure* alone compresses the premium toward—and past—parity even for a Bitcoin bull.

What the default profile contributes. Decomposing the power-law base case (1.74× at 4.5%, ten years), the stochastic-volatility flywheel is by far the largest single contributor: turning it off drops the premium to 1.12× (it adds ≈ 0.4 turns, by letting volatility spikes balloon the premium and fund accretive issuance). Credit-rating migration adds a few hundredths (1.49× with it off), and soft repression is essentially neutral in the base case (1.49× with it off)—unsurprising, since with a neutral SOFR near 3.5% the 4.5% repression threshold engages only in the restrictive-policy tail, so it leaves the base-case premium nearly unchanged and matters mainly in a hawkish or fiscal-dominance scenario. The par-gated ATM is itself a meaningful brake relative to a frictionless-funding assumption: because the program can issue only when STRC is at or above par—only about one-sixth of the time (roughly 17%) under the trend models, since the price now sits below par most of the time—the accretion flywheel runs at a realistic duty cycle rather than continuously.

STRC ATM availability and Bitcoin accumulation. The par-defense model lets us forecast *when* and *how much*. In the base case STRC sits within $97–101 about 40% of the time and within $90–102 about 90% of the time, and because the price now sits *below* par most of the time, the ATM is open—price at or above par—only about 17% of the horizon under the trend models and about 10% under GBM. Through that narrow open window the median path still accumulates Bitcoin equal to roughly 59% of the starting stack under the power law and 53% under diminishing returns over ten years, but only $\approx$ 17% under GBM. The binding constraint is the par gate rather than band collapse: a security held in a tight band that nonetheless trades below par most of the time chokes off issuance. A hawkish-Fed-plus-Bitcoin-bear regime (SOFR to $\approx$ 5% and held, elevated volatility) pushes the price lower still, so the ATM is open only $\approx$ 7% of the time and STRC-funded accumulation collapses to $\approx$ 9% of the stack. This reproduces the snapshot's lived experience, in which downside volatility froze the ATM while STRC remained over-collateralized. It is also the single largest reason the base-case premium is lower than a frictionless-funding model would imply: the cheap, accretive STRC channel is available far less of the time than a naive "low-volatility preferred" reading would suggest.

A belief cube. Figure 1 and Table 3 report the rational mNAV across the $3 \times 3 \times 3$ cube spanned by Bitcoin outlook (GBM, diminishing, power law), wrapper-risk hurdle (2%, 5%, 8%), and holding horizon (4, 8, 12 years), shown as three horizon cross-sections. The premium ranges from 0.17×—a deep, decisive discount in the skeptic's corner—to 2.10× at the most bullish. Three patterns stand out. First, the Bitcoin model is overwhelmingly the dominant axis: every GBM cell prices below net asset value regardless of hurdle or horizon, while every power-law cell prices above it; the conviction about Bitcoin's trend swamps the financing details. Second, horizon *amplifies* conviction in both directions—a longer hold pushes the power law higher (the flywheel compounds) and the GBM lower (volatility drag and the hurdle out-run weak growth), so the cube fans open with time rather than converging; the one exception is the short-horizon GBM corner, lifted by the deterministic catch-up from a mean-lens fair value that today sits above spot. Third, the hurdle matters most exactly where the decision is closest: in the diminishing-returns band it can move a cell from a clear premium to parity and just below, which is where an investor's own counterparty-risk judgment does real work.

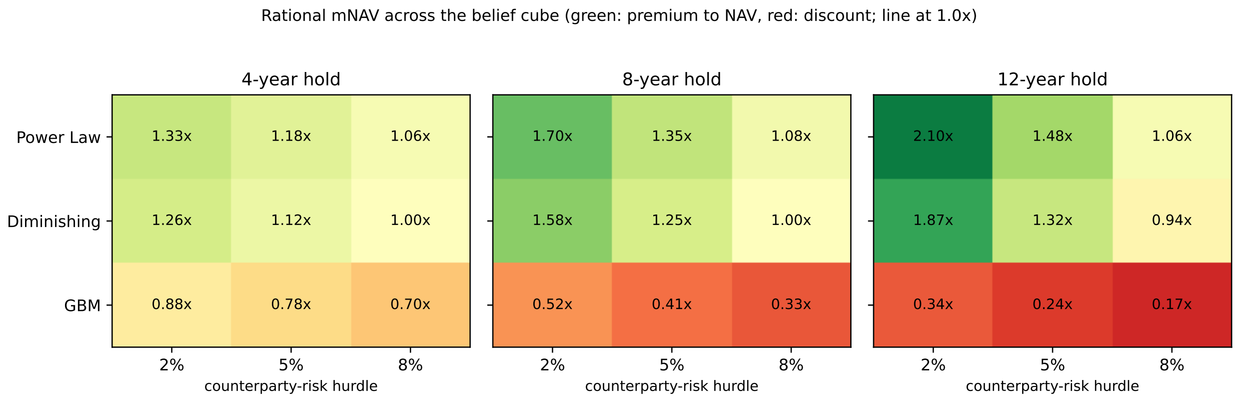

Figure 1: Rational mNAV across the $3 \times 3 \times 3$ belief cube, shown as three horizon cross-sections (default profile). Rows are the Bitcoin model (power law, diminishing, GBM); columns are the counterparty-risk hurdle (2%/5%/8%). Green denotes a premium to net asset value, red a discount; the dividing line is $1.0\times$.

	Counterparty-risk hurdle		
Bitcoin outlook	2%	5%	8%
4-year hold			
Power Law	1.33×	1.18×	1.06×
Diminishing returns	1.26×	1.12×	1.00×
GBM	0.88×	0.78×	0.70×
8-year hold			
Power Law	1.70×	1.35×	1.08×
Diminishing returns	1.58×	1.25×	1.00×
GBM	0.52×	0.41×	0.33×
12-year hold			
Power Law	2.10×	1.48×	1.06×
Diminishing returns	1.87×	1.32×	0.94×
GBM	0.34×	0.24×	0.17×

Table 3: Rational mNAV at all 27 belief sets of the cube (default profile).

For reference, the snapshot's observed $\approx 0.80\times$ premium sits below every trend-model cell and inside the GBM range—it is "fair" only under a thoroughly skeptical (GBM) view of Bitcoin. Under either trend model, at any hurdle or horizon in the grid, the equity reads as cheap relative to self-custody.

10 Limitations

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 repressed 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 soft repression engaged and a power-law Bitcoin trend describe a fiscal-dominance scenario; Vasicek rates reverting to a stable mean with repression off, paired with a GBM or diminishing-returns 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 soft repression, which is dormant at today's rates. The framework therefore takes no stance on the debate; it prices whichever side the user assumes.

The STRC price bands are assumed, not derived. The price-stability bands of Section 4.3 are a *modeling assumption*, not an emergent result of supply, demand, and balance-sheet capacity, and not a published management target. We set them by combining management's stated *intentions* for the instrument, its observed trading history (including the current below-par print), plausible forward scenarios, and a view of how hard the dividend can be flexed to defend the price. The calibration reflects that STRC has in practice traded *below* par: the distribution centers on a persistent modest discount ($\approx$ \$97.6), holds within \$90–102 about ninety percent of the time and within \$97–101 about forty percent of the time, and is skewed to the downside. It is therefore *not* an assumption that management keeps STRC pinned tightly at par—only that the par defense contains the price within a moderately wide, mostly-below-par band. Stress consequently shows up as restricted *ATM access* (with the price below par most of the time, issuance is shut more often than not) and, in the tail, as deeper within-band excursions. What lies outside the model is an outright *failure* of the stabilization mechanism—demand evaporates, the dividend cannot be raised fast enough, and STRC re-rates to a distressed yield far below the band—which would require relaxing the band assumption entirely. The base case should thus be read as "the par defense holds STRC within a wide, persistently-below-par band," the open questions being how often the firm can issue and how deep the excursions run, not whether the instrument can break down.

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 the exit value depends weakly on the entry price through the dislocation term, so the fair value of (19) is a one-pass solution rather than a fixed point—though the residual error is small because the dislocation decays well before a multi-year horizon.

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.

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 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. Markets involve risk of loss. Readers should conduct their own due diligence and consult a qualified, licensed professional before making any investment decision.

11 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, with the Bitcoin process the dominant lever, the wrapper hurdle a powerful and often neglected second, and the horizon an interaction term whose sign depends on conviction. 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.

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 2.

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	

Symbol	Meaning
m_t	mNAV, the market premium: equity market capitalization $P_t N_t$ over gross coin value $H_t B_t$.
$m_0^\star$	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_{floor}	defensive common-ATM floor (issue to defend credit down to this premium).
$m_{\min}$	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, H, B	snapshot share price, diluted 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	
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).
r^*	soft-repression threshold (the policy-rate ceiling leaned against from above).
λ_{rep}	soft-repression strength (per-year pull on the excess above r^*).
W_t^r, W_t^B	Brownian drivers for the rate and for Bitcoin.
Bitcoin dynamics	
F_t	model fair value of Bitcoin (power-law line, or lag-corrected 200-week moving average).
α, A	power-law exponent (5.77) and coefficient, $F = A\, d^\alpha$ in days d since genesis.
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.

Symbol	Meaning
ϕ_{dim}, ϕ_{gbm}	fair-value lag factors for the diminishing-returns and GBM models.
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.
ΔN	common shares issued in a step.
STRC par defense	
P_t^{STRC}	STRC price; $\Pi = \$100$ its par.
q_t^{paid}, q_t^{STRC}	STRC paid dividend and required (market-clearing) yield.
s_0, $\bar{c}$	base STRC spread over SOFR and the hard coupon-rate ceiling.
δ_t	STRC price deviation from par, $\delta_t = P_t^{\text{STRC}}/\Pi - 1$.
δ_{cal}	calm-equilibrium deviation (the persistent modest discount).
$\bar{\delta}_t$	time-varying reversion target for δ_t (shifts below δ_{cal} under stress).
κ_δ, σ_δ	par-revert speed and calm price-deviation volatility.
u_t	downward stress push; γ_v, γ_r its convex vol-stress and rate-stress sensitivities.
ε_t	standard normal shock.
Premium (mNAV) function	
ρ_t	net-backing ratio, $\rho_t = (H_t B_t - L_t + R_t)/(H_t B_t)$; $\widehat{\rho}_t = \text{clip}(\rho_t, 0, 1)$.
π_t	option/convexity premium; π_{floor} its durable floor and π_{cvx} its convexity coefficient.
π_{kick}	coefficient on the convex upside flywheel.
g_0	entry dislocation, $g_0 = m_0 - (\rho_0 + \pi_0)$; κ_π its decay speed.
Simulation	
n_{p}	number of Monte Carlo paths.
Δ	time step (one quarter).