# Mechanical Factor Premia in Automated Market Makers:
# Evidence from Bittensor Subnet Tokens

Philip Z. Maymin*

June 2026

**Abstract**

The constant-product automated market maker that prices Bittensor subnet tokens makes the cross-section of returns mechanically predictable. For a fixed emission staked into a pool, the percentage price impact is inversely proportional to pool size (Proposition 1), so smaller subnets earn higher returns by arithmetic rather than by risk. Using daily data on 128 subnet tokens from February 2025 through June 2026, I document a size premium of 1.0% per day (Newey-West $t = 3.1$) that survives value weighting, the exclusion of newly listed subnets, subsample splits, and, most tellingly, additional lags: it is a persistent property of the cross-section, not a one-day effect. A direct cross-sectional test confirms Proposition 1 through a negative return-on-size slope that strengthens once the contemporaneous emission-flow window is lagged out. The premium is mechanical, not risk-based, in the precise sense of Daniel and Titman (1997):

*Fairfield University, Dolan School of Business. Email: pmaymin@fairfield.edu.

the size characteristic predicts returns while its factor beta carries no out-of-sample premium. Emission yield, the most direct proxy for the staked flow, produces a large same-horizon effect that does not survive a one-day lag, identifying it as the contemporaneous price impact of the flow rather than a harvestable premium. The same AMM mechanics bound arbitrage: because slippage in a constant-product pool is deterministic (Proposition 2), the strategy is unprofitable above roughly \$10,000 in assets under management, which explains why the premium persists.

## Conflict-of-interest disclosure statement

**Philip Z. Maymin**

I am a co-owner of Djinn, a project that operates on the Bittensor network. I have no other potential conflicts to disclose.

Bittensor is a decentralized network that rewards the production of machine intelligence with tokens. The network is organized into 128 specialized subnetworks ("subnets"), each focused on a particular task such as language modeling or image generation. Since February 2025, each subnet has issued its own tradable token (an "alpha token") priced against the network's base currency ("TAO") by an on-chain automated market maker (AMM): a smart contract that holds reserves of both tokens and quotes prices algorithmically in the style of Uniswap V2. The result is a cross-section of 128 assets whose pricing function is fully observable on-chain.

This paper shows that the cross-section of subnet token returns is mechanically predictable, and that the mechanism is the AMM itself. The contribution is not another catalog of crypto factors but a setting in which the generating mechanism is transparent enough to be derived, signed, and tested. In equity markets the size-return relation is a robust regularity whose economic origin remains debated. In Bittensor the constant-product AMM makes the relation arithmetic: for a fixed token emission staked into a pool, the percentage price impact is inversely proportional to pool size, a result I state as Proposition 1. The central implication is a *size* premium that is mechanical rather than compensatory: smaller subnets appreciate more from the same emission, period after period, so the premium is persistent and predictable. A second implication operates at the instant of the flow rather than across days: a subnet's *emission yield* (the ratio of protocol rewards to market capitalization) maps into its same-period return, the AMM's contemporaneous price impact. I show that these two implications are empirically distinct, that only the size premium is persistent and harvestable, and that the same AMM that generates it imposes the slippage that protects it from arbitrage.

I construct daily returns for 128 subnets over 485 days (February 14, 2025, through

June 15, 2026) from blockchain-recorded pool data. I form long-short factor portfolios by daily tercile sorts on lagged characteristics: market capitalization (size), emission yield, past returns (momentum), liquidity, and stake. The construction follows Fama and French (1993) and its cryptocurrency application by Liu, Tsyvinski, and Wu (2022) and Borri and Shakhnov (2022).

Five results emerge. First, the size premium is large and persistent. A small-minus-big (SMB) portfolio earns 1.0% per day (Newey-West $t = 3.1$), and it is robust where a spurious premium would not be: it survives value weighting, the exclusion of newly listed subnets, and both subsample halves, and, most diagnostically, it survives additional lags of the sorting characteristic (0.8% per day at a two-day lag, $t = 2.8$, and similar thereafter). A persistent premium that does not decay when the signal is staled is a property of the cross-section, not of a one-day timing coincidence. Second, the prediction of Proposition 1 holds when tested *directly* in the cross-section: a daily Fama-MacBeth regression of returns on the standardized log size characteristic yields a negative slope ($-0.27$% per day, $t = -1.8$) that strengthens to $-0.50$% ($t = -4.1$) once the contemporaneous emission-flow window is lagged out, the sign and the persistence Proposition 1 requires, with no dependence on portfolio construction or on any single event. Third, the premium is mechanical, not risk-based, in the precise sense of Daniel and Titman (1997): returns track the size *characteristic*, but the corresponding factor *beta* carries no out-of-sample premium. When first-pass betas are estimated on the first half of the sample and priced in the second, the size beta earns an insignificant 0.10% per day ($t = 0.96$), even though sorting on the size characteristic in the same period earns 0.62% per day ($t = 4.1$). A premium that lives in the characteristic and not in the covariance is what a mechanical, non-compensatory premium looks like.

Fourth, the most direct proxy for the staked flow, emission yield, behaves exactly as a contemporaneous mechanical effect rather than a harvestable premium should. A high-minus-low emission-yield sort earns a striking 2.0% per day at the conventional one-day lag (Newey-West $t = 4.1$), but the effect is concentrated entirely at that horizon: it collapses to 0.3% ($t = 1.7$) when the characteristic is lagged a single additional day, and the direct emission-yield slope flips from $+0.70\%$ ($t = 3.7$) to $-0.15\%$ ($t = -2.2$). Such concentration at a single daily lag is the signature of the AMM's contemporaneous price impact (the staked emission moving the pool's price within the period of Proposition 1), or, read more conservatively, of microstructure timing in the daily snapshots; either way it is not a premium an investor can capture, and I do not treat it as one. The lag structure thus separates the two faces of Proposition 1: a within-period flow impact that is real but untradeable, and a persistent size premium that is harvestable up to costs.

Fifth, the same AMM mechanics that generate the premium also bound its arbitrage. Because slippage in a constant-product pool is deterministic (Proposition 2), I compute exact transaction costs: the size strategy is profitable only below roughly $10,000 in assets under management and is sharply loss-making at $100,000. The premium and its non-arbitrageability are two faces of one mechanism, which is why the premium persists.

One further prediction of Proposition 1, that a halving of emissions should attenuate the premium, receives only directional support, and I am careful not to overstate it. The December 2025 halving cut the block reward from 1 to 0.5 TAO, and the size premium does decline afterward, in the predicted direction. But the single event is confounded: it coincides with a change in the emission-allocation rule, and the decline cannot be separated from a general maturation trend (a linear time trend renders the post-halving dummy insignificant, and a placebo two months earlier produces a comparable decline). I therefore report the halving in Section V as corroboration, not identification. The robust test of Proposition 1 is the cross-sectional size slope, which uses every day and does not rest on a single date.

These results speak to two literatures. To cryptocurrency asset pricing, they add a protocol-native cross-section in which the factor structure is derived rather than fit. Liu and Tsyvinski (2021) show crypto returns load on crypto-specific factors; Liu, Tsyvinski, and Wu (2022) document a three-factor (market, size, momentum) description; Borri and Shakhnov (2022) document size and illiquidity premia in the cross-section of cryptocurrencies; Bianchi and Babiak (2021) find liquidity, size, and reversal as latent factors. I extend this work to AMM-priced tokens and, more importantly, identify the mechanism that produces the premia. To market microstructure, they connect the cross-section of returns to AMM design. Lehar and Parlour (2025) provide the first large-scale empirical analysis of the Uniswap constant-product AMM and document the near-absence of long-lived arbitrage; Park (2023) analyzes the inherent costs the constant-product rule imposes on traders; Capponi and Jia (2021) formalize price impact and liquidity provision in these pools. The Bittensor setting adds an endogenous flow, protocol emissions, that interacts with the pricing function to generate a priced, signed, and shock-testable cross-section. The mechanism has a direct equity analogue in Kyle (1985): flow shocks have larger percentage effects on illiquid, small-capitalization assets. What is usually dispersed and opaque, the flow and the market-making function, is here a single observable smart contract.

The remainder of the paper proceeds as follows. Section I describes the network and its pricing mechanism. Section II details the data. Section III presents the factor portfolios and the AMM-implied premia. Section IV contains the cross-sectional tests, including the direct test of Proposition 1 and the characteristics-versus-covariances evidence. Section V addresses robustness, the halving, transaction costs, and entry. Section VI concludes.

# I Institutional Background

## A The Bittensor Network

Bittensor is a decentralized peer-to-peer network in which participants supply machine intelligence services and are paid in TAO, the network's native cryptocurrency.[1] TAO trades on major exchanges at approximately $266 as of June 2026 (issued supply roughly 11 million, market capitalization near $2.9 billion). The network is organized into subnets, each specializing in a task such as language modeling, image generation, protein folding, or inference serving. As of June 2026, 128 subnets are active, with a combined alpha-token market capitalization of approximately $1.1 billion, about 38% of TAO's own market capitalization.

Each subnet contains up to 256 participants in two roles. *Miners* perform the subnet's core task; *validators* score miner output. The protocol distributes newly minted TAO ("emissions") in proportion to these scores. Of each subnet's emissions, 41% flows to miners, 41% to validators, and 18% to the subnet owner. Emissions function as seigniorage income paid to productive participants, analogous to a flow of interest in a traditional system.

## B Dynamic TAO and Subnet Alpha Tokens

Before February 2025, emissions were allocated across subnets by committee vote. On February 13, 2025, the network deployed Dynamic TAO (dTAO), replacing administrative allocation with market pricing.[2]

Under dTAO, each subnet maintains a constant-product AMM pool, a smart contract holding reserves of two tokens that sets prices algorithmically in the style of Uniswap V2

[1] The protocol was introduced by Rao (2020) and further developed in Steeves et al. (2022).

[2] See the dTAO whitepaper at `https://bittensor.com/dtao-whitepaper`.

(Adams, Zinsmeister, and Robinson, 2020). Each pool holds TAO ($\tau_i$) and the subnet's alpha tokens ($\alpha_i$). The alpha price in TAO is

$$p_i = \frac{\tau_i}{\alpha_i}, \tag{1}$$

where $\tau_i$ and $\alpha_i$ are the pool reserves. To invest, a participant deposits TAO and receives alpha ("staking"). Because the product $\tau_i \cdot \alpha_i$ is held fixed by the contract, depositing TAO reduces the pooled alpha and mechanically raises the price. Staking is economically equivalent to buying the alpha token.

Market capitalization in TAO equals $p_i \cdot A_i^{\text{total}}$, where $A_i^{\text{total}}$ is the total outstanding alpha supply (pooled and staked). Because the pool is constant-product, adding TAO reserves raises both the price and the market cap, directly linking size and liquidity.

## C Emission Allocation Under dTAO

The protocol distributes emissions across subnets by a market rule. In the initial dTAO phase (February through November 2025), a subnet's share of emissions was proportional to its alpha price:

$$\Delta\tau_i = \frac{p_i}{\sum_j p_j} \cdot \Delta\bar{\tau}, \tag{2}$$

where $\Delta\bar{\tau}$ is total emission per block. Since November 2025, allocation has tracked net TAO staking inflows through a 30-day exponential moving average (the "Taoflow" mechanism). Under either rule, subnets that attract capital receive more emissions, a reinforcing loop.

Total emission was 1 TAO per block (every 12 seconds) until the first halving on December 14, 2025, after which it fell to 0.5 TAO per block. At current prices this is roughly \$1 million per day across the 128 active subnets. The halving is the exogenous shock to the emission numerator of Proposition 1 that I exploit in Section V.

# II Data and Variable Construction

## A Data Sources

I collect daily subnet pool data from the Taostats API (https://api.taostats.io), which records blockchain state at daily frequency since the dTAO launch. For each subnet-day I observe the alpha price (in TAO), market capitalization, TAO and alpha reserves, alpha staked, and liquidity, together with daily metadata on emission rates, validator and miner counts, and registration parameters. TAO/USD prices come from the Taostats price API. The sample runs from February 14, 2025, through June 15, 2026: 485 daily observations and 484 daily returns.

## B Sample Construction

I exclude subnet 0 (the root network, whose price is identically 1 TAO) and subnets in "startup mode," a bootstrapping phase during which the AMM is inactive. This leaves 128 subnets with varying histories. The daily cross-section of eligible subnets expands from 64 early in the sample to 128 by June 2026 (Figure 6).

I compute daily simple returns in TAO,

$$r_{i,t}^{\mathrm{TAO}} = \frac{p_{i,t}}{p_{i,t-1}} - 1, \tag{3}$$

and in USD by compounding with the TAO/USD return,

$$r_{i,t}^{\mathrm{USD}} = (1 + r_{i,t}^{\mathrm{TAO}})(1 + r_t^{\mathrm{TAO/USD}}) - 1. \tag{4}$$

I winsorize daily returns at ±100%; results are qualitatively unchanged without winsorization. I require at least 7 days of return history for a subnet to enter on a given date, which avoids contamination from initial AMM price discovery.

## C Characteristics

I construct the following lagged characteristics, using information available at the close of day $t-1$ to form portfolios on day $t$:

- **Market capitalization**: $\mathrm{MCAP}_{i,t-1} = p_{i,t-1} \cdot A_{i,t-1}^{\mathrm{total}}$, in TAO.
- **Emission yield**: $\mathrm{EY}_{i,t-1} = E_{i,t-1}/\mathrm{MCAP}_{i,t-1}$, where $E_{i,t-1}$ is the daily emission rate. Numerator and denominator share units (rao, where $10^9$ rao = 1 TAO), so the ratio is unitless, analogous to a dividend yield.
- **Momentum (7-day)**: $\mathrm{MOM7}_{i,t-1} = p_{i,t-1}/p_{i,t-8} - 1$.
- **Momentum (30-day)**: $\mathrm{MOM30}_{i,t-1} = p_{i,t-1}/p_{i,t-31} - 1$.
- **Reversal**: $\mathrm{REV}_{i,t-1} = p_{i,t-1}/p_{i,t-2} - 1$ (1-day past return).
- **Liquidity**: $\mathrm{LIQ}_{i,t-1} = \tau_{i,t-1}$, the TAO reserves in the pool.
- **Stake**: $\mathrm{STAKE}_{i,t-1} = \alpha_{i,t-1}^{\mathrm{staked}}$, alpha staked outside the pool.

# III Factor Portfolios

## A Portfolio Construction

Each day I sort eligible subnets into terciles on each lagged characteristic and compute equal-weighted daily returns for the bottom, middle, and top terciles. Long-short factors are the difference between extreme terciles:

- **MKT**: equal-weighted return of all subnets (TAO terms).
- **SMB**: bottom (small) minus top (big) market cap.
- $\mathbf{HML_{EMIS}}$: top (high) minus bottom (low) emission yield.
- $\mathbf{WML_{7}}$: top (winners) minus bottom (losers) 7-day momentum.
- $\mathbf{WML_{30}}$: same for 30-day momentum.
- **REV**: bottom (past losers) minus top (past winners) 1-day return.
- **LIQ**: bottom (illiquid) minus top (liquid) TAO reserves.
- **STAKE**: top minus bottom stake.

## B Summary Statistics

Table I reports summary statistics. The size factor (SMB) earns 1.04% per day (Newey-West $t = 3.10$, OLS $t = 3.42$, annualized Sharpe 2.97), and Section V shows it survives value weighting, the exclusion of new entrants, both subsamples, and additional lags. Table II reports the underlying size-tercile returns: the smallest subnets carry the premium while the largest lose value outright, and Figure 2 plots the cumulative returns.

The emission-yield sort earns a larger raw return, 2.04% per day (NW $t = 4.15$), and Table III shows it is concentrated in the high-yield tercile. I report this figure for completeness but do not treat it as a premium, because it is a one-day effect: Section A and Table VII show that it collapses to 0.3% per day as soon as the sorting characteristic is lagged one additional day, whereas the size premium does not. As Section C explains, emission yield is the most direct proxy for the staked emission relative to pool depth, the quantity Proposition 1 prices contemporaneously; that the emission-yield effect lives only at the contemporaneous horizon, while the size premium persists across lags, is exactly what the proposition implies and is the central diagnostic of the paper. Both sorts have heavy right tails (skewness 8.3–8.4, kurtosis above 100), so all inference uses Newey-West standard errors with 5 lags, the referee-standard correction for the serial dependence and fat tails evident in Table I.

Momentum is significant but secondary: the 7-day factor earns 0.55% per day (NW $t = 2.89$) and the 30-day factor 0.40% (NW $t = 2.39$). Short-term reversal is negative ($-0.49$% per day, NW $t = -2.00$), confirming momentum rather than contrarian profits. The liquidity factor earns 1.03% per day (NW $t = 3.17$) but is 0.98 correlated with SMB (Table V): the constant-product pool mechanically ties depth to market cap, so SMB and LIQ measure the same object. In spanning tests (Section C) the liquidity alpha vanishes once size is controlled.

## C AMM-Implied Premia

The size premium, the contemporaneous emission-yield effect, and the redundancy of liquidity all have a structural explanation in the constant-product mechanics.

**Proposition 1** (Emission Amplification). *Consider a subnet with TAO reserves $\tau$ and alpha reserves $\alpha$ in a constant-product pool ($k = \tau \cdot \alpha$). If the subnet receives an emission*

*$\Delta\tau$ that is staked into the pool, the percentage price change is approximately*

$$\frac{\Delta p}{p} \approx \frac{2\,\Delta\tau}{\tau}. \tag{5}$$

*For a fixed emission $\Delta\tau$, the return is inversely proportional to the TAO reserve, and hence to market capitalization. Small subnets earn mechanically higher returns from the same emission.*

*Proof.* Before the emission, $p_0 = \tau/\alpha$. After staking $\Delta\tau$, the constant product $k = \tau\alpha$ gives $\alpha' = k/(\tau + \Delta\tau)$, so $p_1 = (\tau + \Delta\tau)/\alpha' = (\tau + \Delta\tau)^2/(\tau\alpha)$. Then $p_1/p_0 - 1 = (\tau + \Delta\tau)^2/\tau^2 - 1 = 2\Delta\tau/\tau + (\Delta\tau/\tau)^2 \approx 2\Delta\tau/\tau$ for small $\Delta\tau/\tau$. Since $\tau$ is proportional to market cap, the return from a given emission is inversely proportional to size. $\square$

Two signed predictions follow, and they operate at different horizons, which organizes the rest of the paper. The first is contemporaneous. At the instant a subnet's emission $\Delta\tau$ is staked, the price moves by $2\Delta\tau/\tau$, so a subnet's same-period return rises with the emission staked relative to pool depth. The closest observable proxy is the emission yield $E/\text{MCAP}$, whose numerator is the emission and whose denominator scales with $\tau$; Proposition 1 predicts a *positive* contemporaneous relation between emission yield and return. This is a statement about the within-period price impact, not about predictability: it says nothing an investor can trade on, because the emission and the price move it causes occur together. The second prediction is persistent. Holding the emission flow roughly fixed across the cross-section, a smaller pool appreciates more from each period's emission than a larger one, period after period, so smaller subnets earn higher returns going forward. This is the size premium, and unlike the contemporaneous emission-yield effect it is predictable from lagged information. The cross-sectional emission dispersion is large (the ratio of the largest to the smallest daily emission is 59, coefficient of variation 1.63), so emission yield is not a relabeling of inverse size: its average daily cross-sectional correlation with 1/MCAP is 0.67, far from one. Section A tests both predictions as characteristic slopes and shows that they separate exactly along the contemporaneous-versus-persistent line the proposition draws.

The proposition shows the premium is not merely an empirical regularity but an arithmetic consequence of the pricing rule. Emission-funded staking moves a small pool's price more, in percentage terms, than a large pool's, generating a size premium that persists as long as emissions flow.

## D Discussion

Proposition 1 captures the mechanical core, but two further channels reinforce it. First, validators and miners who receive emissions often restake them into the same subnet, a feedback loop that amplifies appreciation in small, high-yield pools. Second, part of the size premium may compensate for illiquidity: small pools have wide effective spreads, and an investor risks costly exit. Section E quantifies that cost exactly and shows it is the binding constraint on arbitrage.

# IV Cross-Sectional Asset Pricing Tests

## A Direct Test of Proposition 1

The cleanest test of Proposition 1 does not go through factor portfolios at all. The proposition is a statement about characteristics, and it makes two predictions that differ by horizon: a contemporaneous positive relation between emission yield and return, and a persistent negative relation between size and return. I test both directly by running, each day, a cross-sectional regression of subnet returns on the two standardized characteristics and averaging the slopes over time in the manner of Fama and MacBeth (1973):

$$r_{i,t} = \lambda_{0,t} + \lambda_{\mathrm{EY},t}\, z(\mathrm{EY}_{i,t-\ell}) + \lambda_{\mathrm{SIZE},t}\, z(\log \mathrm{MCAP}_{i,t-\ell}) + \varepsilon_{i,t}, \tag{6}$$

where $z(\cdot)$ denotes a daily cross-sectional standardization and $\ell$ is the lag at which the characteristics are observed. Varying $\ell$ is the key test, because the two predictions of Proposition 1 imply opposite behavior as the signal is staled.

Table VI reports the result. At the conventional one-day lag ($\ell = 1$) both slopes carry the signs Proposition 1 requires: the emission-yield slope is +0.70% per day (Newey-West $t = 3.68$) and the size slope is $-0.27\%$ per day ($t = -1.82$). The two characteristics behave very differently, however, when the signal is lagged a single additional day. The emission-yield slope reverses, to $-0.15\%$ per day ($t = -2.17$) at $\ell = 2$: the positive relation existed only at the contemporaneous horizon, exactly as a within-period price impact must, since the staked emission and the price move it causes occur together and cannot predict the next day's return. The size slope does the opposite. It strengthens, to $-0.50\%$ per day ($t = -4.05$) at $\ell = 2$ and $-0.42\%$ ($t = -3.48$) at $\ell = 3$. Lagging out the contemporaneous emission-flow channel, which at $\ell = 1$ partly offsets the size slope, isolates the persistent size premium and sharpens it. The robust, harvestable prediction of Proposition 1 is the size slope, and it is significant and stable across every lag beyond the contemporaneous one. This test depends on neither portfolio construction nor any single event, and it is the paper's primary structural evidence.

Table VII makes the same point at the portfolio level and is the central diagnostic of the paper. The size factor earns 1.04% per day at $\ell = 1$ and a statistically indistinguishable

0.77% to 0.84% at lags two through five (NW $t$ between 2.75 and 3.29 throughout): a persistent premium that does not decay when the signal is staled. The emission-yield factor earns 2.04% at $\ell = 1$ but only 0.24% to 0.29% at longer lags, where it is no longer reliably significant. A premium an investor could capture would survive a one-day delay in observing the characteristic; the size premium does, and the emission-yield effect does not. This is the empirical content of the distinction between a persistent cross-sectional premium and a contemporaneous mechanical price impact, and it is why the paper is organized around size.

## B Fama-MacBeth Regressions and the Characteristics-Covariances Distinction

I next ask whether the premia are compensation for bearing factor risk. They are not, and the way they fail this test is itself the evidence that they are mechanical.

Table VIII, Panel A, runs the standard two-pass regression with first-pass betas estimated over the full sample. The size beta earns a significant 0.62% per day ($t = 2.89$); other betas do not. But full-sample betas use future information. The honest test estimates betas on the first half of the sample and prices them in the second (Panel B). Out of sample, the size beta-premium collapses to 0.10% per day ($t = 0.96$), and no beta is priced; the intercept, by contrast, remains significant ($t = 2.13$).

The collapse is not a weakness of the result; it is the result. In the same second half, sorting on the size *characteristic* earns 0.62% per day ($t = 4.09$, Table X). Returns track the characteristic, not the covariance with the factor. This is precisely the Daniel and Titman (1997) characteristics-versus-covariances pattern, and it is what a mechanical premium must look like: Proposition 1 ties return to a subnet's own size, not to how its return co-moves with a portfolio. A risk story predicts the opposite, that loadings should price out of sample.

They do not.

## C Spanning Tests

Table IX reports spanning alphas: the intercept from regressing each factor on all others. SMB retains a small but significant alpha (0.09% per day, $t = 2.08$); the liquidity alpha is insignificant ($t = -0.70$) once size is present, confirming that size, not liquidity per se, is the fundamental characteristic. The emission-yield factor carries a large spanning alpha (1.02% per day, $t = 9.28$), but this is the same contemporaneous one-day effect documented in Section A: no other factor spans it precisely because it is a within-period price impact rather than a persistent return, and it disappears once the sorting characteristic is lagged. The reversal factor has a significantly negative alpha, the flip side of momentum.

## D Beyond a Pricing Model

I do not claim that these factors constitute a complete pricing model. Applied to the factor-sort portfolios themselves, a Gibbons, Ross, and Shanken (1989) test cannot reject zero intercepts, but that test is near-circular, since the test assets are the sorts that define the factors. The honest specification test uses individual subnets. There, the GRS test *rejects* every nested model I consider: market only ($F = 1.56$, $p = 0.004$), adding size ($p = 0.030$), adding momentum ($p = 0.048$), and adding emission yield ($p = 0.044$); the average absolute pricing error does not shrink monotonically as factors are added. The factors describe the dominant common variation in subnet returns, and each carries an independent spanning alpha, but they do not price every subnet. That is consistent with the central thesis: the premia are mechanical functions of characteristics, not loadings on a small set of priced risks.

# V Robustness

## A USD-Denominated Returns

The baseline uses TAO-denominated returns to isolate subnet-specific variation from the TAO/USD rate. Recomputing the market factor in USD gives 0.39% per day ($t = 1.17$), with higher volatility and a lower Sharpe (1.02 versus 1.40), because the added exchange-rate volatility is noise. The long-short premia, which difference out the common TAO/USD movement, are essentially unchanged in USD.

## B Value Weighting

Equal-weighted long-short portfolios could in principle be driven by the smallest, least investable subnets. They are not. Table XII, Panel A, recomputes the factors with value (market-cap) weighting. The size premium survives and in fact strengthens, from 1.04% to 1.12% per day (NW $t = 3.31$); the contemporaneous emission-yield sort is likewise not an equal-weighting artifact (2.04% to 2.22%). The size premium is not driven by the smallest subnets.

## C Subsample Analysis

I split the sample at its midpoint (first half February 15 to October 16, 2025; second half October 17, 2025, to June 15, 2026). Table X reports factor means and $t$-statistics for each half. The size premium is significant in both halves: SMB earns 1.45% ($t = 2.49$) then 0.62% ($t = 4.09$). The raw mean declines but significance rises as the cross-section widens and volatility falls. The one-day emission-yield effect is present in both halves as well (2.93% then 1.15%). Momentum attenuates more sharply ($\text{WML}_{30}$ from $t = 2.47$ to

$t = 1.12$), consistent with maturation of return persistence. The decline in the size premium across halves is consistent with Proposition 1 under the December halving, which I examine next.

## D The Halving: Directional Corroboration

Proposition 1 makes the size premium proportional to the emission rate: the premium is the percentage price impact of the staked emission, so an exogenous cut to the emission should attenuate it. The first halving on December 14, 2025, cut the block reward from 1 to 0.5 TAO, and the size premium does decline afterward. Table XI, Panel A, shows the full-sample SMB premium falling from 1.39% to 0.45% per day across the event.

I am deliberately cautious about what this establishes. The original version of this paper leaned on the halving, and the single event cannot bear that weight. Three features prevent a clean causal reading. First, SMB is highly volatile, so the decline, while in the predicted direction, is imprecisely estimated (Welch $p = 0.06$, Newey-West $p = 0.09$), and the symmetric-window estimates in Panel B put the post/pre ratio anywhere between 0.04 and 0.40 across 30- to 90-day windows. Second, the halving coincides with the November 2025 switch from price-based to flow-based emission allocation, so two regime changes overlap. Third, and most important, the decline cannot be separated from a general maturation trend: in a regression of daily SMB on a post-halving dummy, a market control, and a linear time trend, the post dummy becomes insignificant ($t = -0.31$), and a placebo halving two months earlier produces a decline of similar magnitude (the actual event ranks only second of six placebo dates by coefficient size).

I therefore read the halving as directional corroboration, not identification: the size premium moved in the predicted direction, but the single event, confounded with a contemporaneous rule change and a maturation trend, cannot identify the coefficient. This is a deliberate retreat from the stronger causal claim the first draft made. The robust test of Proposition 1 is the cross-sectional size slope of Section A, which uses every day and does not hinge on one date. Figure 9 plots the cumulative SMB return around the event.

## E Transaction Costs and Portfolio Capacity

Are the premia implementable? Because subnet tokens trade on constant-product AMMs, slippage for any trade is deterministic.

**Proposition 2** (AMM Slippage). *In a constant-product pool with TAO reserves $\tau$, the one-way slippage from buying $\Delta\tau$ worth of alpha is exactly*

$$\text{Slippage} = \frac{\Delta\tau}{\tau}, \tag{7}$$

*deterministic and linear in trade size relative to pool depth.*

Using on-chain reserves, Table XIII computes slippage for equal-weighted portfolios at several AUM levels. For the small-tercile subnets that drive SMB, one-way slippage is 0.59% at \$10K, 5.90% at \$100K, and 58.98% at \$1M, rising linearly with size relative to the shallow pools. Figure 7 shows the consequence: at \$10K with daily rebalancing the net-of-cost SMB return is 0.41% per day (annualized net Sharpe 1.56), still profitable; at \$100K slippage of 5.98% round-trip turns the 1.04% gross premium into a 4.97% daily loss; at \$1M and above the strategy is ruinous.

This is why the premium persists. Aggregate daily capacity across the small-tercile pools is on the order of tens of thousands of dollars, far below the threshold that would attract systematic arbitrage. The premium and its non-arbitrageability are the same fact stated twice: the AMM that amplifies emissions into the premium (Proposition 1) is the AMM that imposes the slippage which protects it (Proposition 2). For comparison, Liu, Tsyvinski, and Wu (2022) report crypto factors that survive conservative cost estimates; the difference here is that the constant-product rule imposes deterministic, size-dependent costs orders of magnitude larger than limit-order-book spreads for equivalent trades, the same costs Park (2023) identifies as intrinsic to the mechanism.

## F Volatility and Downside Risk

I sort subnets into terciles on 30-day realized volatility, downside semi-deviation, upside semi-deviation, idiosyncratic volatility (residual from a market model), and market beta. Table XIV shows that high-risk subnets earn higher returns on every measure: high-minus-low total volatility earns 0.58% per day ($t$ = 4.20), high-minus-low downside 0.48% ($t$ = 3.74), and high-minus-low beta 0.52% ($t$ = 4.09). This positive risk-return relation is the opposite of the low-volatility and betting-against-beta anomalies in equities (Ang, Chen, and Xing, 2006; Frazzini and Pedersen, 2014), and it is not evidence of efficient risk pricing. Volatile subnets are small subnets with thin pools, where the constant-product formula amplifies price moves; the low-minus-high volatility factor is strongly correlated with SMB, so size subsumes most of the effect.

Table XV decomposes factor returns into downside and upside semi-deviations. The size premium has favorable asymmetry: SMB's downside deviation (4.01% per day) is well below its upside (8.18%), a ratio of 0.49 and a Sortino ratio of 4.93. The asymmetry means the premium is not simply compensation for crash risk. The emission-yield row is reported for completeness and shows even stronger asymmetry (ratio 0.26), but as Section A establishes its return is a contemporaneous one-day effect, so its Sortino ratio is not investable.

## G Entry and Survivorship

Two features of a young, growing network could in principle manufacture the premia: the recycling of subnet slots and the steady entry of new, small subnets. Neither does.

Subnets can be deregistered when their emissions fall to zero or a new registrant outbids the slot. During the sample, 23 of the 128 integer slots were recycled. Because the network reuses identifiers, a naive panel would splice unrelated projects. I prevent this in two ways. When a subnet re-enters startup mode, the startup filter opens a gap of roughly seven days, and percentage-change returns are missing across the boundary. For momentum signals, I explicitly nullify any value whose lookback window spans an internal gap, masking 7 values for 7-day and 448 for 30-day momentum; all premia are robust to this correction.

The more serious concern, raised naturally by the cross-section growing from 64 to 128, is an entry or "IPO" effect: new subnets are small, so the bottom size tercile is partly fresh listings whose early returns could inflate SMB. I test this directly rather than note it as a caveat. Table XII, Panel B, recomputes the premia under age filters and a fixed-cohort restriction. Excluding subnets younger than 30 days leaves SMB at 0.84% per day (NW $t = 3.36$); excluding those younger than 60 days leaves SMB at 0.98% ($t = 4.71$). Restricting to the fixed cohort of 64 subnets present at the dTAO launch, which by construction admits no entrants, leaves SMB at 0.96% (NW $t = 3.38$). The contemporaneous emission-yield effect survives the same filters (1.31%, 1.13%, and 2.11% respectively). Newly listed subnets are a minority even of the small tercile (15.3% are younger than 30 days, 27.8% younger than 60). The size premium is a property of the pricing mechanism, not of the entry process.

# VI Conclusion

The cross-section of Bittensor subnet tokens is mechanically predictable, and the mechanism is the automated market maker that prices them. For a fixed emission staked into a constant-product pool, the percentage price impact is inversely proportional to pool size (Proposition 1). The harvestable consequence is a size premium of 1.0% per day (Newey-West $t = 3.1$): smaller subnets appreciate more from each period's emission, period after period. Tested directly as a characteristic slope the premium carries the predicted negative sign, and value weighting, the exclusion of new entrants, subsample splits, and additional lags of the sorting characteristic leave it intact. The most direct proxy for the staked flow, emission yield, produces a large effect of its own, but only at the contemporaneous horizon: it does not survive a one-day lag, marking it as the AMM's within-period price impact rather than a tradeable premium. The lag structure separates the two, and only size persists.

The premium is mechanical rather than risk-based, in a sense the data make precise. Returns track the size characteristic, but its factor beta earns nothing out of sample: a premium that lives in the characteristic and not in the covariance is what Daniel and Titman (1997) taught us to look for, and it is what a non-compensatory, arithmetic premium must produce. The same logic explains why the factors do not constitute a pricing model for individual subnets, and why I do not claim one.

The December 2025 halving, an exogenous cut to the emission rate, moves the size premium in the predicted direction, but I report it as directional corroboration only, a deliberate retreat from the stronger causal claim of the first draft: the change is imprecisely estimated, overlaps a change in the allocation rule, and cannot be disentangled from a maturation trend. The result that does not depend on a single date is the cross-sectional size slope, and it is significant at every lag beyond the contemporaneous one.

Finally, the premium and its persistence are one mechanism seen twice. Because slippage in a constant-product pool is exact (Proposition 2), the size strategy is profitable only below roughly \$10,000 in AUM and loss-making by \$100,000. The AMM that amplifies emissions into the premium is the AMM that imposes the cost which keeps the premium unarbitraged. This connects the cross-section of returns to AMM design in the spirit of Lehar and Parlour (2025) and Park (2023), and it gives the Kyle (1985) intuition, that flows move illiquid assets more, a setting where both the flow and the market-making function are a single observable contract. Whether the structure survives as pools deepen and the market matures is the natural question for future work.

Table I: Factor Portfolio Summary Statistics

| Factor | Mean (%/day) | Std (%/day) | Sharpe (ann.) | $t$-stat (OLS) | NW $t$ (5 lags) | Skew | Kurt | $N$ |
|---|---|---|---|---|---|---|---|---|
| MKT | 0.33 | 4.56 | 1.40 | 1.61 | 1.51 | 5.25 | 71.06 | 484 |
| SMB | 1.04 | 6.66 | 2.97 | 3.42 | 3.10 | 8.40 | 127.67 | 484 |
| $\text{HML}_{\text{EMIS}}$ | 2.04 | 6.96 | 5.61 | 6.46 | 4.15 | 8.30 | 108.42 | 484 |
| $\text{WML}_7$ | 0.55 | 3.57 | 2.96 | 3.39 | 2.89 | 0.03 | 11.05 | 477 |
| $\text{WML}_{30}$ | 0.40 | 3.15 | 2.42 | 2.70 | 2.39 | 0.37 | 8.16 | 454 |
| REV | −0.49 | 3.75 | −2.49 | −2.86 | −2.00 | 2.65 | 26.04 | 483 |
| LIQ | 1.03 | 6.52 | 3.00 | 3.46 | 3.17 | 7.97 | 118.29 | 484 |
| STAKE | −0.43 | 4.94 | −1.65 | −1.91 | −1.31 | −4.27 | 33.45 | 484 |

Daily summary statistics for long-short factor portfolios from tercile sorts on lagged subnet characteristics. MKT is the equal-weighted return of all subnets. SMB is small minus big market capitalization. $\text{HML}_{\text{EMIS}}$ is high minus low emission yield. $\text{WML}_7$ and $\text{WML}_{30}$ are winners minus losers on 7- and 30-day past returns. REV is past losers minus past winners (1-day). LIQ is illiquid minus liquid (TAO reserves). STAKE is high minus low alpha staked. Sharpe ratios are annualized ($\times\sqrt{365}$). OLS $t$ assumes i.i.d. returns; NW $t$ uses Newey-West standard errors with 5 lags. Average tercile size: 38 subnets (minimum 21). Sample: February 2025 through June 2026.

Table II: Size-Sorted Tercile Portfolio Returns

| Portfolio | Mean (%/day) | Ann. Return (%) | Ann. Std (%) | Sharpe |
|---|---|---|---|---|
| Small (bottom tercile) | 0.91 | 332.8 | 128.0 | 2.60 |
| Medium | 0.22 | 78.7 | 86.8 | 0.91 |
| Large (top tercile) | −0.12 | −45.3 | 92.4 | −0.49 |
| SMB (Small − Large) | 1.04 | 378.1 | 127.3 | 2.97 |

Subnets are sorted daily into terciles on lagged market capitalization (TAO). Portfolios are equal-weighted, returns in TAO. Annualization uses 365 days. Sample: February 2025 through June 2026.

Table III: Emission-Yield-Sorted Tercile Portfolio Returns

| Portfolio | Mean (%/day) | Ann. Return (%) | Ann. Std (%) | Sharpe |
|---|---|---|---|---|
| Low (bottom tercile) | −0.61 | −223.8 | 87.1 | −2.57 |
| Medium | 0.24 | 86.6 | 89.6 | 0.97 |
| High (top tercile) | 1.43 | 522.1 | 132.1 | 3.95 |
| $\text{HML}_{\text{EMIS}}$ (High − Low) | 2.04 | 746.0 | 133.0 | 5.61 |

Subnets are sorted daily into terciles on lagged emission yield (daily emission rate divided by market capitalization). Portfolios are equal-weighted, returns in TAO. The high-yield tercile carries the effect and the low-yield tercile loses value. These are the one-day-lag figures: Table VII shows the high-minus-low difference collapses when the characteristic is lagged a further day, identifying it as the AMM's contemporaneous price impact (Proposition 1) rather than a harvestable premium. Sample: February 2025 through June 2026.

Table IV: Momentum-Sorted Tercile Portfolio Returns

| Portfolio | Mean (%/day) | Ann. Return (%) | Ann. Std (%) | Sharpe |
|---|---|---|---|---|
| *Panel A: 7-Day Momentum* | | | | |
| Losers (bottom tercile) | −0.08 | −30.7 | 47.0 | −0.65 |
| Middle | 0.04 | 13.9 | 41.9 | 0.33 |
| Winners (top tercile) | 0.47 | 171.5 | 68.9 | 2.49 |
| $WML_7$ | 0.55 | 202.2 | 68.3 | 2.96 |
| *Panel B: 30-Day Momentum* | | | | |
| Losers (bottom tercile) | 0.16 | 57.5 | 39.6 | 1.45 |
| Middle | 0.20 | 71.2 | 33.6 | 2.12 |
| Winners (top tercile) | 0.56 | 203.3 | 57.0 | 3.57 |
| $WML_{30}$ | 0.40 | 145.7 | 60.1 | 2.42 |

Subnets are sorted daily into terciles on past 7-day (Panel A) or 30-day (Panel B) returns. Portfolios are equal-weighted, returns in TAO.

Table V: Factor Correlation Matrix

| | MKT | SMB | EMIS | MOM7 | MOM30 | REV | LIQ | STK |
|---|---|---|---|---|---|---|---|---|
| MKT | 1.00 | 0.35 | 0.38 | 0.34 | 0.29 | −0.06 | 0.37 | −0.13 |
| SMB | 0.35 | 1.00 | 0.60 | 0.00 | −0.22 | −0.12 | 0.98 | −0.49 |
| EMIS | 0.38 | 0.60 | 1.00 | −0.16 | −0.00 | 0.35 | 0.57 | −0.70 |
| MOM7 | 0.34 | 0.00 | −0.16 | 1.00 | 0.39 | −0.43 | 0.10 | 0.24 |
| MOM30 | 0.29 | −0.22 | −0.00 | 0.39 | 1.00 | −0.21 | −0.13 | 0.41 |
| REV | −0.06 | −0.12 | 0.35 | −0.43 | −0.21 | 1.00 | −0.18 | −0.31 |
| LIQ | 0.37 | 0.98 | 0.57 | 0.10 | −0.13 | −0.18 | 1.00 | −0.42 |
| STK | −0.13 | −0.49 | −0.70 | 0.24 | 0.41 | −0.31 | −0.42 | 1.00 |

Pairwise correlations of daily factor returns. EMIS = $HML_{EMIS}$. MOM7 and MOM30 = WML at 7- and 30-day horizons. STK = STAKE. The 0.98 correlation between SMB and LIQ reflects the AMM structure linking market cap to pool depth.

Table VI: Direct Cross-Sectional Test of Proposition 1: Slopes by Characteristic Lag

| | Lag $\ell = 1$ | | Lag $\ell = 2$ | |
|---|---|---|---|---|
| Characteristic | Slope (%/day) | NW $t$ | Slope (%/day) | NW $t$ |
| $z$(emission yield) | +0.70 | +3.68 | −0.15 | −2.17 |
| $z$(log market cap) | −0.27 | −1.82 | −0.50 | −4.05 |

Time-series averages of daily cross-sectional slopes from regressing subnet returns on standardized lagged characteristics, $r_{i,t} = \lambda_{0,t} + \lambda_{\mathrm{EY},t}\, z(\mathrm{EY}_{i,t-\ell}) + \lambda_{\mathrm{SIZE},t}\, z(\log \mathrm{MCAP}_{i,t-\ell}) + \varepsilon_{i,t}$, at characteristic lags $\ell = 1$ and $\ell = 2$ days. Emission yield enters as its cross-sectional percentile rank to limit the influence of near-zero-capitalization outliers; standardization is daily and cross-sectional. Newey-West $t$-statistics (5 lags). Proposition 1 predicts a positive *contemporaneous* emission-yield slope and a *persistent* negative size slope. As $\ell$ increases past the contemporaneous horizon the emission-yield slope reverses sign while the size slope strengthens, separating the within-period flow impact from the persistent premium. Sample: February 2025 through June 2026.

Table VII: Persistence Diagnostic: Factor Premia by Characteristic Lag

| | Lag 1 | | Lag 2 | | Lag 3 | | Lag 5 | |
|---|---|---|---|---|---|---|---|---|
| Factor | Mean (%/d) | NW $t$ | Mean (%/d) | NW $t$ | Mean (%/d) | NW $t$ | Mean (%/d) | NW $t$ |
| SMB (size) | 1.04 | 3.10 | 0.77 | 2.75 | 0.84 | 3.29 | 0.81 | 3.14 |
| $\mathrm{HML}_{\mathrm{EMIS}}$ | 2.04 | 4.15 | 0.26 | 1.74 | 0.24 | 1.53 | 0.29 | 1.99 |

Long-short factor returns when the sorting characteristic is observed at lag 1 through lag 5 days before the return, holding the return date fixed. A persistent, harvestable premium survives the delay; a contemporaneous mechanical effect does not. The size premium is statistically indistinguishable across lags (NW $t \geq 2.75$ throughout). The emission-yield premium falls by roughly 87% from lag 1 to lag 2 and is no longer reliably significant at longer lags, identifying it as the AMM's within-period price impact (Proposition 1 applied to the contemporaneous flow) rather than a tradeable premium. Newey-West $t$-statistics (5 lags). Sample: February 2025 through June 2026.

Table VIII: Fama-MacBeth Cross-Sectional Regressions: Full-Sample and Out-of-Sample Betas

| Variable | Risk Premium (%/day) | Std Error (%) | $t$-statistic | $p$-value |
|---|---|---|---|---|
| *Panel A: Full-sample first-pass betas* | | | | |
| Intercept | 0.02 | 0.11 | 0.18 | 0.859 |
| $\beta_{\text{MKT}}$ | 0.23 | 0.14 | 1.68 | 0.095 |
| $\beta_{\text{SMB}}$ | 0.62 | 0.21 | 2.89 | 0.004 |
| $\beta_{\text{EMIS}}$ | 0.26 | 0.23 | 1.15 | 0.249 |
| $\beta_{\text{MOM}}$ | 0.28 | 0.23 | 1.19 | 0.233 |
| *Panel B: Out-of-sample betas (first-half betas, second-half premia)* | | | | |
| Intercept | 0.17 | 0.08 | 2.13 | 0.034 |
| $\beta_{\text{MKT}}$ | 0.03 | 0.04 | 0.78 | 0.435 |
| $\beta_{\text{SMB}}$ | 0.10 | 0.11 | 0.96 | 0.338 |
| $\beta_{\text{EMIS}}$ | $-0.01$ | 0.08 | $-0.07$ | 0.944 |
| $\beta_{\text{MOM}}$ | 0.06 | 0.10 | 0.60 | 0.549 |

Two-pass Fama-MacBeth regressions. First pass: time-series regressions of each subnet's daily return (TAO) on four factors (MKT, SMB, $\text{HML}_{\text{EMIS}}$, $\text{WML}_7$). Second pass: daily cross-sectional regressions of returns on estimated betas; the table reports time-series averages of the slopes with Fama-MacBeth standard errors. Panel A estimates first-pass betas over the full sample (look-ahead). Panel B estimates first-pass betas on the first half and prices them in the second, eliminating look-ahead. The size beta-premium is significant in Panel A but not in Panel B, while the size *characteristic* earns $t = 4.09$ in the same second half (Table X): returns track the characteristic, not the covariance (Daniel and Titman, 1997).

Table IX: Spanning Tests: Factor Alphas

| Factor | Alpha (%/day) | $t$-statistic | $R^2$ |
|---|---|---|---|
| MKT | −0.19 | −2.90 | 0.47 |
| SMB | 0.09 | 2.08 | 0.95 |
| $HML_{EMIS}$ | 1.02 | 9.28 | 0.54 |
| $WML_7$ | 0.30 | 2.67 | 0.42 |
| $WML_{30}$ | −0.15 | −1.38 | 0.45 |
| REV | −0.50 | −4.65 | 0.26 |
| LIQ | −0.03 | −0.70 | 0.95 |
| STAKE | 0.55 | 6.48 | 0.71 |

Each row reports the intercept (alpha) from regressing the given factor on all others. A significant alpha indicates information beyond the remaining factors. The emission-yield factor has by far the largest alpha. The high $R^2$ for SMB and LIQ (0.95) reflects their 0.98 correlation; the significant SMB alpha and insignificant LIQ alpha indicate size, not liquidity per se, is the fundamental characteristic.

Table X: Subsample Analysis

| | Full Sample | | First Half | | Second Half | |
|---|---|---|---|---|---|---|
| Factor | Mean (%/day) | $t$ | Mean (%/day) | $t$ | Mean (%/day) | $t$ |
| MKT | 0.33 | 1.61 | 0.43 | 1.07 | 0.24 | 3.08 |
| SMB | 1.04 | 3.42 | 1.45 | 2.49 | 0.62 | 4.09 |
| $HML_{EMIS}$ | 2.04 | 6.46 | 2.93 | 4.83 | 1.15 | 7.58 |
| $WML_7$ | 0.55 | 3.39 | 0.82 | 2.71 | 0.29 | 2.32 |
| $WML_{30}$ | 0.40 | 2.70 | 0.69 | 2.47 | 0.14 | 1.12 |
| REV | −0.49 | −2.86 | −0.68 | −2.14 | −0.29 | −2.45 |

The sample is split at its midpoint. First half: February 15 to October 16, 2025 (243 days). Second half: October 17, 2025, to June 15, 2026 (241 days). The halving occurred December 14, 2025, in the second half. $t$-statistics use OLS standard errors.

Table XI: The Halving: Size Premium Around December 14, 2025

| | Pre Mean (%/day) | Post Mean (%/day) | Ratio (Post/Pre) | Welch $p$ | NW $p$ |
|---|---|---|---|---|---|
| *Panel A: Full sample (302 pre, 182 post)* | | | | | |
| SMB | 1.39 | 0.45 | 0.321 | 0.062 | 0.090 |
| *Panel B: SMB, symmetric windows* | | | | | |
| ±30 days | 1.46 | 0.24 | 0.168 | 0.037 | 0.034 |
| ±45 days | 1.03 | 0.04 | 0.041 | 0.023 | 0.010 |
| ±60 days | 1.14 | 0.14 | 0.122 | 0.011 | 0.021 |
| ±90 days | 1.08 | 0.43 | 0.397 | 0.052 | 0.064 |

Mean daily SMB returns before and after the halving. Welch $p$ tests equality of pre- and post-means. In Panel A the NW $p$ is from a regression of daily SMB on a post-halving dummy (Newey-West, 5 lags). In Panel B the NW $p$ is from the same regression with a market control, within the symmetric window. The size premium declines after the halving, in the direction Proposition 1 predicts, but the change is not significant in HAC terms, the windows give noisy magnitudes (post/pre ratio 0.04 to 0.40), and adding a linear time trend to the full-sample regression renders the post dummy insignificant ($t = -0.31$). The halving also coincides with the switch to flow-based emission allocation. The paper does not rely on the halving for identification; see Section D.

Table XII: Robustness of the Structural Premia: Value Weighting and Entry

| Specification | SMB | | $\text{HML}_{\text{EMIS}}$ | |
|---|---|---|---|---|
| | Mean (%/day) | NW $t$ | Mean (%/day) | NW $t$ |
| *Panel A: Weighting* | | | | |
| Equal-weighted (baseline) | 1.04 | 3.10 | 2.04 | 4.15 |
| Value-weighted | 1.12 | 3.31 | 2.22 | 4.35 |
| *Panel B: Entry filters* | | | | |
| Baseline | 1.04 | 3.10 | 2.04 | 4.15 |
| Exclude age < 30 days | 0.84 | 3.36 | 1.31 | 6.01 |
| Exclude age < 60 days | 0.98 | 4.71 | 1.13 | 5.21 |
| Fixed cohort (launch subnets) | 0.96 | 3.38 | 2.11 | 4.03 |

Robustness of the size premium and, for completeness, the contemporaneous emission-yield sort (a one-day effect, not a harvestable premium; see Table VII). Panel A compares equal- and value- (market-cap-) weighted long-short portfolios. Panel B excludes recently listed subnets at age thresholds, and restricts to the fixed cohort of 64 subnets present at the dTAO launch (no entrants by construction). Newly listed subnets are 15.3% of the small tercile at the 30-day threshold and 27.8% at 60 days. Newey-West $t$-statistics (5 lags). Sample: February 2025 through June 2026.

Table XIII: AMM Slippage and Net-of-Cost SMB Returns by Portfolio AUM

| AUM | One-Way Slippage (%) | | | RT Cost | Net SMB | Net |
|---|---|---|---|---|---|---|
| | Small | Medium | Large | (%/day) | (%/day) | Sharpe |
| \$10K | 0.59 | 0.03 | 0.01 | 0.60 | 0.41 | 1.56 |
| \$100K | 5.90 | 0.26 | 0.08 | 5.98 | −4.97 | −19.0 |
| \$1M | 58.98 | 2.57 | 0.81 | 59.79 | −58.8 | — |
| \$10M | 590 | 25.7 | 8.1 | 598 | — | — |

One-way slippage is $\Delta\tau/\tau$, per-subnet investment (TAO) over the subnet's TAO reserves. AUM is converted from USD at the daily TAO/USD price and spread across eligible subnets. Round-trip (RT) cost sums buy-side slippage on small-tercile subnets and sell-side slippage on large-tercile subnets. Net SMB subtracts RT cost from the gross SMB return (1.04%/day). Assumes 100% daily turnover (worst case). Net Sharpe is annualized.

Table XIV: Volatility and Risk-Sorted Tercile Portfolio Returns

| Sort Variable | Low (%/day) | Mid (%/day) | High (%/day) | H−L (%/day) | $t$(H−L) |
|---|---|---|---|---|---|
| Total Volatility | −0.01 | 0.30 | 0.57 | 0.58 | 4.20 |
| Downside Vol | 0.07 | 0.24 | 0.55 | 0.48 | 3.74 |
| Upside Vol | −0.01 | 0.34 | 0.53 | 0.54 | 3.98 |
| Idiosyncratic Vol | 0.00 | 0.30 | 0.56 | 0.56 | 4.09 |
| Market Beta | 0.05 | 0.24 | 0.57 | 0.52 | 4.09 |
| Skewness | 0.19 | 0.21 | 0.46 | 0.28 | 2.57 |

Subnets are sorted daily into terciles on the 30-day rolling value of each characteristic. Returns are daily means in TAO. High-minus-low is positive for every risk measure: riskier subnets earn higher returns, the opposite of the equity low-volatility anomaly. $t$-statistics use OLS standard errors. $N = 454$ days (30-day formation window).

Table XV: Factor Return Risk Decomposition

| Factor | Mean (%/d) | Std (%/d) | Down Dev (%/d) | Up Dev (%/d) | Down/Up Ratio | Sortino (ann.) | % Neg Days |
|---|---|---|---|---|---|---|---|
| MKT | 0.33 | 4.56 | 3.57 | 5.20 | 0.69 | 1.79 | 43 |
| SMB | 1.04 | 6.66 | 4.01 | 8.18 | 0.49 | 4.93 | 42 |
| EMIS | 2.04 | 6.96 | 2.27 | 8.91 | 0.26 | 17.18 | 36 |
| $MOM_7$ | 0.55 | 3.57 | 3.38 | 3.78 | 0.90 | 3.13 | 43 |
| $MOM_{30}$ | 0.40 | 3.15 | 3.04 | 3.26 | 0.93 | 2.51 | 43 |
| REV | −0.49 | 3.75 | 3.30 | 4.40 | 0.75 | −2.82 | 60 |

Downside (upside) semi-deviation is $\sqrt{E[\min(r,0)^2]}$ ($\sqrt{E[\max(r,0)^2]}$). A Down/Up ratio below 1 indicates favorable asymmetry. Sortino uses downside deviation as the risk measure. % Neg is the fraction of days with negative returns. SMB has favorable asymmetry (0.49); the EMIS row is the contemporaneous one-day sort, shown for completeness, and its Sortino is not investable.

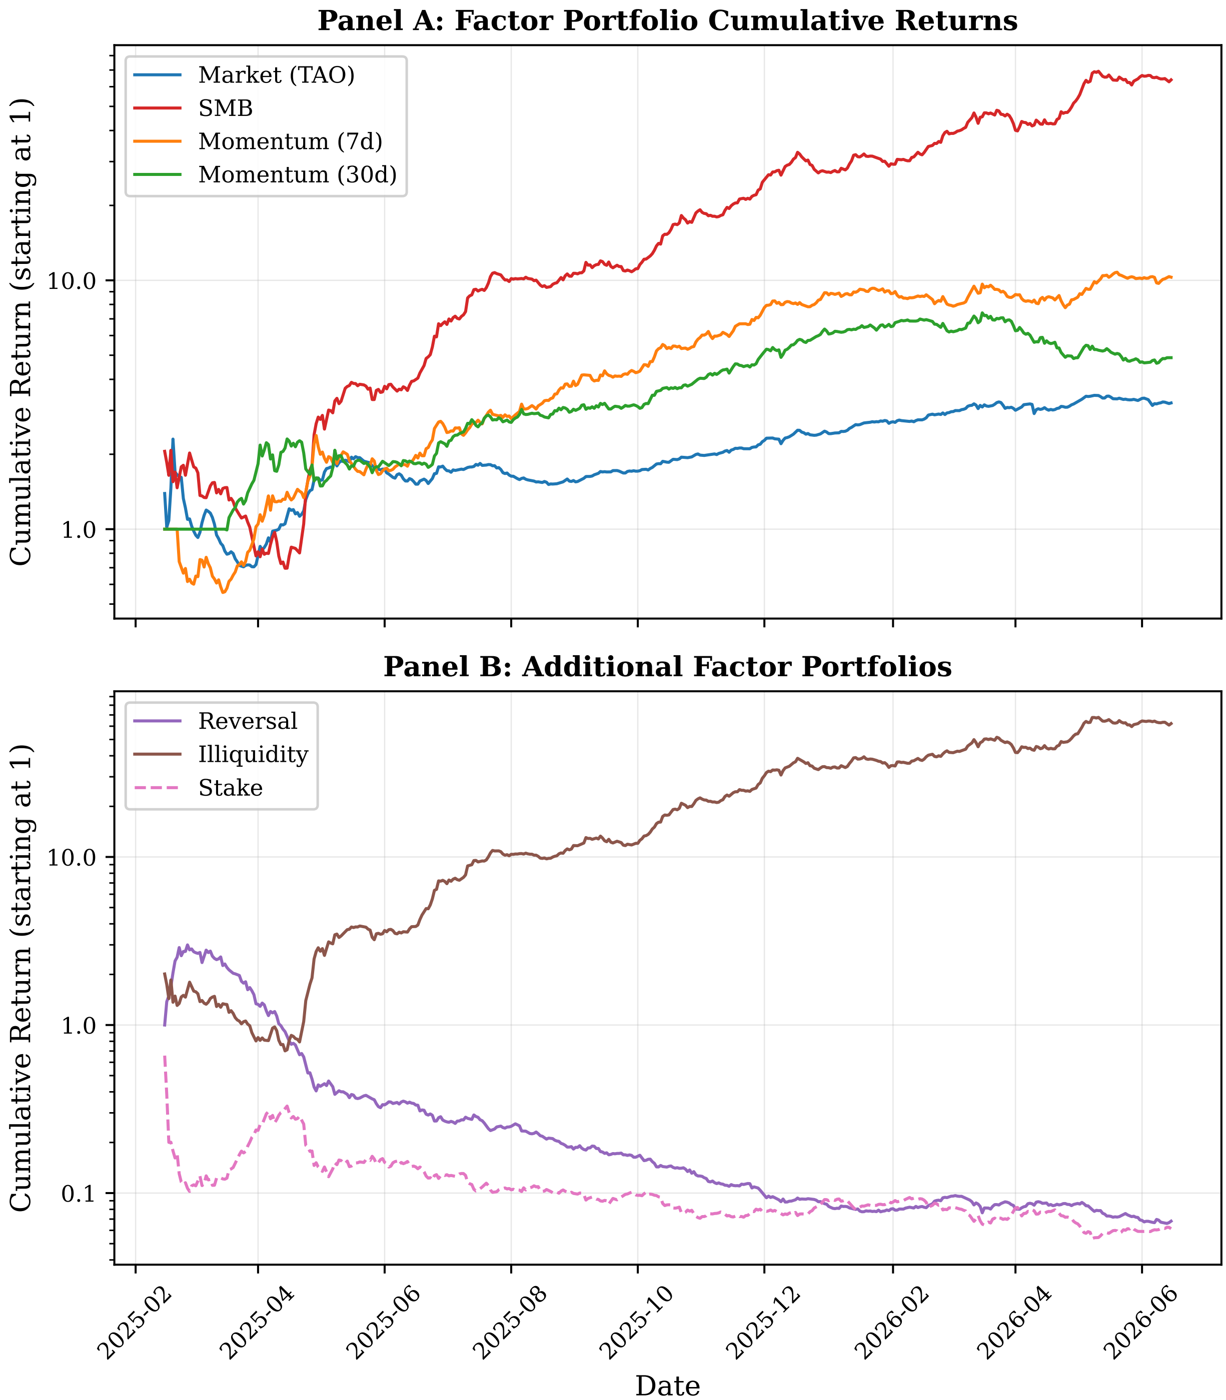


Figure 1: Cumulative Returns of Factor Portfolios

Panel A plots cumulative returns ($1 invested at inception) for the market (MKT), small-minus-big (SMB), 7-day momentum ($\mathrm{WML}_7$), and 30-day momentum ($\mathrm{WML}_{30}$) factors. Panel B plots reversal (REV), illiquidity (LIQ), and stake (STAKE). Returns are in TAO, $y$-axis log scale. Sample: February 2025 through June 2026.

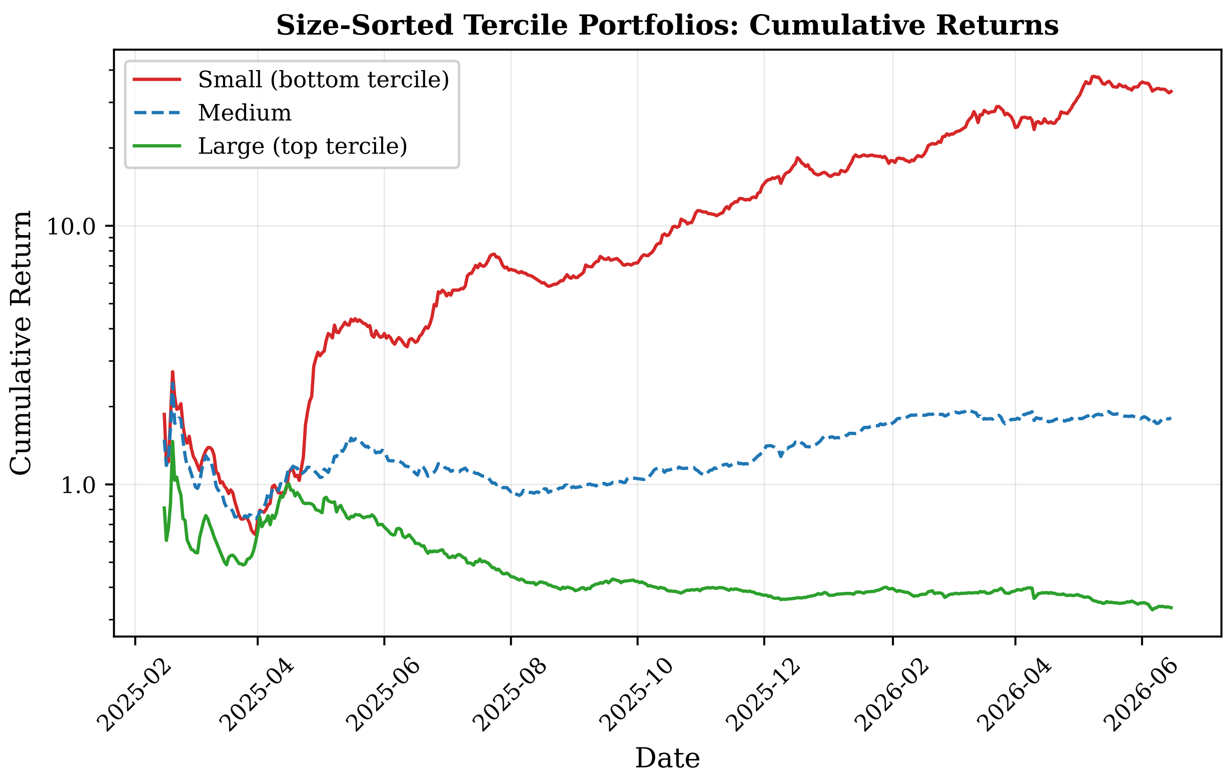


Figure 2: Cumulative Returns of Size-Sorted Tercile Portfolios
Subnets are sorted daily into terciles on lagged market capitalization. Portfolios are equal-weighted. The small tercile outperforms the large tercile by more than an order of magnitude over the sample. Returns are in TAO, log scale.

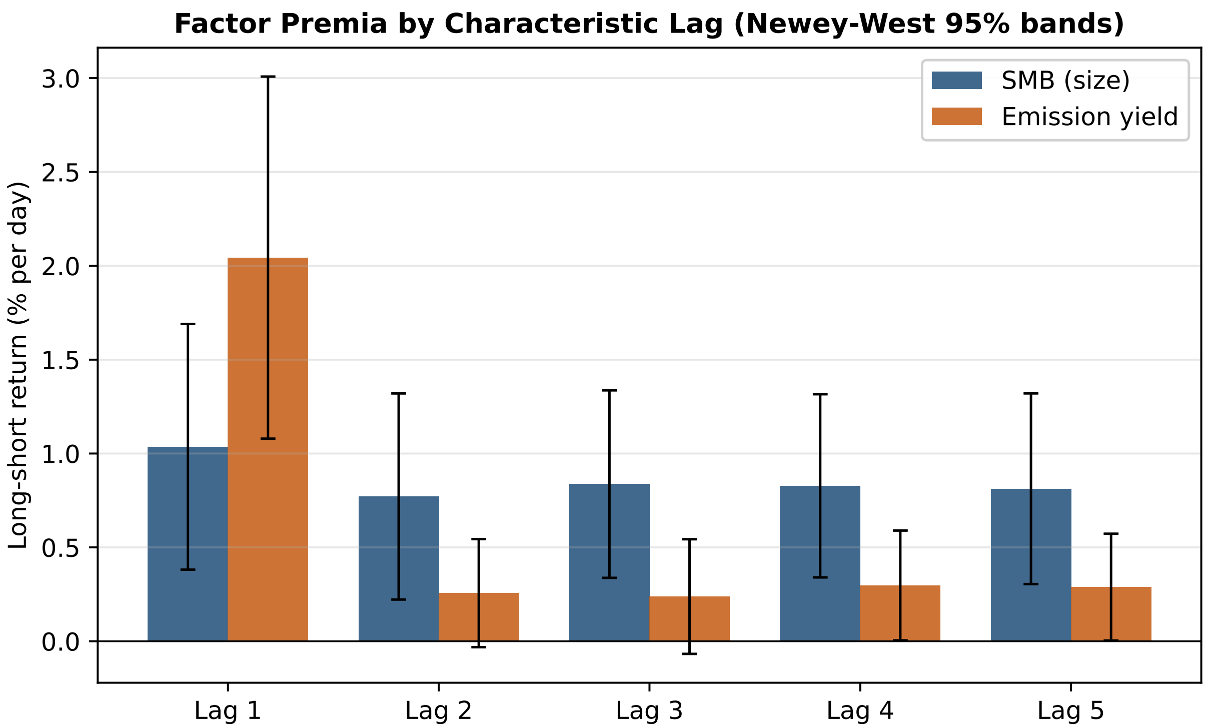


Figure 3: Persistence Diagnostic: Factor Premia by Characteristic Lag
Long-short factor returns (% per day, with Newey-West 95% confidence bands) when the sorting characteristic is observed one to five days before the return. The size premium (SMB) is statistically indistinguishable across lags, the signature of a persistent, harvestable premium. The emission-yield premium is large at the one-day lag but collapses by roughly 87% at a two-day lag and is no longer reliably significant, the signature of a contemporaneous mechanical price impact (Proposition 1 applied to the within-period flow) rather than a tradeable premium. Sample: February 2025 through June 2026.

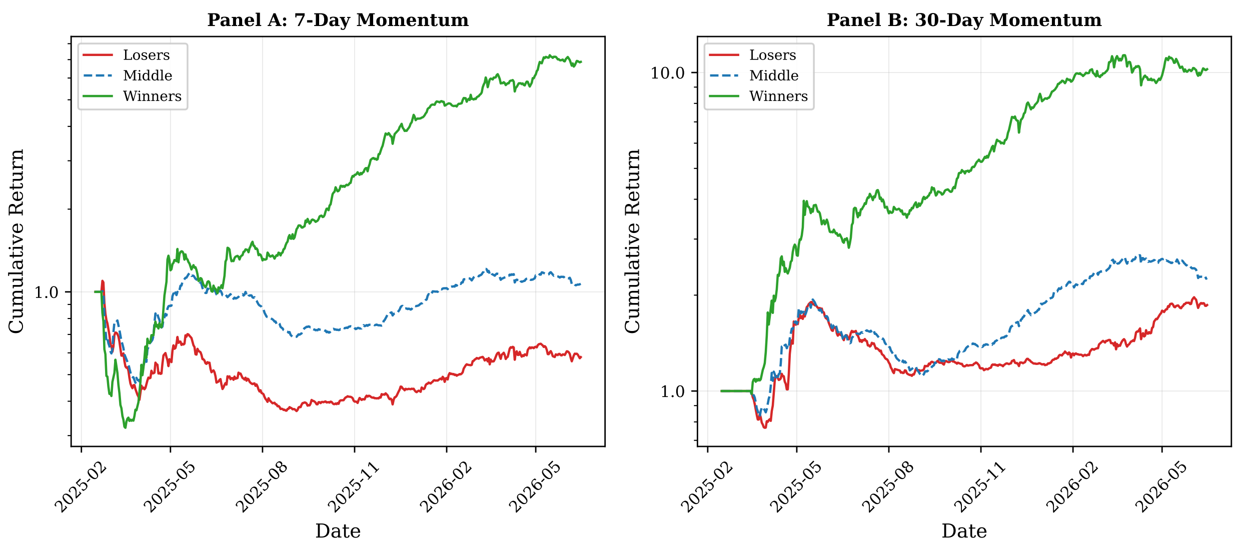


Figure 4: Cumulative Returns of Momentum-Sorted Tercile Portfolios

Panel A sorts subnets daily into terciles on 7-day past returns, Panel B on 30-day past returns. Winner subnets outperform losers. Returns are in TAO, log scale.

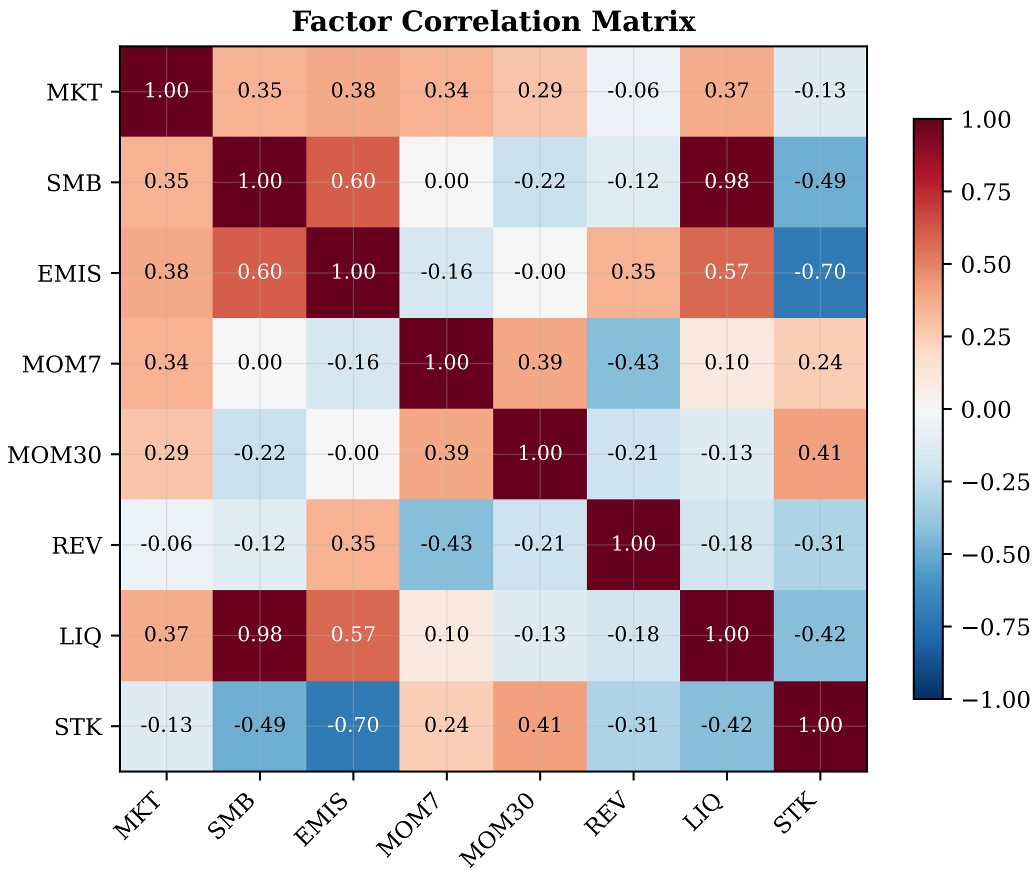


Figure 5: Factor Correlation Matrix

Pairwise correlations among daily factor returns. The 0.98 correlation between SMB and LIQ reflects the AMM structure linking market cap to pool depth. Momentum and reversal are negatively correlated, as expected.

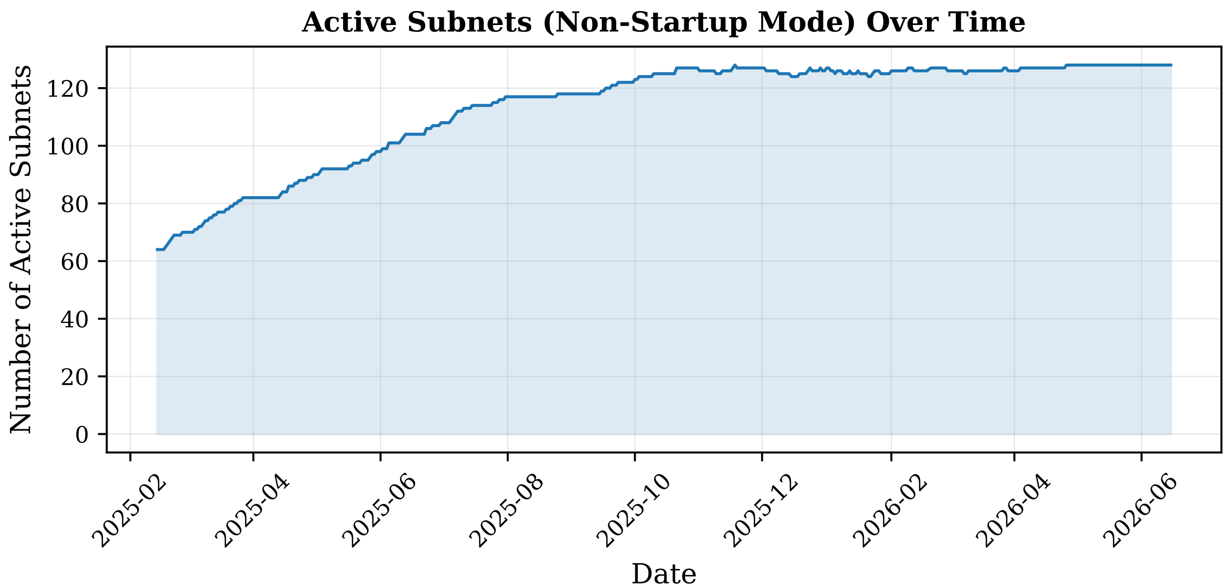


Figure 6: Number of Active Subnets Over Time

Count of non-startup-mode subnets in the daily cross-section. The eligible cross-section expands from 64 early in the sample (February 2025) to 128 by June 2026.

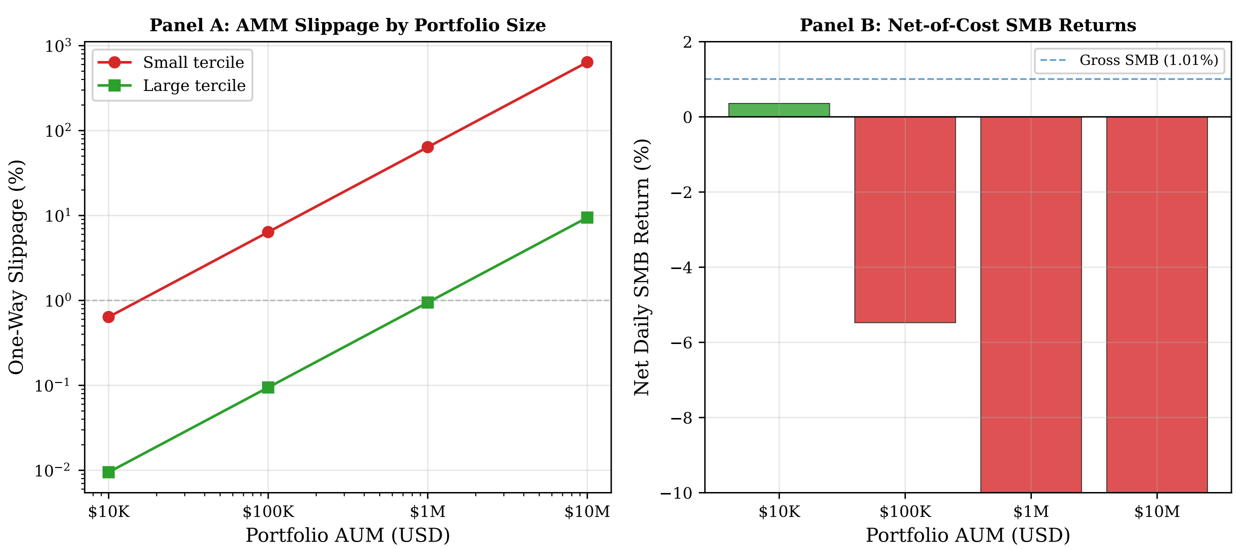


Figure 7: AMM Slippage and Net-of-Cost SMB Returns by Portfolio AUM

Panel A: one-way slippage (%) for small- and large-tercile subnets, log-log scale. Slippage equals $\Delta\tau/\tau$, linear in trade size relative to pool depth. Panel B: net-of-cost daily SMB after subtracting round-trip slippage. The strategy is profitable only near \$10K AUM; at \$100K and above slippage exceeds the gross return. Dashed line shows gross SMB (1.04%/day).

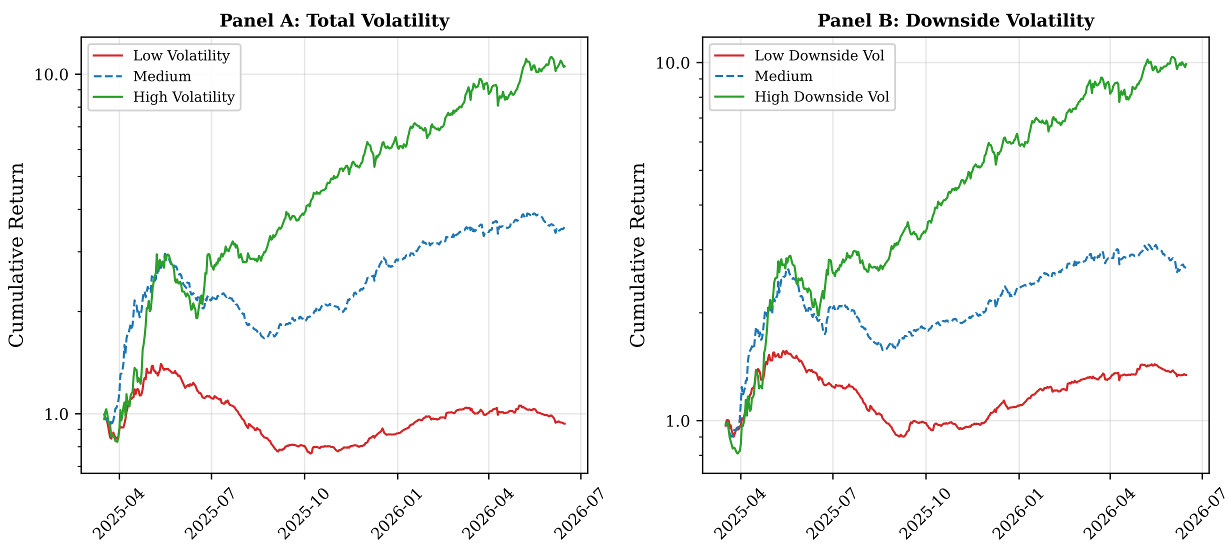


Figure 8: Cumulative Returns of Volatility-Sorted Tercile Portfolios

Panel A sorts subnets daily into terciles on 30-day realized volatility, Panel B on 30-day downside semi-deviation. High-risk subnets outperform low-risk subnets, the opposite of the equity low-volatility anomaly, reflecting the mechanical size-volatility link in constant-product AMMs.

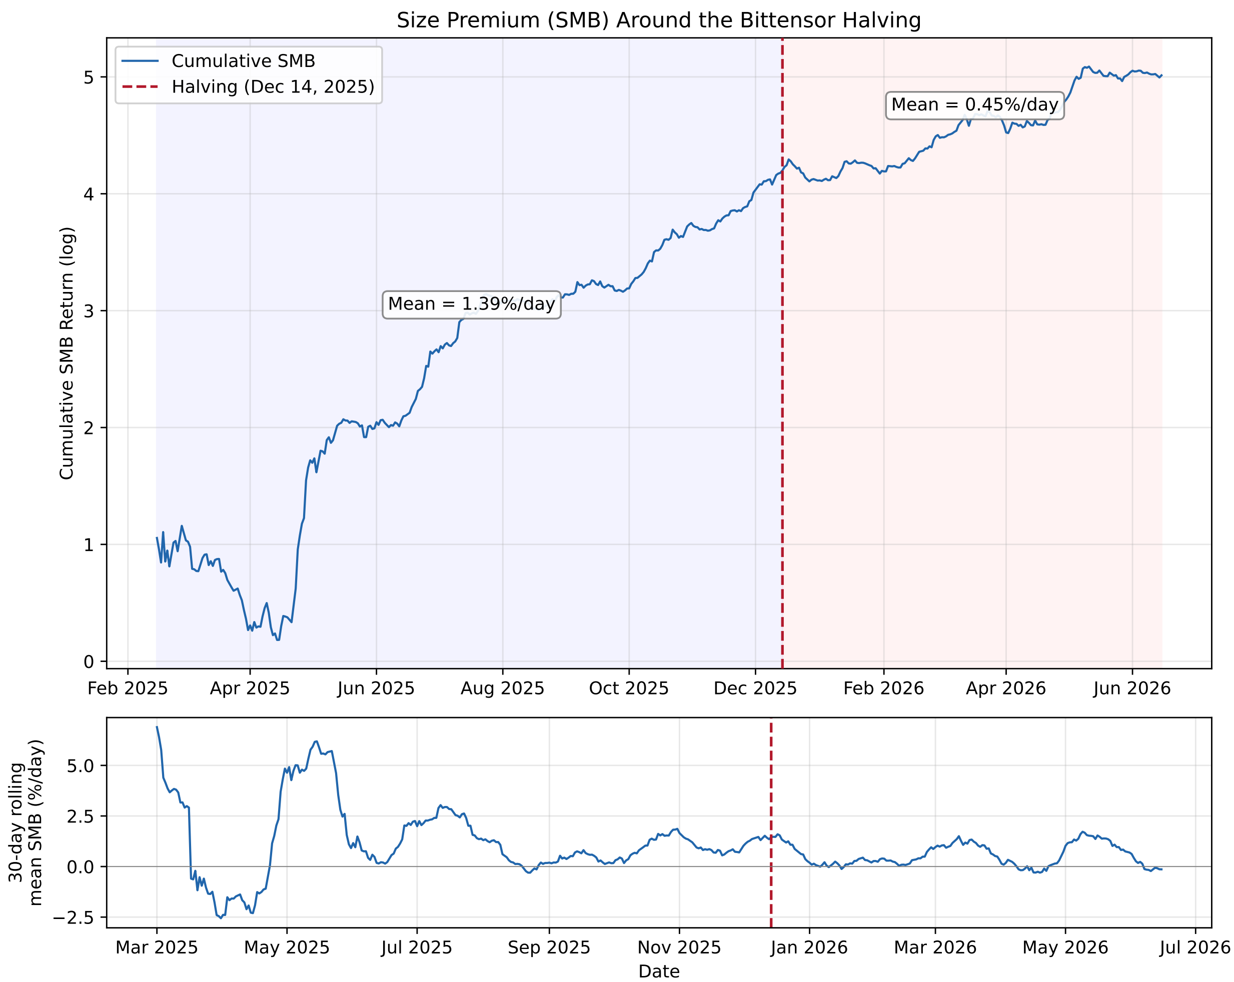


Figure 9: Cumulative SMB Return and the TAO Halving

Cumulative daily SMB (log scale, top) with a vertical line at the halving (December 14, 2025), and a 30-day rolling mean of daily SMB (bottom). The slope flattens after the halving, in the direction Proposition 1 predicts, but Section D shows the SMB decline cannot be cleanly separated from a maturation trend, so the halving is reported as directional corroboration only.

# References


Adams, Hayden, Noah Zinsmeister, and Dan Robinson, 2020, Uniswap v2 core, Technical report, Uniswap.

Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.

Ang, Andrew, Joseph Chen, and Yuhang Xing, 2006, Downside risk, *Review of Financial Studies* 19, 1191–1239.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen, 2013, Value and momentum everywhere, *Journal of Finance* 68, 929–985.

Bianchi, Daniele, and Mykola Babiak, 2021, A factor model for cryptocurrency returns, CERGE-EI Working Paper 710.

Borri, Nicola, and Kirill Shakhnov, 2022, The cross-section of cryptocurrency returns, *Review of Asset Pricing Studies* 12, 667–705.

Capponi, Agostino, and Ruizhe Jia, 2021, The adoption of blockchain-based decentralized exchanges, Working paper, Columbia University.

Daniel, Kent, and Sheridan Titman, 1997, Evidence on the characteristics of cross-sectional variation in stock returns, *Journal of Finance* 52, 1–33.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

Fieberg, Christian, Gerrit Liedtke, Thorsten Poddig, Thomas Walker, and Adam Zaremba, 2025, A trend factor for the cross section of cryptocurrency returns, *Journal of Financial and Quantitative Analysis* 60, 3116–3153.

Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.

Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121–1152.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

Lehar, Alfred, and Christine A. Parlour, 2025, Decentralized exchange: The Uniswap automated market maker, *Journal of Finance* 80, 321–374.

Liu, Yukun, and Aleh Tsyvinski, 2021, Risks and returns of cryptocurrency, *Review of Financial Studies* 34, 2689–2727.

Liu, Yukun, Aleh Tsyvinski, and Xi Wu, 2022, Common risk factors in cryptocurrency, *Journal of Finance* 77, 1133–1177.

Lui, Elizabeth, and Jiahao Sun, 2025, Bittensor protocol: The Bitcoin in decentralized artificial intelligence? A critical and empirical analysis, arXiv:2507.02951.

Moskowitz, Tobias J., Yao Hua Ooi, and Lasse Heje Pedersen, 2012, Time series momentum, *Journal of Financial Economics* 104, 228–250.

Park, Andreas, 2023, The conceptual flaws of decentralized automated market making, *Management Science* 69, 6731–6751.

Rao, Yuma, 2020, Bittensor: A peer-to-peer intelligence market, Technical report, Opentensor Foundation.

Steeves, Jacob, Ala Shaabana, Yuqian Hu, Francois Luus, Sin Tai Liu, and Jacqueline Dawn Tasker-Steeves, 2022, Incentivizing intelligence: The Bittensor approach, Workshop paper, Decentralization and Trustworthy Machine Learning in Web3, NeurIPS 2022.