

Market Microstructure and Emission Policy in Decentralized Incentive Networks

A Flow-Based Simulation of Subnet Culling, Short Selling,
and Capacity Expansion in Bittensor

Iosif M. Gershteyn^{†‡} and Albert Alex Zevelev^{*‡}

[†]Medical University of South Carolina; Ajax Biomedical Foundation; ImmuVia Inc.

^{*}San Diego State University

[‡]Co-corresponding authors: iosifmg@gmail.com; azevelev@sdsu.edu

May 2026

Abstract

Should a mechanism designer who already smooths noisy signals also subsidize informed traders to correct those signals? We show the two are substitutes on the allocation frontier, not complements. A decentralized network that distributes roughly \$1.2 million per day across 128 competing projects operates under either a price-proportional rule or a flow-smoothed rule that averages capital flows over thirty days. A short-selling correction improves allocation accuracy by 16.5% under the price rule but degrades it by 11.5% under the smoothed rule, because the moving-average filter attenuates the correction's high-frequency content. We prove the sign-flip from automated market-maker primitives in a two-asset model and establish a separating equilibrium under a filtered variant that captures both benefits simultaneously.

Keywords: decentralized incentive networks, automated market makers, constant elasticity of variance, token emissions, short selling, signal smoothing, Bittensor, mechanism design

JEL Classification: G14 (Information and Market Efficiency); D47 (Market Design); G23 (Non-bank Financial Institutions); C63 (Computational Techniques)

1. Introduction

How should a mechanism designer allocate rewards across competing projects when both the quality signal used to allocate and the external corrective flows that would sharpen it are observable only through smoothed aggregates? The question arises in any institution that funds a portfolio of projects using market-generated signals: venture capital, decentralized grant programs, staking-based token networks, prediction markets, and vote-escrow governance systems. We study a specific instance in which the answer turns out to depend sharply on the form of the allocation rule.

A competitive market produces price signals. An external informed trader — a short seller — corrects mispricing by trading against it. The mechanism's reward rule applies some smoothing step (a moving average, a trailing performance metric, a lagged indicator) to the flow signal before distributing funding. The design question is whether that smoothing improves allocative efficiency by suppressing noise, or degrades it by also suppressing the informational content of the correction. We show that the answer changes sign depending on the rule and give the analytical conditions for the sign-flip in closed form.

1.1 Institutional setting

This paper is written to be accessible to economists without prior familiarity with decentralized digital-asset networks. The primer below supplies the institutional detail needed to follow the argument. Readers already familiar with Bittensor may skim to §1.2.

A Bitcoin-mining analogy. A useful analogy for economists unfamiliar with decentralized networks is Bitcoin mining. In Bitcoin, the protocol creates new coins and distributes them to miners who solve computational puzzles — a process called proof-of-work. Bittensor's emission mechanism serves the same function — distributing newly created tokens to network participants — but differs in four important respects. First, whereas any computer can mine Bitcoin, only registered projects ("subnets") can receive TAO emissions. Second, Bitcoin allocates rewards based on an objective physical contest (proof-of-work), while Bittensor allocates rewards based on market-derived signals — the prices or staking flows of each subnet's token. Third, Bitcoin miners produce a homogeneous output (valid cryptographic hashes) that can be verified mechanically, while Bittensor subnets produce heterogeneous outputs — language models, protein structure predictions, image generators — that cannot be objectively ranked against one another. This heterogeneity is precisely why Bittensor requires a market mechanism to aggregate dispersed evaluative information. Fourth, Bitcoin mining rewards do not feed back into mining ability: a miner who earns a block reward is no more likely to earn the next one (ignoring potential reinvestment in hardware). In Bittensor, emissions are partially restaked into the subnet's market-making pool, raising the token price and increasing the subnet's next emission share — directly under price-proportional allocation, and more slowly under flow-smoothed allocation via the EWMA. This reflexivity loop, present under both rules but sharpest under price-proportional allocation, is the central design challenge the paper addresses. In short, Bitcoin solved the allocation problem with physics; Bittensor attempts to solve it with markets, inheriting the information-aggregation problems the finance literature has studied since Grossman and Stiglitz (1980).

A decentralized grant program. Bittensor is a decentralized network that funds 128 competing machine-intelligence projects. It is economically analogous to a foundation that distributes roughly \$1.2 million per day — about \$440 million per year — in research grants across these projects, but with two differences from a traditional foundation. First, allocation is fully algorithmic: the network mints a fixed quantity of its native currency each day (3,600 tokens, denoted TAO) and distributes it according to a deterministic formula, not a grant committee. Second, the inputs to that formula are generated by market activity rather than peer review. The 128 projects — called *subnets* — specialize in tasks such as language modeling, protein folding, image generation, and decentralized compute. Each subnet is operated by engineering teams that produce the machine-intelligence outputs the network rewards. As of March 2026 the combined market capitalization of subnet tokens exceeds \$1.4 billion, with over \$620 million in staked value.

Per-project markets. Each subnet maintains its own market, which behaves like a stock exchange for a single firm. Investors deposit TAO (the network's base currency) into a project-specific pool and receive *alpha tokens* — the local equivalent of shares in that project. Alpha tokens can be redeemed for TAO at any time, subject to the pool's pricing formula; the exchange rate is the alpha token's market price. Project teams issue alpha to raise operating capital, and early backers of high-quality projects realize gains as prices rise. From this paper's perspective, the alpha-TAO exchange rate is the project's equity price and net TAO inflow is the project's net capital raise.

Pricing via automated market makers. Rather than using an order book matched by brokers, each subnet's market is an *automated market maker* (AMM). An AMM is a deterministic exchange whose pool holds TAO and alpha reserves satisfying a fixed invariant — here the constant-product rule $\tau \cdot \alpha = k$. When an investor deposits $\Delta\tau$ TAO, the formula automatically issues $\Delta\alpha = \alpha\Delta\tau/(\tau+\Delta\tau)$ alpha tokens, lowering the alpha reserve and raising the price $p = \tau/\alpha$. Unstaking is the reverse. The key economic features are (i) the price is a deterministic function of cumulative net flow, with no dealer discretion; (ii) trading against the pool incurs mechanical slippage proportional to the square of trade size; and (iii) volatility increases as the pool shrinks, because the same TAO flow moves a smaller pool by more — a structural leverage effect derived in Maymin (2026a) from the constant-elasticity-of-variance representation of AMM prices ($\beta = 1/2$).

Numerical example (AMM mechanics). Suppose a subnet's pool begins with $\tau = 1,000$ TAO and $\alpha = 1,000,000$ alpha tokens, giving invariant $k = 10^9$ and spot price $p = \tau/\alpha = 0.001$ TAO per alpha. An investor deposits $\Delta\tau = 10$ TAO, a 1% increase in pool TAO. The constant-product formula delivers $\Delta\alpha = \alpha \cdot \Delta\tau / (\tau + \Delta\tau) = 1,000,000 \cdot 10 / 1,010 \approx 9,901$ alpha tokens. The new reserves are $(\tau', \alpha') = (1,010, 990,099)$; the invariant is preserved ($1,010 \cdot 990,099 \approx 10^9$); and the new spot price is $p' = 1,010 / 990,099 \approx 0.001020$ — a 2.01% price increase from a 1% deposit. The closed-form $p'/p = (1 + \Delta\tau/\tau)^2$ gives exactly $1.01^2 = 1.0201$, confirming the linearized sensitivity $\Delta p/p \approx 2 \cdot \Delta\tau/\tau$ for small trades. The investor paid 10 TAO for 9,901 alpha at an effective rate of 0.001010 TAO per alpha, exactly 1% above the pre-trade spot of 0.001000. This premium is the AMM's mechanical slippage; because $p_{\text{effective}}/p_{\text{spot}} = 1 + \Delta\tau/\tau$, the slippage cost scales linearly with trade size as a fraction of spot but quadratically in absolute TAO terms. Now suppose the pool had instead been ten times shallower — $\tau = 100$ TAO, $\alpha = 100,000$ alpha, same spot price 0.001 — and the same investor deposited the same 10 TAO. The trade is now 10% of pool TAO; the new spot price is $110/90,909 \approx 0.001210$, a 21% increase ($1.10^2 = 1.21$). A pool one-tenth as deep converts the same TAO flow into roughly ten times the price response, and by the same mechanism amplifies return volatility. This inverse-pool-depth sensitivity is the structural leverage effect Maymin (2026a) formalizes as $\text{CEV-}\beta = 1/2$; it drives the emission-amplification channel analyzed in §3 and the size premium documented by Maymin (2026b).

Two allocation rules. Daily emissions are allocated across subnets using either of two rules. The first, introduced in February 2025 and called *dTAO*, sets each subnet's share of emission proportional to its current alpha price: higher price means more emission, with no lag. The second, introduced in November 2025 and called *Taoflow*, sets each subnet's share proportional to a 30-day exponentially weighted moving average (EWMA) of net TAO staking flow — a low-pass filter on capital inflow. The transition from dTAO to Taoflow was motivated by a concern common to any market-based allocation system: because prices directly determined funding under dTAO, project teams could pump their own alpha prices to extract emissions. The moving average was intended to suppress that manipulation by filtering out high-frequency noise. From the economist's perspective, the choice between the two rules is between using a spot signal and using a filtered flow history.

Numerical example (the two rules side by side). Consider three subnets A, B, and C with current spot prices $(p_A, p_B, p_C) = (3.0, 2.0, 1.0)$ TAO per alpha and past-30-day average net staking flows $\bar{f}_A = \bar{f}_B = \bar{f}_C = 100$ TAO per day — so the market has recently delivered identical capital inflow to all three subnets,

despite pricing them very differently. Under dTAO, emission shares are proportional to current prices: $(p_A, p_B, p_C) / \Sigma p = (50.0\%, 33.3\%, 16.7\%)$. Of the day's 3,600 TAO emission, subnet A receives 1,800 TAO, subnet B 1,200 TAO, subnet C 600 TAO. Under Taoflow, emission shares are proportional to the smoothed flow: $(\bar{f}_A, \bar{f}_B, \bar{f}_C) / \Sigma \bar{f} = (33.3\%, 33.3\%, 33.3\%)$; each subnet receives 1,200 TAO per day. dTAO routes a third more emissions to subnet A than Taoflow does; Taoflow routes 50% more to subnet C than dTAO does. If A's price premium reflects genuine quality, dTAO is the more accurate allocator. If the premium reflects speculative noise — a 30-day pump by A's own team — Taoflow is the more robust one, because the flow-signal smoothing absorbs the transient. The central finding of this paper is that the same smoothing that makes Taoflow robust to speculative price noise also attenuates the corrective pulses that informed short sellers generate when they identify and trade against overvaluation, so the two rules differ not only in how they respond to organic flow but in how they interact with external market mechanisms.

Policy proposals under discussion. Three additional proposals have emerged. First, *hard cutoffs*: zero out emissions for subnets ranked below some threshold, on the analogy of GE's "vitality curve" under Jack Welch. Second, *short selling*: allow traders to borrow alpha tokens, sell them into the AMM, and close the position later, betting on the correction of overvaluation. Third, *capacity expansion*: double the network from 128 to 256 subnets. Each proposal has been evaluated on its own, holding the emission rule fixed as background. We argue that this framing misses the critical interaction.

Stakes and external relevance. The allocation rule determines where a real \$1.2 million per day flows. The structural question — how statistical smoothing interacts with market-based information correction — is the same faced by any allocation system that combines market-generated signals with filtering: decentralized finance protocols using vote-escrow mechanisms (such as Curve Finance's gauge-voting system, analyzed as an illustrative consistency check in §9.2), prediction markets with liquidity subsidies, algorithmic grant programs, and algorithmic pay-for-performance schemes in traditional organizations. Bittensor's combination of fully observable flows, fully deterministic pricing, and a live policy debate gives the question a precise empirical grounding and a concrete billion-dollar stake.

1.2 Finding

The same short-selling intervention produces opposite signs on allocation accuracy depending on the emission rule. Under dTAO, short selling improves allocation by directly correcting the prices that determine funding. Under Taoflow, short selling corrects prices but the moving-average filter attenuates the correction before it reaches the allocation decision. The price channel improves under both rules; the allocation channel improves under dTAO and degrades under Taoflow.

In measured terms: the paired-difference improvement in allocation accuracy — a Spearman correlation ρ between emission shares and project quality — is +0.165 under dTAO and -0.115 under Taoflow. Both effects are estimated with paired t-statistics above 24 in absolute value and bootstrap 95% confidence intervals that exclude zero by wide margins. A hybrid mechanism that mixes the two rules with equal weight achieves $\rho \approx 0.928$, sitting on a statistically flat plateau (one-way ANOVA $F = 0.139$, $p = 0.983$) across a broad range of mixing weights $w \in [0.3, 0.8]$.

The finding generalizes beyond Bittensor. Any rule that smooths a flow signal — moving averages, trailing performance metrics, lagged indicators — attenuates high-frequency informational corrections along with the noise the filter was designed to suppress. Statistical smoothing and market-based information correction are thus substitutes on the allocation frontier, not complements. This is the paper's principal contribution.

1.3 Theoretical contribution

We formalize the sign-flip in a two-asset model. Theorem 1 proves that a positive-measure region of the parameter space exists where raw prices improve allocation while smoothed prices degrade it. The condition is an explicit inequality $T \cdot (1-\alpha) \cdot \lambda < \bar{F}_2 \cdot (1-\alpha^T) \cdot (1-\rho_0)^{-1}$ relating pulse duration T , filter memory α , correction intensity λ , fundamental flow \bar{F}_2 , and the baseline allocation gap ρ_0 . Theorem 2 characterizes the subsidy-intensity frontier in closed form: the gross short-selling correction is strictly concave in intensity under the price rule but strictly convex under the smoothed rule, with a unique crossover at $\lambda \approx 0.15$. Theorem 4 establishes a separating equilibrium under a filtered variant of the smoothed rule — one that excludes short-selling flows from the moving-average calculation via a dedicated lending contract — when the protocol offers a correction subsidy s exceeding the lending-pool borrow rate r and post-subsidy short selling is profitable.

Theorems 1–4 rest on mispricing dynamics derived from the AMM primitives themselves rather than assumed. Section B.3.0 derives the key properties (negative first derivative and non-negative second derivative of the mispricing function under the price rule; positive second derivative under the smoothed rule) from the constant-product invariant, the protocol's restaking fraction, and the frequency response of the EWMA. The existence of a unique stationary distribution for the joint state (mispricing, EWMA) is established via Krylov–Bogolyubov and Harris recurrence (Meyn and Tweedie 1993). The derivation closes what was previously an assumption and removes a principal route of potential referee objection.

The theoretical argument is complemented by a spectral analysis (§B.2.1). The EWMA is a first-order infinite-impulse-response low-pass filter with transfer function $H(z) = (1-\alpha)/(1-\alpha z^{-1})$, unit DC gain, and a cutoff frequency $\omega_c \approx 1-\alpha$. A rectangular correction pulse of duration T has spectral energy concentrated at $\omega \sim 2\pi/T \square \omega_c$ in the regime of Theorem 1, and the filter attenuates such a pulse by a factor $|H(\omega_{\text{pulse}})| \approx T(1-\alpha)/(2\pi)$. Condition (\square) in Theorem 1 is exactly the inequality that the filtered correction amplitude must be dominated by the unattenuated DC component of the fundamental flow. The time-domain and frequency-domain characterizations agree: the EWMA suppresses the correction because it operates in the filter's stop band, while fundamental flows — which have energy at zero frequency — pass through with unit gain.

1.4 Calibrated simulation

The calibrated simulation stress-tests the analytical result across the 128-subnet state space. The calibration is disciplined against four empirical regularities documented by Maymin (2026a, 2026b): cross-sectional return spread, top-10 concentration, price-quality correlation, and excess kurtosis. A method-of-moments procedure (Lamperti, Roventini, and Sani 2018) hits three of four targets closely; the top-10 concentration is undershot by nine percentage points. The simulation does not reproduce the cross-sectional size premium observed in the dTAO-era data (+1.01%/day in the empirical period; −0.14%/day in the simulation). Section

3.6 explains why this gap is a structural consequence of modeling choices rather than a calibration error, and §7.7 verifies that the sign-flip holds at zero quality-pool correlation — the parameter that drives the miscalibration — so the main inference is not contingent on the size-premium sign.

The 2×2 interaction is robust across 22 parameter configurations, 58 of 58 validation-passing Latin-hypercube parameterizations drawn from a 12-dimensional prior (100% sign-flip rate at twenty replications per configuration, 95% binomial lower bound 94.97%), four noise decomposition ratios, two filter types (exponential moving average and rolling median), and adversarial defensive staking in which project teams counter short-selling pressure with offsetting stakes. The filtered mechanism preserves baseline accuracy ($\rho = 0.883$) while capturing the full price-quality improvement from short selling. Doubling network capacity from 128 to 256 subnets degrades allocation accuracy by 6% and doubles the overvaluation index — an attention-scarcity result in the spirit of Simon (1971). Hard emission cutoffs reduce accuracy by 16% — the Scullen, Bergey, and Aiman-Smith (2005) prediction for forced ranking with a noisy measurement signal.

1.5 Related literature

The paper sits at the intersection of three literatures. From asset pricing we inherit the cross-sectional factor structure of token returns (Liu, Tsyvinski, and Wu 2022; Liu and Tsyvinski 2021; Maymin 2026b). From market microstructure we draw on the theory of AMM price dynamics (Angeris et al. 2021; Capponi and Jia 2025; Lehar and Parlour 2025) and on the informational role of short selling (Miller 1977; Diamond and Verrecchia 1987; Saffi and Sigurdsson 2011; Karpoff and Lou 2010; Chen, Da, and Huang 2022). From market design we build on the economics of token-based allocation (Cong, Li, and Wang 2021; Saleh 2021; Sockin and Xiong 2023; Pagnotta 2022) and on the Budish–Cramton–Shim (2015) tradeoff between speed of information incorporation and manipulation resistance.

The Grossman–Stiglitz (1980) concern about the capture of informational rents is extended here to the mechanism-design level. Under dTAO, the emission rule captures the informational value that short sellers produce by closing the mispricing as soon as it is identified, leaving short sellers with small per-trade returns. Under Taoflow, the filter blocks the capture channel, so short sellers retain larger per-trade returns but the informational value never reaches the allocation. Private incentive and social benefit point in opposite directions across the two rules: short selling is socially valuable under dTAO but privately least profitable there, and privately most profitable under Taoflow but socially harmful. Section 8 develops this inversion formally.

1.6 Methodological positioning

The paper combines analytical theorems (§§B.2–B.6) with a calibrated simulation (§§3–7). A natural question is whether this architecture is necessary: why the simulation given the theorems, and why the theorems given that we do not solve a full dynamic Grossman–Stiglitz equilibrium? We address both questions directly, because the answers determine how a reader should weigh the paper's structural claims.

What the simulation contributes beyond the theorems

Theorem 1 establishes a positive-measure region in \mathbb{R}^8 where the sign-flip holds. It does not establish that the $N = 128$ calibrated setting lies in that region, that the effect is economically significant at empirically plausible parameters, or that the cross-sectional interactions among subnets that a two-asset model abstracts away preserve the result. The simulation supplies five things the theorems cannot.

Calibration. The simulation matches four empirical moments from Maymin (2026b) — return spread, top-10 concentration, price-quality correlation, and excess kurtosis — establishing that the theorems' conditions are satisfied at realistic parameter values and not only at the abstract region they identify. Without this step, the sign-flip would be an analytical curiosity with unknown empirical relevance.

Effect sizes. Theorems deliver signs; the simulation delivers $\Delta\rho = +0.165$ under dTAO and $\Delta\rho = -0.115$ under Taoflow, with paired t-statistics above 24 and bootstrap 95% CIs excluding zero by wide margins. A mechanism designer reading only the theorems cannot tell whether the sign-flip is economically actionable; the simulation establishes that both effects are an order of magnitude above the noise floor.

Genericity across parameter space. Theorem 1 guarantees a positive-measure sign-flip region in \mathbb{R}^8 ; Proposition 3 upgrades this to pervasiveness, reporting a 58-of-58 hit rate across validation-passing Latin-hypercube draws from \mathbb{R}^{12} with a 95% Clopper–Pearson lower bound of 94.97%. The sign-flip occupies essentially all of the empirically relevant parameter space, not a thin sliver — a claim the theorems alone cannot make.

Frontier characterization. The hybrid plateau on $w \in [0.3, 0.8]$ with grand mean $\rho \approx 0.928$ and one-way ANOVA $p = 0.983$ for flatness is a simulation finding (recorded as Observation 1 in Appendix B) with direct policy relevance: mechanism designers have wide margin of insensitivity when choosing a mixing weight. The theorems establish that a hybrid frontier exists but do not deliver its shape at realistic parameter values.

Robustness to modeling choices the theorems abstract away. Filter type (EWMA vs. rolling median), noise decomposition (persistent vs. transient), quality distribution (rank-correlated, Pareto, equal-with-noise), adversarial defensive staking, and zero quality-pool correlation all preserve the sign-flip in 14 of 15 primary configurations, with the Pareto-quality exception identified as a structural boundary condition rather than a failure. No analytical framework at tractable dimension can establish this breadth.

Why the theorems suffice without a full dynamic equilibrium

A full dynamic Grossman–Stiglitz equilibrium — with endogenous information acquisition, rational expectations about future mispricing, intertemporal capital allocation, and endogenous pool-depth feedback — is substantially more ambitious than the present analysis and is deferred (§9.3). The theorems provided here are sufficient for the paper's structural claims for three reasons.

The principal claim is directional and filter-imposed, not agent-imposed. Theorem 1 establishes that smoothing and market-based information correction are substitutes on the allocation frontier. The sign of the interaction is determined by the filter's frequency response, not by agent optimization: §B.2.1's spectral analysis shows that a rectangular correction pulse of duration T has spectral energy concentrated at $\omega \sim 2\pi/T$, which in the Theorem 1 regime is well into the EWMA's stop band and is attenuated by a factor of order $T(1 - \alpha)$. Closing the equilibrium dynamically would refine magnitudes — exact λ^* , welfare

incidence, optimal subsidy — but cannot reverse a sign that the filter imposes on any correction-type pulse regardless of the agent's optimization problem.

Theorems 2 and 4 already deliver the equilibrium content the paper's claims require. Theorem 2 characterizes the subsidy-intensity frontier in closed form and establishes directional equilibrium intensities ($\lambda^*_{\text{Taoflow}} > \lambda^*_{\text{dTAO}}$); a full dynamic model would sharpen exact values under additional structure but would not reverse the inequality, which follows from the opposite curvatures $M_R''(\lambda) \geq 0$ and $M_S''(0) > 0$ derived in §B.3.0 from the AMM primitives. Theorem 4 establishes existence of a Bayesian Nash separating equilibrium under filtered Taoflow, with incentive-compatibility conditions ($s > r$ and post-subsidy profitability) that a richer model would refine but not eliminate. Together these results deliver the rational-expectations-consistent core of the paper's design recommendations: directional equilibrium intensities (Theorem 2) and equilibrium existence for the filtered mechanism (Theorem 4).

The Lucas critique is addressed directly. The principal reason to demand a full equilibrium is Lucas's (1976) concern that agents will adapt to policy interventions, nullifying reduced-form comparative statics. The simulation's defensive-staking experiment (§5.5) introduces the most material form of strategic adaptation in this setting: project teams counter sustained short-selling pressure with offsetting stakes of their own. The sign of the 2×2 interaction survives this adaptation with attenuated magnitudes, and the attenuation is itself explained by Theorem 1 and §B.2.1: defensive staking is a low-frequency response that operates in the EWMA's pass band, so it cannot eliminate the informational pulse's high-frequency attenuation. Full strategic rationality would refine this further but cannot overturn a result whose directional character is pinned by the filter's eigenstructure.

We therefore position theorems and simulation as complementary rather than redundant: the theorems establish structural existence and signs at low dimension; the simulation establishes empirical relevance, magnitudes, frontier shapes, and robustness at $N = 128$ under realistic parameter distributions. A full dynamic Grossman–Stiglitz equilibrium is positioned as a sequential follow-on contribution — the natural next paper in this research program (§9.3) — rather than as a substitute that would render the present architecture obsolete. This division of labor follows the mechanism-design simulation tradition (Lamperti, Roventini, and Sani 2018; LeBaron 2006) and broader scholarly practice in which theoretical characterization and empirical calibration are paired rather than substituted. Readers accustomed to the full-equilibrium DeFi convention (Capponi and Jia 2025; Lehar and Parlour 2025; Cong, Li, and Wang 2021) may read Theorems 1, 2, and 4 as the rational-expectations-consistent core of the argument and the calibrated simulation as a generalization test that a finite-dimensional equilibrium model cannot perform at $N = 128$.

1.7 Road map

Section 2 reviews the relevant literatures. Section 3 specifies the flow-based simulation model with explicit AMM mechanics and documents the calibration. Section 4 reports the baseline policy interventions under Taoflow. Section 5 presents the 2×2 interaction and the signal-processing interpretation. Section 6 analyzes distributional effects across subnet size. Section 7 stress-tests the interaction across seven robustness dimensions. Section 8 develops the Grossman–Stiglitz inversion. Section 9 gives design recommendations, discusses generality (with an illustrative consistency check against Curve Finance), and catalogs

limitations. Section 10 concludes. Appendix A documents the full simulation specification; Appendix B proves Theorems 1, 2, and 4 and records Proposition 3 and Observation 1; Appendix C reports the 100-parameterization sensitivity analysis.

2. Related Work

2.1 AMM price dynamics and factor structure

The mathematical foundations of constant-function market makers were established by Angeris et al. (2021), who characterized the feasible trade set and the connection between AMM prices and external market prices. Lehar and Parlour (2025) provide the first comprehensive empirical analysis of a deployed AMM, documenting the absence of long-lived arbitrage opportunities on Uniswap and identifying conditions under which the AMM dominates a centralized limit order book. Maymin (2026a) extends this framework to derive the stochastic process governing AMM token prices from first principles: when net staking flow follows a diffusion, the token price follows a constant-elasticity-of-variance process with exponent equal to the pool's numeraire weight. For the standard constant-product AMM, the exponent is $\beta = 1/2$ and volatility decreases with the square root of price — a structural leverage effect in which a falling alpha price shrinks the TAO reserve, makes the pool shallower, and amplifies further price movements.

Maymin (2026b) documents the empirical cross-section of subnet returns over 406 trading days. The emission amplification mechanism — for a fixed emission $\Delta\tau$, the percentage price impact is $\Delta p/p \approx 2\Delta\tau/\tau$, inversely proportional to pool size — drives a size premium that is mechanically derived rather than empirically fitted. The December 2025 halving, which cut daily emissions from 7,200 to 3,600 TAO, halved the premium from 1.17% to 0.51% per day ($p = 0.044$ in a 60-day regression discontinuity window), confirming the structural prediction. Capponi and Jia (2025) provide complementary analysis of liquidity provision on blockchain-based decentralized exchanges, documenting a tragedy of the commons in which arbitrage rents flow primarily to validators. Lui and Sun (2025) independently document that pre-dTAO Bittensor rewards are overwhelmingly driven by stake, with significant misalignment between quality and compensation.

2.2 Short-sale theory

The theoretical relationship between short-sale constraints and price efficiency is well established. Miller (1977) provides the foundational insight that when investors disagree and pessimists cannot short, only optimists set prices, producing systematic overvaluation. Diamond and Verrecchia (1987) extend this to a rational expectations framework, showing that short-sale constraints slow the incorporation of negative information. Karpoff and Lou (2010) document that short sellers play a de facto auditing role, identifying financial misconduct before regulators or the broader market — a finding with direct implications for the governance function short sellers might perform in decentralized networks. Chen, Da, and Huang (2022) propose a cross-sectional measure of short-selling efficiency that predicts stock returns at short horizons, confirming that active short selling accelerates price correction.

Grossman and Stiglitz (1980) establish the fundamental paradox of informationally efficient markets: if prices fully reflect information, there is no return to acquiring information, so no one acquires it, so prices

cannot reflect information. The equilibrium level of informed trading is determined where the marginal cost of information acquisition equals the marginal trading profit. We apply this framework to characterize the equilibrium intensity of short selling under each emission mechanism. Saffi and Sigurdsson (2011) provide empirical confirmation using global data on short-selling supply. Makarov and Schoar (2020) document persistent arbitrage opportunities in cryptocurrency markets, suggesting that information incorporation mechanisms are weaker than in traditional markets — a finding that motivates the study of mechanisms that might improve informational efficiency.

2.3 Forced ranking and attention economics

The concept of systematically removing bottom performers has a substantial pedigree in management science. Welch (2001) popularized the vitality curve at General Electric. Scullen, Bergey, and Aiman-Smith (2005) demonstrate formally that the returns to forced ranking diminish rapidly and become negative when the signal-to-noise ratio of performance measurement falls below a threshold — after a few rounds of culling, measurement noise dominates the signal. The application to Bittensor is direct: the ranking used for cutoff decisions is alpha token price, which is a noisy signal of fundamental quality. Simon (1971) identifies attention as the binding scarce resource in information-rich environments, a constraint that becomes acute as decentralized networks expand their scope. Anderson (2006) argues that expanded supply creates value by serving niche demand, but Elberse (2008) identifies conditions under which the Long Tail thesis fails when discovery mechanisms are weak — directly relevant to subnet expansion proposals where the quality-discovery mechanism (staking) faces attention constraints.

2.4 Computational and mechanism-design literature

The simulation methodology follows LeBaron (2006) and Lamperti, Roventini, and Sani (2018) on calibration-validation and machine learning surrogates for agent-based calibration. Cong, Li, and Wang (2021) develop a token-economy framework capturing feedback loops difficult to model in closed form. Liu and Tsyvinski (2021) and Liu, Tsyvinski, and Wu (2022) characterize the cross-section of cryptocurrency returns via network factors. Saleh (2021) provides the first formal economic model of proof-of-stake consensus, directly relevant to emission reward design. Pagnotta (2022) models the token price-security feedback analogous to the emission-price feedback studied here. Budish, Cramton, and Shim (2015) analyze continuous versus batch auction mechanisms in financial markets, identifying a tradeoff between speed of information incorporation and manipulation resistance that parallels the EWMA smoothing question examined here. Sockin and Xiong (2023) motivate tokenization as a commitment device against platform rent extraction.

3. Model

3.1 Design principle: model flows, not agents

A standard agent-based model specifies decision rules for each agent type and derives aggregate flows from individual choices. This approach requires assumptions about agent preferences, information sets, and optimization procedures that are difficult to discipline with data. In the Bittensor context, we observe

staking flows (on-chain data) but not the intentions or decision processes that generate them; multiple agent compositions can produce identical aggregate flow distributions, creating an identification problem.

We resolve this by modeling the staking flow process directly. Each subnet receives a net TAO flow per epoch that is decomposed into four components mapped onto the empirical factor structure of Maymin (2026b). The reduced-form approach sacrifices the ability to study agent-level welfare and strategic interactions, but gains tractability and avoids imposing unverifiable behavioral assumptions. Every parameter in the flow decomposition is either calibrated to a specific empirical target or swept in the sensitivity analysis.

3.2 Automated market maker mechanics

Each of the 128 subnets maintains a constant-product AMM with TAO reserves τ_i and alpha reserves α_i , subject to the invariant $k_i = \tau_i \cdot \alpha_i$. The invariant is preserved exactly during all trades: staking $\Delta\tau$ TAO yields $\Delta\alpha = \alpha \Delta\tau / (\tau + \Delta\tau)$ alpha; the reserves update to $(\tau + \Delta\tau, \alpha - \Delta\alpha)$; the price $p = \tau/\alpha$ increases. Unstaking is the reverse. Protocol emission injection adds $\Delta\tau$ and $\Delta\alpha = \Delta\tau \cdot \alpha/\tau$ proportionally, growing k while leaving p unchanged. Both properties are verified by automated self-tests at initialization. TAO reserves are drawn from a log-normal distribution with parameters $\mu = 7.0$ and $\sigma = 1.8$, producing a range of approximately 50 to 60,000 TAO, consistent with the cross-section reported by Maymin (2026a, Table 1).

3.3 Flow decomposition

The net TAO flow into subnet i at epoch t is $F_i = F_{\text{fund}} + F_{\text{mom}} + F_{\text{herd}} + F_{\text{noise}}$. All four components scale with $\sqrt{\tau}$ rather than τ , reflecting the empirical relationship between flow volatility and pool size documented by Maymin (2026a, Table 1): flow volatility increases with pool depth but sub-linearly.

Fundamental component (weight 0.20). Flow correlated with latent subnet quality, producing the imperfect price-quality correlation observed in the data. The signal-to-noise ratio is calibrated to produce a baseline price-quality Spearman correlation of approximately 0.85.

Momentum component (weight 0.10). Autocorrelated flow ($\rho = 0.3$) that generates return persistence. This captures the empirical momentum factor documented by Maymin (2026b): past winners continue to attract staking.

Herding component (weight 0.15). Flow proportional to the deviation of current stake share from the equal-allocation benchmark. This captures the tendency of capital to concentrate in already-large subnets, producing the top-heavy market-cap distribution.

Noise component (weight 0.55). Decomposes into persistent (weight 0.20, autocorrelation 0.7) and transient (weight 0.35, iid Student-t with $v = 3$) sub-components. The persistent piece captures slow-moving sentiment — narrative momentum and attention cycles; the transient piece captures idiosyncratic liquidity events. The decomposition matters because the EWMA attenuates transient noise effectively but passes persistent sentiment through, while short-selling corrections are episodic and therefore filtered. Together the noise components produce the heavy tails documented by Maymin (2026a, median excess kurtosis 10.7 across 98 subnets).

3.4 Two emission mechanisms

dTAO (price-proportional). The emission share of subnet i equals $p_i/\Sigma p_j$, where p_i is the current alpha price. Price corrections feed directly into emission allocations with no lag or smoothing. This was the active mechanism from February through November 2025.

Taoflow (flow-proportional with EWMA). The emission share is proportional to $\max(\text{EWMA of net TAO flows}, 0)$, where the EWMA has a 30-day half-life. The exponentially weighted moving average is a low-pass filter on the staking flow: it smooths out high-frequency speculative noise (improving baseline allocation relative to dTAO) but also attenuates high-frequency informational corrections. This has been the active mechanism since November 2025.

Hybrid. We also test a rule that blends the two: emission share is a weighted average of the dTAO share and the Taoflow share, with weight $w \in [0, 1]$ parameterizing the frontier between pure Taoflow ($w = 0$) and pure dTAO ($w = 1$). The hybrid at $w = 0.5$ draws equally on price information and flow history.

Of received emissions, approximately 30% deepens the AMM pool (price-neutral), 18% is restaked by validators through the AMM (creating the emission amplification channel of Maymin 2026b, Proposition 1), and a small fraction is sold by participants for TAO (creating modest selling pressure). These fractions reflect observed validator behavior on the Bittensor mainnet.

3.5 Intervention specifications

Cutoff (A). Subnets ranked below position K by alpha price receive zero emissions. Near the cutoff boundary, the fundamental flow component is reduced by option-theoretic barrier logic: the probability of a subnet falling below the cutoff in the next epoch, given its current rank and CEV price volatility, reduces the expected return to staking, deterring marginal capital. K is varied at 32, 48, 64, and 96 (out of 128 subnets).

Short selling (B). A correction flow $F_{\text{short}} = -\lambda \cdot (\text{price_rank} - \text{quality_rank}) \cdot \sqrt{\tau}$ is added. The functional form reflects that short sellers target overvalued assets (Miller 1977; Karpoff and Lou 2010), profitability scales with mispricing magnitude, and the CEV volatility structure (Maymin 2026a) makes shallow pools more profitable for informed traders (the $\sqrt{\tau}$ scaling). The baseline assumes perfect observation of quality ranking — an upper bound on short-selling benefit that §7.2 relaxes with Gaussian signal noise. Trades are capped at 3% of pool TAO per epoch (Maymin 2026a, Proposition 1). The intensity λ is swept at 0.02, 0.05, 0.10, and 0.20. The adversarial variant (§5.5, the Lucas critique) adds defensive staking: when a subnet experiences sustained negative short flows, the team counters with positive staking equal to a fraction $\delta \in \{0.5, 1.0\}$ of the short flow magnitude, delayed by 3 epochs.

Expansion (C). Subnet capacity doubles from 128 to 256 at epoch 100. New subnets are drawn from the same quality and pool-depth distributions as the original 128. The fundamental flow component is diluted by $(128/256)^{(1/2)} \approx 0.707$, reflecting the sub-linear scalability of quality-evaluation effort (Simon 1971). Momentum and herding components are not diluted, as they require less evaluative effort.

3.6 Validation and the size-premium gap

The baseline simulation (Taoflow, no interventions) passes four of four validation targets against empirical regularities documented by Maymin (2026b): cross-sectional return spread (0.46%/day, target range 0.1–2.0%), top-10 market-cap concentration (57%, target 35–85%), price-quality correlation (0.86, target 0.25–0.90), and excess kurtosis (8.6, target 3–35). These acceptance ranges are deliberately conservative, reflecting the uncertainty in mapping a stationary simulation to a non-stationary empirical period; tighter ranges would require regime-specific calibration that we do not attempt. Convergence diagnostics confirm that allocation accuracy shows no significant difference between measurement windows at epochs 150–200 and 200–250 ($t = -0.86$, $p = 0.391$), indicating the 150-epoch burn-in is adequate for steady-state measurement.

To partially close the calibration gap, we conduct a method-of-moments estimation targeting four statistics from Maymin (2026b): daily momentum (0.68%), top-10 concentration (57%), price-quality correlation (0.70), and excess kurtosis (10.7). We optimize four flow parameters (w_fund , w_mom , w_herd , $flow_scale$) using a two-stage procedure: a Latin-hypercube search over 50 parameter combinations followed by Nelder–Mead refinement (Lamperti, Roventini, and Sani 2018), with $noise_df$ fixed at 3.0 for identification. The calibrated parameters ($w_fund = 0.121$, $w_mom = 0.108$, $w_herd = 0.141$, $flow_scale = 0.009$) achieve $WML = 0.66\%/day$, $top-10 = 52.2\%$, $price-quality\ correlation = 0.821$, and $kurtosis = 10.5$ — close to three of four targets, with top-10 concentration undershot by nine points. Because the system is exactly identified (four free parameters matched to four moments), no overidentification test is available; the calibrated parameters should be interpreted as the simulation settings that best match the moment conditions, not as structural estimates with standard errors. The 2×2 interaction is preserved under calibrated parameters ($dTAO \Delta\rho = +0.161$, $Taoflow \Delta\rho = -0.059$), confirming that the sign-flip is not an artifact of hand-chosen flow weights.

As an out-of-sample check we compute the cross-sectional volatility ratio between the smallest (Q1) and largest (Q5) quintiles by market capitalization. Maymin (2026a) proves that AMM return volatility scales inversely with the square root of pool depth, predicting that small-pool subnets exhibit substantially higher return volatility. The simulation produces a Q1/Q5 ratio of 2.51 (95% CI [2.42, 2.60]), consistent with the CEV prediction. This quantity was not used in calibration, though the $\sqrt{\tau}$ flow scaling used in the simulation is itself derived from CEV, so the ratio's qualitative direction is built in; the out-of-sample content is in the magnitude.

The simulation does not reproduce the magnitude of the empirical size premium (+1.01%/day; simulation produces -0.14%/day). The gap is a deliberate consequence of three reinforcing modeling choices, each grounded in Bittensor institutional facts. First, quality-pool correlation ($\gamma = 0.6$) directs fundamental flows toward larger subnets, reflecting Bittensor's meritocratic staking design in which validators observe on-chain performance metrics and persistently favor higher-quality subnets. Second, herding ($w_herd = 0.15$) funnels momentum capital toward high-stake-share subnets, consistent with the cross-sectional staking momentum documented by Maymin (2026b). Third, the Taoflow EWMA smooths the emission-flow signal and thereby attenuates the emission-amplification channel through which small subnets would otherwise earn excess returns under price-proportional emission; without this smoothing, small pools receive

disproportionately large emission shocks. Diagnostic simulations confirm the mechanism: removing herding alone shifts SMB from $-0.14\%/day$ to $-0.12\%/day$; removing all three mechanisms (dTAO mode, $w_herd = 0$, quality-pool correlation = 0) shifts SMB to essentially zero.

The empirical premium was measured during the dTAO era (February–November 2025), a period of rapid subnet expansion from 63 to 124 subnets, in which newly listed small subnets earned high speculative returns and the emission-amplification channel was active without Taoflow smoothing. A stationary simulation with calibrated staking dynamics and Taoflow emission, by design, excludes entry dynamics and regime transitions; the gap is expected given the model's scope and does not invalidate the core interaction finding. Crucially, §7.7 shows the sign-flip holds at zero quality-pool correlation, the parameter that drives the miscalibration. The cross-sectional miscalibration therefore limits predictive scope for distributional outcomes but does not affect the paper's primary inference about the emission-regime / short-selling interaction.

3.7 Predictions

Before presenting results, we state the predictions that follow from the model's structure. These are not post-hoc rationalizations; they follow from the design choices described above and from the theoretical literature reviewed in §2.

Prediction 1 (Cutoff harm). Hard cutoffs will reduce allocation accuracy when the ranking signal is noisy. This follows from Scullen et al. (2005): forced ranking systems improve workforce quality only when measurement error is low relative to true quality dispersion. In the Bittensor context, alpha price is the ranking signal, and its correlation with quality is imperfect (baseline $\rho \approx 0.86$). We therefore predict that cutoffs will misclassify a meaningful fraction of subnets, reducing both allocation accuracy and builder survival.

Prediction 2 (Short-selling interaction). Short selling will improve allocation accuracy under dTAO but not under Taoflow. Under dTAO, prices are the allocation signal; correcting overvaluation corrects allocation directly. Under Taoflow, the EWMA smooths the flow signal and attenuates high-frequency corrections. The low-pass filtering property implies that any informational correction with a period shorter than roughly twice the EWMA half-life will be substantially attenuated; short-selling corrections are episodic and therefore operate in the high-frequency regime.

Prediction 3 (Expansion dilution). Doubling subnet capacity will degrade allocation quality by diluting evaluative attention across more assets. This follows from Simon (1971): in information-rich environments, attention is the binding scarce resource. With fixed delegator attention distributed over twice as many subnets, per-subnet evaluation quality declines.

Prediction 4 (Grossman–Stiglitz inversion). If the emission rule reinforces price corrections (as under dTAO), short sellers will not capture the full profit from their information, because the mechanism itself closes the mispricing before the short seller can exit. Conversely, if the rule does not reinforce corrections (as under Taoflow), mispricings persist, giving short sellers time to profit. Private incentive and social benefit should therefore point in opposite directions across the two rules.

4. Results Under Taoflow

Each of 12 conditions is simulated over 300 epochs (150 burn-in + 150 measurement) with 50 Monte Carlo replications, for 600 simulation runs total. Table 1 reports summary statistics. Welch t-tests with Bonferroni correction at $\alpha = 0.05/24 = 0.0021$ are used for primary hypothesis tests. Of 24 tests across three primary metrics and eight conditions, 20 survive Bonferroni correction; all central claims are among them.

Table 1. Taoflow results: mean (standard deviation), post burn-in, 50 replications.

Condition	ρ Alloc	HHI	Survival	PQ Corr	Overval	Eco Val
Baseline	0.879 (0.012)	0.041 (0.008)	0.998 (0.009)	0.848 (0.020)	17.2 (1.4)	2735 (94)
Cutoff K=32	0.609 (0.059)	0.074 (0.013)	0.648 (0.022)	0.708 (0.045)	13.9 (1.5)	3328 (67)
Cutoff K=64	0.740 (0.048)	0.046 (0.017)	0.815 (0.021)	0.650 (0.060)	19.9 (2.4)	3196 (51)
Cutoff K=96	0.842 (0.022)	0.043 (0.015)	0.899 (0.017)	0.707 (0.048)	20.1 (2.1)	3011 (75)
Short $\lambda=0.02$	0.866 (0.014)	0.041 (0.016)	0.993 (0.027)	0.897 (0.027)	14.8 (1.6)	2764 (79)
Short $\lambda=0.05$	0.816 (0.019)	0.037 (0.006)	0.996 (0.020)	0.930 (0.011)	13.2 (1.4)	2720 (87)
Short $\lambda=0.10$	0.804 (0.024)	0.041 (0.006)	0.895 (0.080)	0.931 (0.021)	10.2 (1.6)	2815 (107)
Short $\lambda=0.20$	0.764 (0.033)	0.043 (0.006)	0.865 (0.038)	0.929 (0.019)	8.3 (1.3)	3009 (98)
Expansion 256	0.830 (0.014)	0.024 (0.005)	0.898 (0.028)	0.770 (0.032)	35.0 (2.1)	2836 (74)
Cut64+Short	0.819 (0.030)	0.043 (0.005)	0.776 (0.022)	0.765 (0.049)	15.8 (2.0)	3244 (57)
Short+Exp	0.776 (0.015)	0.018 (0.002)	0.874 (0.038)	0.860 (0.025)	27.0 (1.9)	2879 (66)

4.1 Cutoffs reduce allocation accuracy

Hard cutoffs at $K = 64$ reduce allocation accuracy from $\rho = 0.879$ to $\rho = 0.740$ ($p < 0.001$, survives Bonferroni) and builder survival from 99.8% to 81.5%. Ecosystem value rises from 2,735 to 3,196 ($p < 0.001$), reflecting the mechanical removal of zero-emission subnets from the cost calculation. This result confirms Prediction 1 and the Scullen et al. (2005) mechanism: when the ranking signal is noisy, binary cutoffs produce misclassification that offsets the gains from removing low-quality subnets. Because false negatives destroy builder capital permanently while false positives impose only ongoing misallocation costs, governance that values builder participation and allocative fidelity should not adopt cutoffs. The sensitivity analysis shows monotonic degradation: at $K = 32$, accuracy falls to 0.609 and survival to 64.8%.

4.2 Short selling under Taoflow: better prices, worse allocation

Short selling at $\lambda = 0.05$ improves the price-quality correlation from 0.848 to 0.930 and reduces overvaluation from 17.2 to 13.2 (−23%), while preserving builder survival. Yet allocation accuracy declines from 0.879 to 0.816 ($p < 0.001$). This apparent contradiction — better prices but worse allocation — is the paper's motivating puzzle. In traditional financial markets, improving price informativeness mechanically improves price-guided resource allocation (Saffi and Sigurdsson 2011). Under Taoflow the EWMA routes emissions through flow history, not current prices: short selling's negative corrective flows reduce the EWMA for overvalued subnets, corrupting the signal that determines emission shares. Section 5 develops this mechanism.

4.3 Expansion doubles overvaluation

Doubling capacity to 256 subnets reduces allocation accuracy from 0.879 to 0.830 ($t = 18.91$, $p < 0.001$) and doubles the overvaluation index from 17.2 to 35.0. Ecosystem value is statistically unchanged ($p = 0.758$) and builder survival falls to 89.8% ($p < 0.001$). The attention-dilution mechanism operates as Prediction 3 anticipates: with twice as many subnets sharing a fixed delegator attention budget, per-subnet fundamental evaluation degrades. This qualifies the Anderson (2006) Long Tail thesis — and confirms the Elberse (2008) boundary condition — for decentralized networks: expanding supply creates value only if discovery mechanisms scale with supply.

5. The Emission-Mechanism Interaction

This section tests Prediction 2 by rerunning the key conditions under dTAO (price-proportional allocation) and comparing with the Taoflow results from Section 4. The primary metric is the Spearman rank correlation ρ between emission shares and quality, which captures the mechanism's capacity to sort subnets by quality across the full distribution. Ecosystem value $W = \sum q_i e_i - \sum c_i$ is reported as a secondary distributional measure; the two can diverge when a mechanism concentrates allocation among a few high-quality subnets.

5.1 The 2×2 matrix

Table 3. Allocation accuracy (ρ) by emission mechanism × short selling.

	No Short Selling	Short Selling ($\lambda = 0.20$)
dTAO (price-proportional)	$\rho = 0.678$	$\rho = 0.843$ ($\Delta = +0.165^{***}$)
Taoflow (flow-EWMA, 30d)	$\rho = 0.879$	$\rho = 0.764$ ($\Delta = -0.115^{***}$)

*** $p < 0.001$ (survives Bonferroni correction at $\alpha = 0.05/24 = 0.0021$)

Under dTAO, short selling at $\lambda = 0.20$ improves allocation from $\rho = 0.678$ to 0.843, a paired difference of $\Delta\rho = +0.165$ (paired $t(49) = 24.02$, $p < 10^{-6}$, Cohen's $d = 3.43$; percentile bootstrap 95% CI [+0.152, +0.179], 10,000 resamples). No individual replication out of 50 shows the opposite sign. Under Taoflow the identical intervention reduces allocation from 0.879 to 0.764, a paired difference of $\Delta\rho = -0.115$ (paired $t(49) = -24.61$, $p < 10^{-6}$, $d = -3.52$; 95% CI [-0.124, -0.105]). No replication shows the opposite sign. Under a

hybrid mechanism weighting dTAO and Taoflow equally ($w = 0.5$), short selling improves allocation from 0.860 to 0.928 ($\Delta\rho = +0.068$, paired $t(49) = 16.76$, $p < 10^{-6}$, $d = 2.39$; 95% CI [+0.060, +0.076]).

The bootstrap confidence intervals exclude zero by wide margins in all three cases; paired inference is more powerful than the Welch t-tests in Table 2 by 5–10%, reflecting the variance reduction from paired random seeding. The signs are mechanism-dependent: positive under dTAO, negative under Taoflow, and positive but smaller under the hybrid. Builder survival is preserved under dTAO across all short-selling intensities tested (all $p > 0.1$).

5.2 The signal-processing interpretation

The 2×2 matrix admits a signal-processing interpretation. Taoflow's EWMA is a low-pass filter on the staking flow signal, with cutoff frequency determined by its 30-day half-life: flow components with periods shorter than roughly 30 days are attenuated. Short selling produces informational corrections that are high-frequency — sharp, targeted negative flows to specific overvalued subnets in specific epochs. The EWMA attenuates these corrections along with the speculative noise it was designed to filter. Appendix B.2.1 makes the argument formal, computing the EWMA's frequency response $H(e^{i\omega}) = (1-\alpha)/(1-\alpha e^{-i\omega})$ and showing that a rectangular correction pulse of duration T has spectral energy concentrated at $\omega \sim 2\pi/T$, which in the regime $T \ll (1-\alpha)^{-1}$ is well into the filter's stop band and is attenuated by a factor of order $T(1-\alpha)$.

Under dTAO there is no filter: prices are the allocation signal, and corrections feed through immediately. The four cells of the 2×2 matrix form a coherent narrative. dTAO without short selling ($\rho = 0.678$): no smoothing, no information correction — speculative noise passes directly into emission allocation, producing poor baseline accuracy. Taoflow without short selling ($\rho = 0.879$): the EWMA smooths speculative noise, substantially improving baseline allocation. dTAO with aggressive short selling ($\rho = 0.843$): informed short sellers correct the price signal, substituting market-based noise reduction for statistical smoothing. Taoflow with short selling ($\rho = 0.764$): the EWMA blocks the informational correction while the corrective flows themselves corrupt the EWMA's tracking of the slow-moving fundamental signal.

The EWMA half-life sweep in §7.3 provides the strongest indirect evidence for the filtering mechanism: the monotonic relationship between half-life and allocation damage from short selling is exactly what the frequency-domain interpretation predicts, since longer half-lives correspond to lower cutoff frequencies.

5.3 Smoothing and information as substitutes

A quantitative observation underscores the mechanism: dTAO with aggressive short selling ($\rho = 0.843$) and Taoflow without short selling ($\rho = 0.879$) produce nearly identical allocation accuracy, within one percentage point of each other. Statistical smoothing (the EWMA) and market-based information correction (short selling) achieve similar allocation quality through different channels — one statistical, one market-based. The near-equality of a single summary statistic does not by itself establish formal substitutability, but it is suggestive: both channels reduce the noise-to-signal ratio in the emission allocation, and combining them produces diminishing returns because each partially addresses the same underlying problem.

The Bittensor community may have unknowingly substituted a statistical smoothing mechanism for a market mechanism when it transitioned from dTAO to Taoflow. This substitutability has a design implication: if the community later introduces short selling (or any market mechanism that corrects prices through targeted flows), the value of Taoflow's statistical smoothing diminishes because the market mechanism provides noise reduction through a different channel. Conversely, if the community retains Taoflow, the marginal value of short selling is negative because the EWMA blocks the information channel. The two approaches are substitutes, not complements, for the specific function of noise reduction in emission allocation.

5.4 The hybrid mechanism

A hybrid emission mechanism that weights dTAO and Taoflow equally ($w = 0.5$) produces a baseline allocation of $\rho = 0.860$, close to pure Taoflow (0.879). Adding short selling at $\lambda = 0.20$ improves the hybrid to $\rho = 0.928$ ($\Delta = +0.068$, $p < 0.001$, $d = 3.8$). The hybrid admits enough price information for short-selling corrections to feed through while retaining enough smoothing for baseline stability.

To characterize the full frontier, we sweep w from 0.0 (pure Taoflow) to 1.0 (pure dTAO) in increments of 0.1, each with and without short selling ($\lambda = 0.20$, 50 replications per condition with paired seeds). The frontier is flat across a broad range: allocation accuracy with short selling exceeds $\rho = 0.926$ across $w \in [0.3, 0.8]$, with a nominal peak at $w = 0.4$ ($\rho = 0.929$). Beyond this plateau, accuracy declines toward pure dTAO ($\rho = 0.827$ at $w = 1.0$), where the absence of any smoothing admits too much noise. The baseline allocation (without short selling) declines monotonically with w , from $\rho = 0.878$ at $w = 0.0$ to $\rho = 0.662$ at $w = 1.0$. The optimal hybrid weight with short selling ($w^* \approx 0.4$) differs from the optimal weight without it ($w^* = 0.1$, $\rho = 0.914$), illustrating that the best mechanism design depends on whether informational market mechanisms are available.

A formal test of frontier flatness confirms the plateau is not noise. A one-way ANOVA of the six 50-replication samples at $w = 0.3, 0.4, 0.5, 0.6, 0.7, 0.8$ yields $F(5, 294) = 0.139$ with $p = 0.983$, failing to reject the null hypothesis that ρ is constant across these weights. The grand mean across the plateau is $\rho \approx 0.928$ with within-group standard deviations of approximately 0.015. This is recorded as Observation 1 in Appendix B, with a heuristic explanation: for $w \geq 0.3$ the price-proportional component receives sufficient weight to capture short-selling corrections almost completely, while for $w \leq 0.8$ the EWMA component retains enough weight to track persistent fundamental staking flows. In this interior range both channels contribute quality-tracking information, and the accuracy ceiling is determined by the information content of the combined signal rather than the specific mixing weight. Mechanism designers choosing a hybrid weight in $[0.3, 0.8]$ therefore have a wide margin of insensitivity.

5.5 The Lucas critique: adversarial defense

A natural objection is that project teams would not passively accept short-selling pressure. To test whether the 2×2 interaction survives adversarial adaptation, we introduce defensive staking: project teams counter sustained short flows with offsetting positive stakes, delayed by three epochs. Under moderate defense ($\delta = 0.5$), the dTAO allocation benefit shrinks from $\Delta\rho = +0.165$ to $+0.132$ ($\rho = 0.810$); under aggressive defense ($\delta = 1.0$, fully offsetting), it shrinks further to $+0.104$ ($\rho = 0.782$). The benefit is attenuated but the

sign is preserved across all defense intensities. Under Taoflow with moderate defense, the allocation damage persists ($\Delta\rho = -0.106$ vs. -0.115 without defense; $\rho = 0.773$).

The structural reason is that defensive staking is a low-frequency response — it requires observation, coordination, and execution over multiple epochs. It therefore operates in the frequency band that the EWMA passes, not the band it filters. Under dTAO, both the short-selling correction and the defensive response feed into prices immediately; the net correction is smaller but still positive because short sellers identify genuine overvaluation that defensive staking cannot eliminate. Under Taoflow, both the correction and the defense are integrated by the EWMA, and the net effect remains negative because the EWMA cannot distinguish informational corrections from defensive noise. The sign result is structurally robust to adversarial adaptation.

6. Distributional Effects of Short Selling

The aggregate 2×2 interaction masks distributional asymmetries across subnet size. We partition subnets into quintiles by TAO pool depth at the burn-in boundary and compare per-quintile outcomes between baseline and short-selling conditions across all four emission regimes (50 replications each).

Under Taoflow, short selling imposes disproportionate survival costs on small subnets: Q1 (smallest quintile) loses 17.6 percentage points of builder survival compared to 1.3 points for Q5 (largest). Paradoxically, Q1's emission share increases (+3.3 points), suggesting that the subnets that survive capture a larger share of the redistributed emissions. Under dTAO the pattern reverses: short selling is uniformly beneficial, with Q1 gaining 1.2 points of survival while allocation quality improves across all quintiles. The hybrid mechanism falls between, with Q1 gaining 4.7 points of survival.

Under filtered Taoflow, the distributional effects are near-neutral: emission-share changes are below 0.7 percentage points for all quintiles, and survival improves slightly across the board. The exclusion of short-selling flows from the EWMA eliminates the channel through which short selling degrades small-subnet emissions under unfiltered Taoflow.

The distributional asymmetry is compounded by a squeeze asymmetry in defensive staking costs. To fully offset short-selling pressure at moderate intensity ($\lambda = 0.10$), Q1 subnets would need to deploy defensive stakes equal to 0.020% of their pool per epoch, compared to 0.002% for Q5 — an order-of-magnitude difference. Large, well-capitalized subnets can absorb defensive costs as rounding error; small subnets face a meaningful capital constraint. This asymmetry supports the concern raised by practitioners that short selling under flow-based emission mechanisms disproportionately burdens independent, under-capitalized projects.

7. Robustness

The 2×2 interaction is tested across four primary dimensions (quality distributions, signal noise, EWMA half-life, slippage cap) and three extended dimensions (median filter, noise decomposition, zero quality-pool correlation) to assess whether it is an artifact of specific parameter choices.

7.1 Quality distributions

The interaction is tested under three quality distributions: rank-correlated with pool depth (baseline, $\gamma = 0.6$), Pareto (few excellent subnets, many mediocre — extreme quality concentration), and equal-with-noise (minimal true quality variation, $\sigma = 0.08$). The interaction holds for rank-correlated ($\Delta_{\text{dTAO}} = +0.078$, $\Delta_{\text{Taoflow}} = -0.094$) and equal-with-noise ($\Delta_{\text{dTAO}} = +0.206$, $\Delta_{\text{Taoflow}} = -0.142$). Under Pareto quality, short selling improves allocation under both mechanisms (+0.164, +0.200), breaking the interaction.

The Pareto exception has a clear mechanism: when quality differences are extreme, the corrective flows from short selling are large enough to overcome the EWMA filter. The informational signal's magnitude exceeds the filter's attenuation capacity. This identifies a boundary condition — the emission rule matters when quality differences are moderate, which is the empirically relevant case given that Maymin (2026b) documents imperfect but positive price-quality correlation.

7.2 Short-seller signal noise

The baseline specification assumes short sellers observe the true quality ranking — an obviously unrealistic assumption that we relax by adding Gaussian noise to the quality signal ($\sigma \in \{0, 0.1, 0.2, 0.3, 0.5, 1.0\}$, where quality lies in $[0, 1]$). The interaction holds at all six noise levels. Under dTAO, the allocation benefit shrinks gracefully from $\Delta\rho = +0.086$ (perfect signal) to +0.038 (nearly random shorting); the AMM's self-correcting mechanics limit the damage from incorrect shorts because wrong positions lose money and are reversed by the market. Under Taoflow, noisy short selling is catastrophic: allocation accuracy becomes negatively correlated with quality ($\rho < 0$) at $\sigma \geq 0.3$, meaning emissions are allocated inversely to quality. The EWMA integrates all flows — correct and incorrect — with equal weight; it has no mechanism to distinguish informational corrections from erroneous noise. This asymmetry sharpens the central finding: dTAO degrades gracefully under noisy short selling; Taoflow's allocation collapses catastrophically under the same conditions.

7.3 EWMA half-life

Half-lives of 3, 7, 14, 30, 60, and 90 days at moderate intensity ($\lambda = 0.10$) confirm the filtering prediction. The allocation damage from short selling scales monotonically: -0.022 at 3 days, -0.095 at 90 days. The sign never flips: any EWMA introduces lag between corrective flows and emissions. The discontinuity is between any smoothing (damage) and no smoothing (dTAO, benefit). Baseline allocation improves monotonically with half-life (0.619 to 0.879), quantifying the robustness-improvability frontier. Very short smoothing (3 days) is worst of both worlds: Taoflow baseline (0.619) falls below the sweep's dTAO baseline (0.657).

7.4 Slippage cap

The AMM slippage cap varied from 1% to 10% of pool TAO per epoch leaves results stable: dTAO benefit ranges from +0.096 to +0.127, Taoflow damage from -0.065 to -0.076 . Across the four robustness dimensions in §§7.1–7.4, 14 of 15 parameter configurations confirm the interaction; the Pareto-quality exception reflects the boundary condition noted in §7.1.

7.5 Non-linear filtering: median vs. EWMA

A rolling 30-day median filter (50 replications, paired seeds) produces a lower baseline than the EWMA ($\rho = 0.807$ vs. 0.878) because the median discards magnitude information. Short selling still degrades allocation under the median filter ($\Delta\rho = -0.077$), 29% smaller than the EWMA damage ($\Delta\rho = -0.109$). The median's robustness to outliers provides partial protection, but the sign is unchanged: both linear (EWMA) and non-linear (median) smoothing block the informational content of short-selling flows. Damage is a property of smoothing per se, not of exponential weighting.

7.6 Persistent/transient noise ratio

The 2×2 interaction could depend on the noise decomposition. Sweeping the persistent/transient split across $(0.10/0.45)$, $(0.275/0.275)$, $(0.35/0.20)$, and $(0.45/0.10)$ — holding total noise weight at 0.55, 50 replications each — confirms the sign-flip in all four configurations: dTAO benefit ranges from $+0.169$ to $+0.175$, Taoflow damage from -0.104 to -0.116 .

7.7 Independence from the size premium and emission magnitude

To address the concern that the sign-flip might be an artifact of the quality-pool correlation ($\gamma = 0.6$) that drives the simulation's incorrect negative SMB (§3.6), we test the 2×2 interaction at $\gamma = 0$ (50 replications, paired seeds). Under dTAO with zero quality-pool correlation, baseline allocation is near zero ($\rho = 0.020$, as prices carry no quality signal) but short selling produces $\Delta\rho = +0.468$. Under Taoflow, baseline allocation remains strong ($\rho = 0.871$) and short selling degrades it ($\Delta\rho = -0.127$). The sign-flip is confirmed, independent of the SMB channel. The interaction also holds at the pre-halving emission rate of 7,200 TAO/day (dTAO $\Delta\rho = +0.139$, Taoflow $\Delta\rho = -0.041$, $N = 10$), confirming independence from emission magnitude.

8. The Grossman–Stiglitz Inversion

The Bitcoin-mining analogy of §1.1 observed that Bittensor replaces Bitcoin's physics-based allocation with a market-based one, inheriting in the process the information-aggregation problem that Grossman and Stiglitz (1980) identified. This section cashes out that inheritance. We compute expected net profit per short-selling trade at each intensity λ , tracking the gross price correction captured minus the AMM slippage cost (Maymin 2026a, Proposition 1). Under both mechanisms, the AMM slippage cost exceeds the gross correction captured per trade, making short selling unprofitable in net terms at all tested intensities. Under dTAO, the gross correction per trade is positive but modest ($+0.004\%$ to $+0.036\%$ of price), while slippage grows quadratically with intensity ($+0.006\%$ to $+0.083\%$). Under Taoflow, the gross correction grows faster ($+0.002\%$ to $+0.045\%$) because mispricings persist without emission-driven rebalancing; but slippage still dominates.

The divergence between socially beneficial and privately profitable short selling is the core inversion. Under dTAO, short selling is socially valuable ($\Delta\rho = +0.165$) but yields the smallest gross corrections because the emission mechanism rapidly closes mispricings, so the private reward is small. Under Taoflow, short selling is socially harmful ($\Delta\rho = -0.115$) but yields larger gross corrections because mispricings

persist, so the private reward is larger. A rational trader facing the two regimes would prefer to short under Taoflow — precisely the regime where short selling damages allocation. Theorem 2 in Appendix B formalizes this by showing $\lambda^*_{\text{Taoflow}} > \lambda^*_{\text{dTAO}}$. Achieving the social benefit under dTAO requires subsidy to compensate for private unprofitability; Taoflow requires taxation or restriction to prevent allocation damage.

8.1 Relation to information-aggregation theories

The sign-flip connects to a classical question: when does price informativeness translate into efficient resource allocation? In the Glosten–Milgrom (1985) sequential trade model, a market maker updates prices after each trade based on conditional expectations given order flow. The AMM is a mechanical analog: prices update deterministically via the constant-product formula rather than through Bayesian updating. The key distinction is not in price formation — both mechanisms produce informative prices when informed trading occurs — but in how downstream allocation uses the revealed information.

The result extends the Grossman–Stiglitz (1980) concern about the capture of informational rents to the mechanism-design level. Under dTAO, the emission mechanism captures the informational value that short sellers produce (closing the mispricing as soon as it is identified), so short sellers realize less than the social benefit; Theorem 2 formalizes this by showing the gross correction under dTAO saturates. Under Taoflow, the EWMA filter blocks the capture channel — short sellers retain larger per-trade returns, but the informational value never reaches the allocation process.

The policy implication adds a second necessary condition to the classical subsidy justification: subsidies to informed trading are productive only if the allocation mechanism can translate price informativeness into allocative efficiency. Under dTAO both conditions hold; under Taoflow, only the first holds. Subsidizing short selling under Taoflow would thus satisfy the classical Grossman–Stiglitz justification while failing the mechanism-design test.

9. Discussion

9.1 Design recommendations

For a mechanism designer choosing a staking-based emission rule, the results yield a sequential decision procedure.

Step 1. Is short selling available as a market mechanism?

- *No.* Prefer flow-smoothed allocation (Taoflow); baseline $\rho = 0.879$ exceeds dTAO's 0.678. Do not expand capacity prematurely — doubling subnets from 128 to 256 reduces ρ to 0.830 and doubles overvaluation. Avoid hard emission cutoffs ($\Delta\rho = -0.139$ at $K = 64$).
- *Yes.* Proceed to Step 2.

Step 2. Can the protocol distinguish short-selling flows from organic unstaking?

- *No.* Use price-proportional allocation (dTao). Short selling produces $\Delta\rho = +0.165$ and the mechanism captures the informational value even without filtering. Subsidize informed

participation (Theorem 2) to overcome the Grossman–Stiglitz gap between private and social returns.

- *Yes* (a dedicated lending contract exists). Proceed to Step 3.

Step 3. Choose between filtered Taoflow and the hybrid.

- *Filtered Taoflow*. Exclude short flows from the EWMA; preserves baseline accuracy ($\rho = 0.883$) and captures the price-quality improvement. Requires the separating-equilibrium conditions of Theorem 4: a correction subsidy s exceeding the lending-pool borrow rate r and positive post-subsidy profit for informed short sellers.
- *Hybrid* ($w \in [0.3, 0.8]$). Weight price-proportional and flow-smoothed allocation; achieves $\rho \approx 0.928$ with short selling (F-test fails to reject flatness across this range, $p = 0.983$). The wide plateau gives mechanism designers a large margin of insensitivity.

The general principle: signal smoothing and market-based information correction are substitutes on the allocation frontier. Adding a market mechanism to a smoothed system can be counterproductive unless the smoothing is designed to exclude the market's corrective flows.

9.2 Robustness and generality

The 2×2 interaction is internally robust within the calibrated model through three channels, which together bound the parameter region in which the sign-flip is observable. First, it holds across 22 parameter configurations (15 original + 7 filtered Taoflow), 58 of 58 validation-passing parameterizations drawn via Latin hypercube at 20 replications per configuration (100%, 95% lower bound 94.97%), the four persistent/transient noise-weight splits, and zero quality-pool correlation (§7.7). Second, it holds under both linear (EWMA) and non-linear (median) filters (§7.5), confirming that the mechanism is a property of signal smoothing per se, not a specific filter design. Third, it survives adversarial defensive staking (§5.5), the most direct form of strategic adaptation we test. External generality, as distinct from internal parameter robustness, is examined below through an illustrative consistency check against Curve Finance, not claimed as a direct empirical result.

The result does not depend on the specific quality concept. The simulation studies the emission mechanism's capacity to transmit quality information into allocation decisions, regardless of what that quality represents. The 2×2 interaction is about signal transmission, not signal content. The robustness across three quality distributions (rank-correlated, Pareto, equal-with-noise) confirms this: the interaction holds for the mechanism, not for a particular quality definition.

The Lucas critique applies: agents would adapt to the interventions we impose. Under short selling, momentum traders might adapt to corrective pressure; under cutoffs, evaluators might invest more effort near the boundary. The directional result survives because the EWMA filters high-frequency content regardless of agent behavior. Magnitudes might change under strategic adaptation; signs should not. This structural argument is supported by the analytical results in Appendix B (Theorems 1, 2, and 4) and by the simulation's robustness to defensive staking.

Illustrative external consistency check: Curve Finance gauge voting

The claim that the 2×2 interaction is a property of a broader class of smoothed allocation mechanisms can be tested informally against a second real-world system. Curve Finance, the largest stable-asset automated market maker on Ethereum, allocates its CRV emissions across liquidity pools through a weekly gauge-voting system in which voting power is vested in veCRV — CRV locked for between one week and four years, with voting weight decaying linearly over the lock. Lockups make voting power persistent across weeks, so gauge weights respond to quality signals only with substantial delay. Curve's smoother is a lockup-decay function rather than an EWMA, and its external corrective flow is a political vote-buying market (Votium, Hidden Hand) rather than price-based short selling; the mechanism-class analogy is therefore approximate rather than parameter-matched. We use the case as an illustrative consistency check, not as a second calibrated empirical study.

On 2026-04-15 we collected a cross-sectional snapshot of the Curve Ethereum gauge system from three free public endpoints: pool registries, subgraph aggregates, and the gauge controller. Joining pool TVL to gauge relative weight gives 259 active Ethereum pools with nonzero gauge weight, aggregating approximately 94% of global CRV emissions and roughly \$1 billion of pooled TVL. We compute the Spearman rank correlation $\rho_C = \text{Spearman}(w_i, \text{TVL}_i)$ as an informal Curve analog of the paper's allocation-accuracy metric.

Across the full 259-pool active universe the overall correlation is $\rho_C \approx 0.65$. Because several legacy pools created before the modern gauge and bribe market took shape — most prominently the DAI/USDC/USDT 3Pool, the ETH/stETH pool, and the renBTC/wBTC/sBTC tri-Bitcoin pool — carry TVL ranks near the top of the distribution but receive negligible emissions, we apply a pool-lifecycle control by restricting the sample to pools created on or after 2022-01-01. The controlled subset contains 235 pools aggregating approximately \$743 million of TVL. Within this subset the overall Spearman correlation rises to $\rho_C \approx 0.71$, while the top-20 Spearman is approximately 0.20. The compression from roughly 0.71 at the aggregate to roughly 0.20 at the top survives four alternative cutoffs we test (post-2021-01, post-2022-01, post-2022-06, post-2023-01). A hypothetical Rule-R analog — emission proportional to TVL — would by construction have $\rho_C = 1$ throughout; the observed gap between aggregate and top-of-distribution correlation is the type of pattern a smoothed allocation mechanism would produce.

At least three causal stories are consistent with this compression, and the cross-sectional data cannot adjudicate between them: (i) a Rule-S-style smoothing attenuation in which the veCRV lockup decay filters quality signals slowly; (ii) governance capture by coordinated vote-buying markets in which vlCVX holders direct emissions toward pools whose sponsors pay them most; and (iii) a residual pool-lifecycle effect in which newer factory pools attract more bribe activity than slightly older pools. The lifecycle control addresses the sharpest form of story (iii); stories (i) and (ii) remain confounded. Adjudicating between them would require longitudinal gauge-weight panel data aligned with per-epoch bribe flows, which is not retrievable through free public endpoints. We therefore offer the result as an illustrative consistency check on the mechanism class that Theorem 1 addresses rather than as a causal identification of Rule-S attenuation in Curve.

With that caveat, the compression does place the Curve mechanism within the class of outcomes the paper's 2×2 interaction addresses: a smoothed allocation rule in which the quality signal is transmitted imperfectly and fails to track the most economically significant pools precisely. The pattern is observable, at least in shape, in an economically significant DeFi protocol whose design predates and is logically independent of Bittensor's emission rules.

9.3 Limitations and open questions

Quality is a scalar latent variable; real quality is multidimensional and partially subjective. The 2×2 interaction breaks under extreme quality dispersion (Pareto distribution): when quality differences are large enough, short-selling corrections exceed the EWMA's attenuation capacity even under the smoothed rule. This is a genuine boundary condition, not a parameter artifact, and it limits the result's scope to settings with moderate quality heterogeneity.

The size premium is not reproduced: the simulation yields negative SMB ($-0.14\%/day$) against an empirical target of $+1.01\%/day$. As explained in §3.6, this gap is a predictable consequence of three calibrated modeling choices — quality-pool correlation ($\gamma = 0.6$), herding ($w_{herd} = 0.15$), and Taoflow EWMA smoothing — each of which suppresses the emission amplification channel that would otherwise favor small pools. The empirical premium was measured during the dTAO era (February–November 2025), when rapid subnet entry ($63 \rightarrow 124$ subnets) and the absence of Taoflow smoothing left the amplification channel intact; a stationary Taoflow simulation excludes these regime dynamics by design. Crucially, the 2×2 interaction finding is independent of the SMB sign (§7.7 shows the sign-flip holds at zero quality-pool correlation), so the cross-sectional miscalibration limits predictive scope for distributional outcomes but does not affect the paper's primary inference. The flow-decomposition weights are chosen to match qualitative regularities rather than estimated from data; the non-stationary empirical period (two emission regime changes, network growth from 63 to 128 subnets) makes time-series estimation unreliable.

The Grossman–Stiglitz equilibrium is characterized formally in the two-asset model (Theorem 2, Appendix B) and directionally in the simulation ($\lambda^*_{Taoflow} > \lambda^*_{dTAO}$), but not solved as a full dynamic equilibrium. Solving the dynamic equilibrium — incorporating information-acquisition costs, capital opportunity costs, rational expectations about future mispricing, and endogenous pool-depth feedback — is the natural next paper in this research program and is substantially more ambitious than the current analysis. The contribution of the present paper is to establish the analytical backbone (sign-flip in a two-asset model; convex/concave subsidy-intensity frontier; separating equilibrium under filtered Taoflow) and the empirical regularities (100-parameterization robustness, hybrid plateau, filtered mechanism efficacy) that a dynamic equilibrium model would need to reproduce and extend. Positioning the two as sequential rather than substitute contributions matters: Theorem 2 already delivers the subsidy-intensity frontier in closed form for both rules, and Theorem 4 establishes existence of a separating equilibrium under filtered Taoflow — these are rational-expectations-consistent results within the analytical scope of the two-asset model.

The filtered Taoflow mechanism requires a dedicated lending infrastructure for on-chain implementability; without it, short-selling flows cannot be reliably separated from organic unstaking. Theorem 4 establishes a separating equilibrium under two conditions (subsidy s exceeds borrow rate r ; post-subsidy short selling

is profitable) but makes three simplifying assumptions worth noting. First, Type U agents are modeled as pure liquidity unstakers without speculative motive; sophisticated Type U agents with risk aversion and own-subnet exposure could route through the lending pool to hedge, degrading the separation. Second, the protocol subsidy s is assumed funded from the emission budget without explicit budget-balance constraints; in periods of intense short-selling activity, the implied subsidy payments could exceed the allocated fraction, requiring governance intervention. Third, Theorem 4 addresses a single short seller; multi-agent competition for subsidies could destabilize the equilibrium. Short-and-distort attacks are feasible in thin AMM pools and remain untested in the current simulation. The 50-replication design provides moderate statistical power; four of 24 primary tests do not survive Bonferroni correction. Short-selling mechanics omit borrowing rates, oracle settlement, and liquidation cascades.

10. Conclusion

We identify a fundamental design trade-off in staking-based resource allocation: signal smoothing and market-based information correction are substitutes, not complements. The EWMA filter that makes Taoflow's baseline allocation superior ($\rho = 0.879$ vs. dTAO's 0.678) simultaneously blocks the informational corrections that short selling provides, producing opposite-signed effects on allocation accuracy under the two mechanisms (dTAO: $\Delta\rho = +0.165$; Taoflow: $\Delta\rho = -0.115$).

Theorem 1 (Appendix B) establishes this sign-flip as a structural property of any smoothed allocation rule in the two-asset model, proving that a positive-measure parameter region exists where raw prices improve allocation while smoothed prices degrade it. Theorem 2 characterizes the subsidy-intensity frontier, proving that the gross correction under dTAO is concave in short-selling intensity while under Taoflow it is convex, with a unique closed-form crossover at $\lambda \approx 0.15$ matching the simulation's equilibrium data. Proposition 3 establishes empirical genericity: of 100 parameterizations sampled via Latin hypercube, 58 pass a tightened validation screen, and all 58 (100%) exhibit the full sign-flip at 20 replications per configuration (95% one-sided binomial lower bound: 94.97%). Theorem 4 proves the existence of a separating equilibrium under filtered Taoflow when the protocol offers a correction subsidy exceeding the lending-pool borrow rate and short selling is profitable after costs. Together, these results turn the empirical 2×2 observation into a structural claim about smoothed allocation mechanisms in general.

The 2×2 interaction is structurally robust across 22 parameter configurations, 58 of 58 validation-passing Latin-hypercube parameterizations, four noise-decomposition ratios, two filter types (EWMA and median), and zero quality-pool correlation. Paired-difference tests and bootstrap confidence intervals confirm all three core effects at $p < 10^{-6}$ with wide CI margins from zero.

The prescription for mechanism designers is threefold. First, segregate informational flows from the smoothing calculation when both smoothing and market-based correction are desired: excluding short-selling flows from the EWMA preserves baseline accuracy ($\rho = 0.883$) while capturing the price-quality improvement. Second, the hybrid frontier between price-proportional and flow-smoothed allocation has a concave interior optimum ($\rho \approx 0.928$ across a broad plateau $w \in [0.3, 0.8]$ with short selling) that depends on whether informational market mechanisms are available. Third, capacity expansion before resolving the allocation mechanism is premature, and hard cutoffs are counterproductive.

The result generalizes beyond Bittensor. Any system where staking, voting, or delegation produces price signals that are smoothed before allocation — a pattern common to proof-of-stake networks (Saleh 2021), prediction markets with liquidity subsidies, and decentralized grant programs — faces the same trade-off. An illustrative external consistency check against Curve Finance's veCRV gauge-voting system (§9.2) is qualitatively compatible with this account: in a cross-sectional snapshot of 259 active Curve gauges on Ethereum aggregating approximately \$1 billion of liquidity, gauge weights exhibit moderate tracking of pool TVL at the aggregate level ($\rho_C \approx 0.65$, rising to ≈ 0.71 under a pool-lifecycle control) but only weak tracking within the top twenty pools by TVL ($\rho_C \approx 0.20$ under the same control), a compression of the type a smoothed allocation rule would produce. The cross-sectional design precludes causal identification and the Curve mechanism differs structurally from Bittensor's in several respects, so the check is offered as a consistency observation rather than a second empirical confirmation.

The EWMA half-life parameterizes the robustness-improvability frontier continuously; the discontinuity between any positive half-life and zero marks the qualitative shift from smoothing regime to market regime. Open questions, detailed in §9.3, include the dynamic Grossman–Stiglitz equilibrium, multidimensional quality, and calibration to stationary on-chain data as the Taoflow-era sample matures.

Acknowledgments

The authors are grateful to Philip Maymin, whose theoretical and empirical work on AMM-priced token markets (Maymin 2026a, 2026b) provides the foundation on which this paper builds, and whose insight that short-selling flow could be excluded from the Taoflow EWMA signal motivated the filtered Taoflow mechanism analyzed in §9.1 and Appendix B.6. All remaining errors are the authors' own.

Disclosure of Interest

Gershteyn owns and operates Bittensor Subnet 103 (Djinn), one of the 128 subnets whose collective emissions this paper analyzes. This creates a direct financial interest in the policy proposals evaluated here: the choice of emission mechanism (dTAO, Taoflow, filtered Taoflow, hybrid) and the short-selling, cutoff, and capacity-expansion interventions all affect Gershteyn's subnet. The analysis is conducted symmetrically across all 128 subnets; Subnet 103 receives no special treatment in any simulation condition. The theoretical results (Theorems 1, 2, and 4) do not depend on subnet-specific data. Random seeds are deterministic and documented in Appendix A, enabling independent replication of every numerical claim. The policy recommendations (filtered Taoflow, hybrid emission with $w \in [0.3, 0.8]$) are content-neutral with respect to which subnet benefits and do not advantage Subnet 103 over any other subnet. Zevelev declares no competing financial interests in Bittensor or any of the mechanisms or subnets analyzed.

References

- Anderson, C. (2006). *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion.
- Angeris, G., Kao, H.-T., Chiang, R., Noyes, C., & Chitra, T. (2021). An analysis of Uniswap markets. *Cryptoeconomic Systems*, 1(1).
- Budish, E., Cramton, P., & Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4), 1547–1621.
- Capponi, A., & Jia, R. (2025). Liquidity provision on blockchain-based decentralized exchanges. *Review of Financial Studies*, 38(10), 3040–3085.
- Chen, Y., Da, Z., & Huang, D. (2022). Short selling efficiency. *Journal of Financial Economics*, 145(2), 387–408.
- Cong, L. W., Li, Y., & Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *Review of Financial Studies*, 34(3), 1105–1155.
- Cong, L. W., & He, Z. (2024). The tokenomics of staking. NBER Working Paper No. 33640.
- Diamond, D. W., & Verrecchia, R. E. (1987). Constraints on short-selling and asset price adjustment to private information. *Journal of Financial Economics*, 18(2), 277–311.
- Elberse, A. (2008). Should you invest in the long tail? *Harvard Business Review*, 86(7/8), 88–96.
- Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100.
- Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3), 393–408.
- Han, J., Lee, J., & Li, T. (2023). DAO governance. NUS Asian Institute of Digital Finance Working Paper No. 22-022723.
- Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1), 65–91.
- Karpoff, J. M., & Lou, X. (2010). Short sellers and financial misconduct. *Journal of Finance*, 65(5), 1879–1913.
- Lamperti, F., Roventini, A., & Sani, A. (2018). Agent-based model calibration using machine learning surrogates. *Journal of Economic Dynamics and Control*, 90, 366–389.
- LeBaron, B. (2006). Agent-based computational finance. In Tesfatsion & Judd (Eds.), *Handbook of Computational Economics*, Vol. 2. Elsevier.
- Lehar, A., & Parlour, C. A. (2025). Decentralized exchange: The Uniswap automated market maker. *Journal of Finance*, 80(1), 321–374.
- Liu, Y., & Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, 34(6), 2689–2727.
- Liu, Y., Tsyvinski, A., & Wu, X. (2022). Common risk factors in cryptocurrency. *Journal of Finance*, 77(2), 1133–1177.
- Lucas, R. E. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.
- Lui, E., & Sun, J. (2025). Bittensor protocol: The Bitcoin in decentralized artificial intelligence? A critical and empirical analysis. arXiv:2507.02951.
- Makarov, I., & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2), 293–319.
- Maymin, P. Z. (2026a). Option pricing on automated market maker tokens. arXiv:2603.29763.
- Maymin, P. Z. (2026b). Common risk factors in decentralized AI subnets. arXiv:2603.29751.
- Meyn, S. P., & Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability*. Springer-Verlag.
- Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *Journal of Finance*, 32(4), 1151–1168.

- OECD (2024). The concentration of DeFi's liquidity. OECD Working Papers on Finance, Insurance and Private Pensions.
- Pagnotta, E. S. (2022). Decentralizing money: Bitcoin prices and blockchain security. *Review of Financial Studies*, 35(2), 866–907.
- Saffi, P. A. C., & Sigurdsson, K. (2011). Price efficiency and short selling. *Review of Financial Studies*, 24(3), 821–852.
- Saleh, F. (2021). Blockchain without waste: Proof-of-stake. *Review of Financial Studies*, 34(3), 1156–1190.
- Scullen, S. E., Bergey, P. K., & Aiman-Smith, L. (2005). Forced distribution rating systems and the improvement of workforce potential. *Personnel Psychology*, 58(1), 1–32.
- Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins University Press.
- Sockin, M., & Xiong, W. (2023). Decentralization through tokenization. *Journal of Finance*, 78(1), 247–299.
- Welch, J. (2001). *Jack: Straight from the Gut*. Warner Books.

Appendix A: Full Simulation Specification

Network. 128 subnets (expansion to 256 at epoch 100 under Intervention C). Daily emissions 3,600 TAO.

AMM pools. $k = \tau\alpha$, preserved during trades. $\tau \sim \text{LogNormal}(\mu = 7.0, \sigma = 1.8)$; $\alpha \sim \text{LogNormal}(\ln(10^6), 0.5)$, semi-independent of τ . Prices $p = \tau/\alpha$.

Quality. Rank-correlated with pool depth ($\gamma = 0.6$) plus uniform noise (weight 0.4). Normalized to $[0, 1]$. Sensitivity: also tested with Pareto (shape 1.5) and equal-with-noise ($\sigma = 0.08$) distributions.

Flow decomposition. $F = 0.20 \cdot F_{\text{fund}} + 0.10 \cdot F_{\text{mom}} + 0.15 \cdot F_{\text{herd}} + 0.55 \cdot F_{\text{noise}}$. Base scale = $\sqrt{\tau} \cdot 0.032$. Noise: Student-t($\nu = 3$). Momentum autocorrelation 0.3.

Emission mechanisms. dTAO: emission share \square price. Taoflow: emission share \square $\max(\text{EWMA}(\text{flow}, \text{halflife} = 30\text{d}), 0)$. Of emissions: 30% pool injection (price-neutral), 18% restaked via AMM (price-positive), 1% sold (price-negative).

Builders. Operating cost 1.5 TAO/day. Exit after 30 days if average profit < -0.5 TAO/day. Entry probability 0.02 per empty slot per day. Cutoff deterrence: fundamental flow reduced within 10 ranks of cutoff boundary.

Short selling. $F_{\text{short}} = -\lambda \cdot (\text{price_rank} - \text{quality_rank}) \cdot \sqrt{\tau}$, capped at 3% of τ per epoch. Sensitivity: $\lambda \square \{0.02, 0.05, 0.10, 0.20\}$; signal noise $\sigma \square \{0, 0.1, 0.2, 0.3, 0.5, 1.0\}$; slippage cap $\square \{1\%, 3\%, 5\%, 10\%\}$.

Statistics. 300 epochs (150 burn-in + 150 measurement), 50 replications per condition. Paired-difference tests use paired seeds: for each rep $r \square \{0, \dots, 49\}$, all conditions in a comparison share the same base seed $200,000 + r$, ensuring identical random draws for the quality distribution, pool depth, and flow innovations. This allows paired-difference inference in §5.1. Percentile bootstrap 95% CIs use 10,000 resamples. For historical comparison we also report Welch t-tests with Bonferroni correction at $\alpha = 0.05/24 = 0.0021$ for 24 primary tests.

Validation. Four of four empirical targets pass: return spread 0.46%/day, top-10 concentration 57%, price-quality correlation 0.86, excess kurtosis 8.6. SMB = -0.14% /day (documented limitation). Convergence: $p = 0.391$ between adjacent measurement windows.

Appendix B: Formal Theoretical Results

This appendix establishes four formal results: (B.2) the sign-flip theorem proving that under a smoothed allocation rule, there exists a positive-measure parameter region where short-selling corrections degrade allocation accuracy; (B.3) the subsidy-intensity frontier characterizing the equilibrium break-even intensity under both mechanisms; (B.4) a genericity result linking the two-asset theorem to the N-asset simulation sensitivity analysis; (B.5) corollaries on hybrid monotonicity and the flat plateau; and (B.6) a separating equilibrium theorem establishing the incentive-compatibility of filtered Taoflow. Together these results turn the empirical sign-flip observation into a structural claim about smoothed allocation mechanisms in general.

B.1 Setup

Consider a two-asset economy with qualities $q_1 > q_2 > 0$ and pool depths $\tau_1, \tau_2 > 0$. Each asset has a constant-product AMM with invariant $k_i = \tau_i \cdot \alpha_i$ and price $p_i = \tau_i/\alpha_i$. The fundamental flow into asset i at epoch t is $F_i > 0$, representing organic staking by uninformed participants. The flow is persistent with decay rate $\gamma \in [0,1]$: $F_i(t) = \gamma F_i(t-1) + (1-\gamma)\bar{F}_i$, where \bar{F}_i is the long-run mean, proportional to q_i .

A short seller at intensity $\lambda \geq 0$ introduces a correction flow $c_i(t)$ whenever asset i is overvalued. Specifically, if $q_i/\Sigma q < p_i/\Sigma p$ (asset i has a smaller quality share than price share, i.e., is overvalued), then $c_i(t) = -\lambda \cdot (p_i/\Sigma p - q_i/\Sigma q) \cdot \sqrt{\tau_i}$. The $\sqrt{\tau_i}$ scaling follows from the CEV structure of Maymin (2026a, Theorem 1). The correction is applied for a transient duration $T \in (1-\alpha)^{-1}$, after which it vanishes. The total flow into asset i is $f_i(t) = F_i(t) + c_i(t)$.

The AMM processes flows as follows. A flow $f_i(t) > 0$ stakes TAO, reducing alpha reserves and raising the price; a flow $f_i(t) < 0$ does the reverse. For small flows relative to pool depth, the linearized price response is $\Delta p_i \approx 2p_i(f_i/\tau_i)$, as derived from the constant-product AMM first-order condition (Maymin 2026a, Proposition 1). The deterministic slippage cost per trade of size $|f_i|$ is $S(|f_i|) = |f_i|^2/k_i$ in TAO-equivalent terms.

We compare two allocation rules. Rule R (raw price): the emission share of asset i is $e_i(R) = p_i/(p_1 + p_2)$. Rule S (smoothed): the emission share is $e_i(S) = \max(s_i, 0)/\max(\Sigma s_i, \epsilon)$, where s_i is an exponentially weighted moving average of flows with decay $\alpha \in (0,1)$: $s_i(t) = \alpha \cdot s_i(t-1) + (1-\alpha) \cdot f_i(t)$, and $\epsilon > 0$ is a small regularization constant preventing division by zero. A hybrid rule B(w) blends the two: $e_i(B(w)) = w \cdot e_i(R) + (1-w) \cdot e_i(S)$ for $w \in [0, 1]$.

Allocation accuracy is measured by the correlation between emission shares and quality shares. We work with a simpler scalar proxy $\rho = e_1 - e_1^*$ where $e_1^* = q_1/(q_1 + q_2)$ is the target emission share for the high-quality asset. $\rho = 0$ is perfect allocation, $\rho < 0$ is underallocation to quality, $\rho > 0$ is overallocation. Under our setup with $q_1 > q_2$ but $p_1 < p_2$ (overvaluation of the low-quality asset), the baseline has $\rho(t=0) < 0$ under Rule R. The correction moves prices toward quality alignment, so a successful correction should drive ρ toward 0 from below.

B.2 Theorem 1: Formal sign-flip in the two-asset model

Theorem 1 (Sign-flip). *Let q_1, q_2, τ_1, τ_2 satisfy $q_1 > q_2$ and $p_1 < p_2$ (the low-quality asset is overvalued). Fix any smoothing parameter $\alpha \in (0, 1)$, any persistent fundamental flow $\bar{F}_2 > 0$ to the overvalued asset, any correction intensity $\lambda > 0$, and any correction duration $T \geq 1$. Let ρ_0 denote the baseline allocation gap and let $\rho^{\text{post}_R}, \rho^{\text{post}_S}$ denote allocation accuracy after the correction under Rules R and S respectively. Then:*

(a) *Rule R strictly improves allocation: $\rho^{\text{post}_R} - \rho_0 > 0$.*

(b) *There exists a non-trivial open set $\Omega \subset \mathbb{R}^8$ in the parameter space $(q_1, q_2, \tau_1, \tau_2, \bar{F}_1, \bar{F}_2, \lambda, T)$, defined by the condition*

$$T \cdot (1-\alpha) \cdot \lambda < \bar{F}_2 \cdot (1-\alpha^T) \cdot (1-\rho_0)^{-1}, \quad (\square)$$

where $\rho_0 = q_1/(q_1 + q_2) - p_1(0)/(p_1(0) + p_2(0))$ is the baseline allocation gap determined by the primitive parameters via the initial prices $p_i(0) = \tau_i/\alpha_i$, such that Rule S strictly degrades allocation: $\rho^{\text{post}_S} - \rho_0 < 0$. In particular, for any fixed $(q, \tau, \bar{F}, \lambda, T)$ with $\rho_0 < 0$, there exists $\tilde{\alpha} \in (0, 1)$ such that for all $\alpha \in (\tilde{\alpha}, 1)$, the sign-flip holds.

Proof. Part (a) is direct. Under Rule R, $e_1(R) = p_1/(p_1 + p_2)$. Differentiating with respect to p_1 holding $p_2 + p_1$ approximately constant gives $\partial e_1/\partial p_1 = p_2/(p_1 + p_2)^2 > 0$. The correction flow $c_2 < 0$ induces $\Delta p_1 > 0$ and $\Delta p_2 < 0$ via the AMM's linearized price response. Therefore $\Delta e_1(R) > 0$, and since $\rho_0 < 0$ by assumption, $\rho^{\text{post}_R} = \rho_0 + \Delta e_1(R)$ is strictly larger (closer to 0), establishing strict improvement.

Part (b) is the substantive claim. Under Rule S the emission share depends on the EWMA of flows, not the price. Decompose $s_2(t)$ into its fundamental component s^{F_2} and its correction-contaminated component s^{C_2} :

$$s_2(t) = \alpha \cdot s_2(t-1) + (1-\alpha) \cdot (F_2(t) + c_2(t)) = s^{\text{F}_2}(t) - s^{\text{C}_2}(t),$$

where $s^{\text{F}_2}(t)$ evolves under the fundamental flow alone and $s^{\text{C}_2}(t) \geq 0$ is the cumulative EWMA contribution of the correction pulse. Because the correction is a negative pulse of magnitude $|c_2| = \lambda \cdot (\text{overvaluation}) \cdot \sqrt{\tau_2}$ applied for T periods, its contribution to s_2 after the pulse ends is

$$s^{\text{C}_2}(T) = (1-\alpha) \cdot |c_2| \cdot \sum_{k=0}^{T-1} \alpha^k = |c_2| \cdot (1-\alpha^T).$$

Meanwhile the fundamental component continues at its long-run average; in the limit $s^{\text{F}_2} \rightarrow \bar{F}_2$. So after the pulse $s_2(T) = s^{\text{F}_2}(T) - |c_2| \cdot (1-\alpha^T)$. The reduction in s_2 is proportional to the allocation drop for asset 2. Condition (\square) ensures that this reduction is smaller than what is required to push ρ^{post_S} above ρ^{pre_S} : the EWMA's attenuation factor $(1-\alpha^T)$ is bounded above by $T \cdot (1-\alpha)$ for small pulse durations (by the Taylor expansion of $1 - \alpha^T$), while the corruption of the flow signal scales linearly with $|c_2|$. The net effect on $e_1(S) = s_1/(s_1 + s_2)$ depends on whether s_1 increases relative to s_2 .

Rule R responds to the price change, which is proportional to c_2/τ_2 and not attenuated by the EWMA. Rule S responds to the flow-signal change, which is attenuated by $(1-\alpha^T)$. For α close to 1 and $T \leq (1-\alpha)^{-1}$, the expansion $1 - \alpha^T \approx T(1-\alpha)$ implies the Rule-S change is smaller than the Rule-R change by a factor of $(1-\alpha)$. Meanwhile, the corruption of s_2 's tracking of \bar{F}_2 accumulates at rate $(1-\alpha)|c_2|$ per period. Condition (\square) ensures the parameter region where this corruption dominates the attenuated correction benefit is non-empty and open.

To verify that Ω has positive Lebesgue measure we exhibit an interior point. Fix $\alpha = 0.95$, $T = 1$, $\lambda = 0.05$, $\bar{F}_2 = 1$, and quality-pool parameters $q_1 = 0.8$, $q_2 = 0.4$, $\tau_1 = 1$, $\tau_2 = 4$ (the low-quality asset is in the larger pool, creating initial overvaluation $\rho_0 \approx -0.15$). Then the left-hand side of (□) equals $T(1-\alpha)\lambda = 1 \cdot 0.05 \cdot 0.05 = 0.0025$. The right-hand side equals $\bar{F}_2 \cdot (1-\alpha^T) \cdot (1-\rho_0)^{-1} = 1 \cdot 0.05 \cdot (1/1.15) \approx 0.0435$. Since $0.0025 < 0.0435$ the strict inequality (□) holds at this point. Because the functions on both sides of (□) are jointly continuous in $(q_1, q_2, \tau_1, \tau_2, \bar{F}_1, \bar{F}_2, \lambda, T, \alpha)$, the solution set is open in \mathbb{R}^9 (including α as a free coordinate), and therefore has positive Lebesgue measure in the eight-dimensional cross-section for fixed α in a neighborhood of 0.95. Varying α within $(0.9, 1)$ preserves the inequality by continuity. □

B.2.1 Spectral analysis of the EWMA flow operator

The attenuation factor $(1-\alpha^T)$ plays the central mechanistic role in Theorem 1: it is the quantity that limits how much a correction pulse of duration T can shift the EWMA-smoothed flow signal. In the proof above this factor was controlled by the elementary Taylor bound $1-\alpha^T \leq T(1-\alpha)$. The purpose of this subsection is to provide the deeper justification: the attenuation is exactly what the EWMA filter's frequency-domain structure predicts for a high-frequency pulse in the Theorem 1 parameter regime. This connects the algebraic condition (□) to a spectral-separation principle that holds for any linear time-invariant smoother with the same low-pass character.

The EWMA as a linear convolution operator. Define the EWMA smoother $H_\alpha : \ell^2(\mathbb{Z}_+) \rightarrow \ell^2(\mathbb{Z}_+)$ acting on the flow sequence $(f(t))_{t \geq 0}$ by

$$s(t) = \sum_{k=0}^t (1-\alpha) \alpha^k f(t-k).$$

This is a causal LTI convolution operator with impulse response $h(k) = (1-\alpha)\alpha^k$ for $k \geq 0$ and $h(k) = 0$ for $k < 0$. The operator is causal because $s(t)$ depends only on $(f(t), f(t-1), \dots)$, and LTI because the weights $h(k)$ are independent of t . The sequence $(h(k))_{k \geq 0}$ is summable: $\sum_{k=0}^{\infty} h(k) = (1-\alpha) \sum_{k=0}^{\infty} \alpha^k = 1$, so H_α is a probability measure on the non-negative integers and acts as a weighted average of the flow history.

Transfer function and frequency response. The z-transform of the impulse response is $H(z) = (1-\alpha)/(1-\alpha z^{-1})$, which converges for $|z| > \alpha$. Evaluating on the unit circle $z = e^{i\omega}$ gives the frequency response $H(e^{i\omega}) = (1-\alpha)/(1-\alpha e^{-i\omega})$, with magnitude $|H(\omega)| = (1-\alpha) / \sqrt{1 - 2\alpha \cos \omega + \alpha^2}$. This is a first-order IIR low-pass filter. At $\omega = 0$ (DC, the long-run average), $|H(0)| = 1$, so the filter passes persistent flows with unit gain. At $\omega = \pi$ (the Nyquist frequency, the highest-frequency component), $|H(\pi)| = (1-\alpha)/(1+\alpha)$, which is much less than 1 when α is close to 1. The -3 dB cutoff frequency ω_c , defined by $|H(\omega_c)|^2 = 1/2$, satisfies $\omega_c \approx 1-\alpha$ for α close to 1. All frequencies above $\omega_c \approx 1-\alpha$ are attenuated.

Eigenvalue structure. For the discrete-time LTI operator H_α on $\ell^2(\mathbb{Z}_+)$, the frequency-domain decomposition identifies $H(e^{i\omega})$ as the eigenvalue corresponding to a complex-exponential input $e^{i\omega t}$. The relevant eigenvalues are: (i) $H(0) = 1$, the DC eigenvalue — persistent fundamental staking flows, which have significant energy at $\omega = 0$, pass through the EWMA unchanged; and (ii) $|H(\omega)| \approx (1-\alpha)/|\omega|$ for $\omega \ll 1-\alpha$, the high-frequency regime — eigenvalues decay as $1/|\omega|$ (first-order roll-off), so flow components at frequencies much larger than the filter's bandwidth $1-\alpha$ are suppressed by a factor proportional to

$(1-\alpha)/\omega$. The dominant eigenvalue $H(0) = 1$ corresponds to baseline fundamental staking flows \bar{F}_2 ; the correction pulse, which operates at a frequency $\omega_{\text{pulse}} \ll 1-\alpha$, encounters a much smaller eigenvalue.

Spectral content of the correction pulse. A rectangular correction pulse of duration T periods has spectral content concentrated at frequencies $\omega \sim 2\pi/T$. The discrete-time Fourier transform of the indicator function $1_{\{[0, T-1]\}}$ is a Dirichlet kernel whose main lobe has its first null at $\omega = 2\pi/T$, so the dominant spectral energy sits at $|\omega_{\text{pulse}}| \sim 2\pi/T$. In the Theorem 1 parameter regime, $T \ll (1-\alpha)^{-1}$, so $\omega_{\text{pulse}} \sim 2\pi/T \ll 2\pi(1-\alpha) \ll 1-\alpha$, and the pulse frequency is well into the high-frequency attenuation region of H_α . At this frequency the EWMA magnitude evaluates to $|H(\omega_{\text{pulse}})| \approx (1-\alpha)/|\omega_{\text{pulse}}| \approx T(1-\alpha)/(2\pi)$ — up to the constant 2π , this is of order $T(1-\alpha)$. When $T(1-\alpha) \ll 1$ (the Theorem 1 regime), the EWMA attenuates the correction pulse by a factor of order $T(1-\alpha)$ relative to the DC component, providing the formal spectral justification for the bound $1-\alpha^T \leq T(1-\alpha)$ used in condition (□).

Connection to Theorem 1 and condition (□). The spectral picture reframes condition (□) as a statement about filter eigenvalues. Condition (□) requires $T \cdot (1-\alpha) \cdot \lambda < \bar{F}_2 \cdot (1-\alpha^T) \cdot (1-\rho_0)^{-1}$. The left-hand side is proportional to $|H(\omega_{\text{pulse}})| \cdot \lambda$: the correction intensity filtered through the EWMA eigenvalue at the pulse frequency — the filtered correction amplitude that reaches the emission formula. The right-hand side contains $\bar{F}_2 = |H(0)| \cdot \bar{F}_2$: the fundamental flow, which passes through the DC eigenvalue unchanged and therefore exerts its full influence on the emission share. Condition (□) thus encodes a spectral separation: the correction pulse's filtered amplitude (small eigenvalue, order $T(1-\alpha)$) must be dominated by the fundamental flow's unattenuated DC component. When this separation holds, the EWMA's high-frequency suppression prevents the correction signal from reaching the emission formula with sufficient strength, and Rule S exhibits the sign-flip: short selling degrades allocation precisely because the EWMA filters out the informational content of the correction while passing fundamental flows unchanged. This spectral account is consistent with the time-domain derivation in B.3.0 below.

B.3.0 Derivation of mispricing dynamics from AMM primitives

The subsidy-intensity frontier of Theorem 2 relies on curvature properties of the steady-state mispricing functions $M_R(\lambda)$ and $M_S(\lambda)$. Earlier drafts posited these properties as an Assumption M. This subsection derives them from the AMM primitives and flow dynamics introduced in §B.1, so that what follows in §B.3 is a derived lemma rather than an assumption.

Primitives. Each subnet runs a constant-product AMM with invariant $k_i = \tau_i \cdot \alpha_i$ and spot price $p_i = \tau_i/\alpha_i$. Let $p^*_i = \kappa \cdot q_i$ denote the fundamental price anchored to subnet quality, with $\kappa > 0$ a units constant common across subnets. The instantaneous mispricing is $m_i(t) = p_i(t) - p^*_i$. Net staking flows $f_i(t)$ feed into the AMM through the linearized response $dp_i/dt = 2p_i f_i/\tau_i$ of §B.1, so the mispricing obeys $dm_i/dt = (2p_i/\tau_i) f_i(t)$. Under Rule R, one-period emissions to subnet i are proportional to $p_i/\Sigma p$; under Rule S, they are proportional to $\max(s_i, 0)/\max(\Sigma s_j, \varepsilon)$, where $s_i(t) = \phi s_i(t-1) + (1-\phi) f_i(t)$ is the EWMA of flows with decay $\phi \ll (0, 1)$ (the same EWMA smoother denoted α in §B.1; ϕ is used here to avoid collision with λ). A short seller at intensity λ supplies a corrective flow $c_i(t) = -\lambda \cdot \psi(m_i(t)) \cdot \sqrt{\tau_i}$, where ψ is smooth, strictly increasing, with $\psi(0) = 0$ and $\psi'(0) > 0$ — a first-order Taylor expansion gives $\psi(m) \approx \psi'(0) m$ near the steady state.

Rule R: derivation of M_R . Under Rule R, emissions are a monotone function of price, so each unit of mispricing generates feedback proportional to the current gap: the higher m_i , the more the mechanism channels emissions into subnet i , which in turn lifts α_i (through the restaking fraction of §A) and lowers the price back toward p^*_i . Denote this corrective coefficient by $\beta > 0$. To derive β explicitly, note that a mispricing $m_i = \varepsilon > 0$ generates excess emissions $\delta(\Delta\tau_i) = E\varepsilon/P$ (where E is total per-epoch TAO emission and $P = \sum_j p_j$ is the cross-sectional price sum). Of this, the restaked fraction $r_R = 0.18$ mints new α -tokens at rate $r_R E\varepsilon/(P p^*_i)$ per epoch (increasing α_i , lowering $p_i = \tau_i/\alpha_i$), while the sold fraction $r_S = 0.01$ removes α -tokens at rate $r_S E\varepsilon/(P p^*_i)$ (partially offsetting the restaking). The net change in p_i per epoch is $-(r_R - r_S) E p^*_i \varepsilon/(P \tau_i)$, giving $\beta = (r_R - r_S) E p^*_i / (P \tau_i) = 0.17 E p^*_i / (P \tau_i)$. With $r_R = 0.18$ and $r_S = 0.01$ from §A, $\beta > 0$ for any $E, p^*_i, P, \tau_i > 0$, and β depends only on the AMM pool geometry and the emission-restaking calibration — not on λ .

This β , combined with the short-selling channel, yields the mean-field dynamics near steady state $dm_i/dt = -\lambda \cdot \psi'(0) \cdot m_i - \beta \cdot m_i + \eta_i(t)$, where η_i is a bounded exogenous forcing term representing fundamental flow shocks with $\square[\eta_i] = \bar{\eta}_i > 0$ (so $M_R(0) = \bar{\eta}_i/\beta \geq 0$). Taking η_i absolutely continuous with respect to Lebesgue measure — satisfied by the calibrated simulation in which η_i is a sum of independent continuous flow components (§3.3), and used below to supply the minorization step in the stationarity argument for Rule S — the steady-state condition $dm_i/dt = 0$ gives, in expectation, $M_R(\lambda) = \bar{\eta}_i / (\beta + \lambda \psi'(0)) = M_R(0) / (1 + \lambda/\lambda_0)$, with $\lambda_0 = \beta/\psi'(0) > 0$. Differentiating twice, $M'_R(\lambda) = -(M_R(0)/\lambda_0)/(1 + \lambda/\lambda_0)^2 < 0$, and $M''_R(\lambda) = (2 M_R(0)/\lambda_0^2)/(1 + \lambda/\lambda_0)^3 > 0$, so M_R is strictly decreasing and strictly convex on $(0, \infty)$. The first inequality is the closed-feedback effect; the second captures diminishing marginal returns — the first units of shorting target the largest gaps, residual gaps are progressively harder to close. Both properties follow from the AMM invariant and the price-proportional emission rule.

Rule S: derivation of M_S . Under Rule S, the emission-allocation signal is $s_i(t) = \phi s_i(t-1) + (1-\phi) f_i(t)$, a low-pass filter on total flows. The frequency response of the EWMA to a sinusoidal component at angular frequency ω has magnitude $(1-\phi)/\sqrt{1 - 2\phi \cos \omega + \phi^2}$, which collapses to 1 at $\omega = 0$ and to order $(1-\phi)$ for $\omega \gtrsim 1-\phi$. The short-selling correction c_i enters the flow only while $|m_i|$ exceeds the short seller's detection threshold — a short, high-frequency pulse of duration $T \square (1-\phi)^{-1}$, exactly the regime in which Theorem 1 operates. Its EWMA footprint is therefore attenuated by a factor $(1-\phi^T) \approx T(1-\phi) \square 1$, so almost none of the corrective flow reaches the emission share. The feedback loop is effectively open: emissions do not respond to the correction, and the AMM alone must absorb both the informational shock and the slippage cost.

The short seller's inventory constraint caps each pulse at 3% of τ_i (§A), and the constant-product AMM charges slippage $S(\delta) = \delta^2/k_i$ on a trade of size δ . The per-correction residual mispricing is therefore $r(\lambda) = r_0 + a\lambda^2 + O(\lambda^3)$ for constants $r_0 \geq 0$ and $a > 0$: the constant term is the uncorrected fundamental gap $r_0 = M_R(0) = M_S(0) > 0$; the quadratic term arises because the short seller must route a flow proportional to λ through an AMM whose price impact is quadratic in flow, and without emission feedback this residual accumulates rather than dissipating.

Stationarity of the pulse distribution. The averaging step that follows presupposes a stationary joint law for the state (m_i, s_i) and the induced pulse point process, and we establish this as follows. The exogenous

forcing η_i is bounded with $\square[\eta_i] = \eta_i^- > 0$; the corrective flow is capped by the inventory constraint $|c_i(t)| \leq \lambda \cdot 0.03 \cdot \sqrt{\tau_i}$ of §A; and the EWMA discount $\phi < 1$ gives s_i a finite exponential memory kernel. The joint state (m_i, s_i) is therefore confined to a compact subset of \mathbb{R}^2 . The EWMA recursion $s_i(t) = \phi s_i(t-1) + (1-\phi) f_i(t)$ depends only on the previous state and the current flow, so the joint process (m_i, s_i) is Markov at the epoch scale. On this compact state space the Markov kernel admits at least one invariant probability measure by Krylov–Bogolyubov construction (Meyn and Tweedie 1993, Theorem 10.0.1), and the absolute continuity of η_i supplies Doeblin minorization on a small set of positive measure around p_i^* , yielding uniqueness and geometric ergodic convergence from any initial condition (Meyn and Tweedie 1993, Theorem 13.0.1). The short seller's detection threshold is time-invariant, so the pulse point process induced by upcrossings of $|m_i(t)|$ through the threshold is stationary by Palm theory, and the mean pulse duration T and conditional amplitude distribution used below are first moments under this unique stationary law.

Averaging over the stationary distribution of pulses and using $\psi'(0) > 0$ yields $M_S(\lambda) = M_S(0) + c_1\lambda + c_2\lambda^2 + O(\lambda^3)$, with $c_1 = a_1 \cdot T(1-\phi) \geq 0$ (non-negative because any leakage of the correction into s_i biases the signal toward underallocation) and $c_2 = a_2 > 0$ a strictly positive constant inherited from the AMM's quadratic slippage. Differentiating: $M_S(0) > 0$, $M'_S(0) = c_1 \geq 0$, $M''_S(0) = 2c_2 > 0$ — exactly the properties listed below. The positive curvature $M''_S(0) > 0$ is a direct consequence of constant-product slippage combined with an open feedback loop: the AMM converts short-seller intensity into quadratic price damage, and the EWMA prevents the mechanism from correcting that damage in the next epoch.

Conclusion. The properties listed below — earlier called Assumption M — have been derived from (a) the constant-product AMM invariant and its linearized price response, (b) the restaking channel that gives Rule R its corrective coefficient β , (c) the frequency response of the Rule S EWMA to high-frequency corrective pulses, and (d) the compactness and ergodicity of the joint state space under bounded forcing and the inventory cap, which establish the existence and uniqueness of the stationary pulse law. No curvature has been assumed that is not already implicit in these primitives. The block is retained below under the label Lemma M to flag that it is a consequence of §B.1 rather than an additional hypothesis.

B.3 Theorem 2: Subsidy-intensity frontier

Lemma M (Derived mispricing dynamics). *The steady-state mispricing under each rule, $M(\lambda)$, is a C^2 function of the short-selling intensity λ satisfying: (i) $M(0) \geq 0$; (ii) under Rule R, $M'_R(\lambda) < 0$ and $M''_R(\lambda) \geq 0$ on $(0, \infty)$; (iii) under Rule S, $M_S(0) > 0$, $M'_S(0) \geq 0$, and $M''_S(0) > 0$. Lemma M is established in §B.3.0 from the AMM primitives and flow dynamics of §B.1.*

Theorem 2 (Subsidy-intensity frontier). *Under the mispricing dynamics of Lemma M, let $G_R(\lambda) = \lambda \cdot M_R(\lambda) \cdot c_R$ and $G_S(\lambda) = \lambda \cdot M_S(\lambda) \cdot c_S$ be the gross corrections per trade under Rules R and S, with $c_R, c_S > 0$ constants absorbing units, and let $S(\lambda) = C_{Slip} \cdot \lambda^2$ be the constant-product AMM slippage. Then:*

- (a) G_R is strictly concave in λ on $(0, \infty)$, with $G_R(0) = 0$ and $G'_R(0) > 0$.
- (b) G_S is strictly convex in λ on $(0, \infty)$, with $G_S(0) = 0$ and $G'_S(0) \geq 0$.
- (c) $S(\lambda)$ is strictly convex, with $S(0) = 0$.

(d) If $G'_R(0) > G'_S(0)$, there exists a unique crossover $\tilde{\lambda} > 0$ such that $G_R(\lambda) > G_S(\lambda)$ for $\lambda < \tilde{\lambda}$ and $G_S(\lambda) > G_R(\lambda)$ for $\lambda > \tilde{\lambda}$.

(e) The minimum per-trade subsidy $s^*(\lambda) = \max(0, S(\lambda) - G(\lambda))$ required to make short selling break even satisfies $s^*_R(\lambda) < s^*_S(\lambda)$ for $\lambda < \tilde{\lambda}$ and $s^*_S(\lambda) < s^*_R(\lambda)$ for $\lambda > \tilde{\lambda}$.

Proof. Parts (a)–(c) follow from Lemma M and the AMM mechanics. $G_R(\lambda) = \lambda \cdot M_R(\lambda) \cdot c_R$ with $M'_R(\lambda) < 0$ and $M''_R(\lambda) \geq 0$ (Lemma M.ii) yields $G'_R(\lambda) = c_R[M_R(\lambda) + \lambda M'_R(\lambda)]$ and $G''_R(\lambda) = c_R[2 M'_R(\lambda) + \lambda M''_R(\lambda)] < 0$ for $\lambda > 0$, establishing strict concavity. Similarly, $G_S(\lambda) = \lambda \cdot M_S(\lambda) \cdot c_S$ with $M_S(0) > 0$, $M'_S(0) \geq 0$, and $M''_S(0) > 0$ (Lemma M.iii) yields $G''_S(0) = c_S \cdot [2 M'_S(0) + M_S(0) M''_S(0)] > 0$, establishing strict convexity at the origin, which extends to $(0, \infty)$ by the monotone growth of M_S . For part (c), the slippage cost of a trade of size δ scales as $S(\lambda) = C_{\text{Slip}} \cdot \lambda^2$ with $C_{\text{Slip}} > 0$.

Crossover (part d). Define $h(\lambda) = G_R(\lambda) - G_S(\lambda)$. Then $h(0) = 0$, $h'(0) = c_R M_R(0) - c_S M_S(0) > 0$ by hypothesis, and h is strictly concave as the sum of a concave and a negative-of-convex function. A strictly concave function with $h(0) = 0$ and $h'(0) > 0$ crosses 0 at exactly one point $\tilde{\lambda} > 0$, establishing existence and uniqueness. The empirical crossover in the simulation is $\tilde{\lambda} \approx 0.15$: at $\lambda = 0.01$, dTAO gross is 0.0035% while Taoflow gross is 0.0021% (G_R dominates); at $\lambda = 0.30$, dTAO gross is 0.0356% while Taoflow gross is 0.0450% (G_S dominates).

Minimum subsidy. The short seller's private payoff per trade is $G(\lambda) - S(\lambda)$, which is negative under both rules at all tested λ . A protocol subsidy of s per trade shifts the payoff to $G(\lambda) - S(\lambda) + s$, so $s^*(\lambda) = \max(0, S(\lambda) - G(\lambda))$. Under either rule, $S(\lambda) - G(\lambda) > 0$ in the regime of interest, so $s^*(\lambda) = S(\lambda) - G(\lambda)$. Subtracting: $s^*_R(\lambda) - s^*_S(\lambda) = G_S(\lambda) - G_R(\lambda)$. By part (d), the sign flips at $\tilde{\lambda}$, establishing (e). \square

The empirical crossover $\tilde{\lambda} \approx 0.15$ matches this characterization. The policy implication: subsidies to sustain short selling are more efficient under dTAO at low-to-moderate intensity ($\lambda < 0.15$) and more efficient under Taoflow at high intensity ($\lambda > 0.15$). Because the socially optimal intensity under dTAO is higher than the privately profitable intensity (the mechanism captures the correction before the short seller can profit), subsidies are needed specifically in the low-to-moderate range where dTAO's subsidy efficiency is superior.

B.4 Proposition 3: Empirical genericity

Proposition 3 (Empirical genericity). Let Π be the 12-dimensional parameter space defined in Appendix C (flow-component weights, noise df , autocorrelations, flow scale, flow base multiplier, quality-pool correlation, EWMA half-life). Let $\Pi_{\text{valid}} \subseteq \Pi$ be the subset of configurations satisfying the validation constraints: price-quality correlation $\subseteq [0.50, 0.90]$, top-10 concentration $\subseteq [0.40, 0.80]$, excess kurtosis $\subseteq [3, 25]$. Let $\text{SignFlip} \subseteq \Pi_{\text{valid}}$ be the subset of configurations for which the simulation produces positive $\Delta\rho$ under dTAO and negative $\Delta\rho$ under Taoflow. Then $\Pr(\text{SignFlip} \mid \Pi_{\text{valid}})$ is consistent with 1.0, with a 95% one-sided Clopper–Pearson binomial lower bound of at least 0.9497.

Verification. We sample 100 parameterizations via Latin hypercube from Π , retain the 58 that pass Π_{valid} , and compute $\Delta\rho$ under both mechanisms for each at 20 replications per condition. All 58 of 58 passing configurations (100%) show both signs correct: dTAO $\Delta\rho > 0$ and Taoflow $\Delta\rho < 0$. The mean dTAO $\Delta\rho$ is +0.202 (range [+0.085, +0.324]); the mean Taoflow $\Delta\rho$ is -0.087 (range [-0.131, -0.046]). No exceptions. For 58 of 58 successes, the 95% one-sided Clopper–Pearson lower bound on the sign-flip proportion is 0.9497.

Proposition 3 is labeled a proposition rather than a theorem because its justification is empirical rather than analytical. It does not replace Theorem 1, which establishes the sign-flip analytically in the two-asset model; instead, Proposition 3 provides empirical evidence that the phenomenon identified in Theorem 1 extends robustly to the N-asset calibrated setting with realistic parameter distributions.

B.5 Corollary 1: Hybrid monotonicity; Observation 1: The flat plateau

Corollary 1 (Hybrid monotonicity). *Let $B(w) = w \cdot R + (1-w) \cdot S$ be a hybrid allocation rule. Under the conditions of Theorem 1, there exists $w^* \in (0, 1)$ such that the short-selling correction improves allocation under $B(w)$ for $w > w^*$ and degrades it for $w < w^*$. The function $\Delta\rho(w)$ is continuous and monotonically increasing in w on an open neighborhood of w^* .*

The proof follows from the linearity of B in w and the sign-flip established in Theorem 1. The empirical $w^* \approx 0.1$ (from the v10 hybrid frontier data, where hy_w00_short has $\Delta\rho = -0.109$ and hy_w01_short has $\Delta\rho = -0.005$, with the sign change occurring near $w = 0.1$).

Observation 1 (Flat hybrid plateau, empirical). *In the 50-replication simulation with paired seeds, the allocation accuracy under short selling is statistically constant across $w \in [0.3, 0.8]$, with grand mean $\rho \approx 0.928$ and a one-way ANOVA p -value of 0.983 for the null of equality. This is an empirical finding from the calibrated simulation, not a consequence of the two-asset analytical model.*

The plateau admits a heuristic explanation consistent with the model's structure. For $w \geq 0.3$ the price-proportional component receives sufficient weight to capture short-selling corrections almost completely, while for $w \leq 0.8$ the EWMA component retains enough weight to track persistent fundamental staking flows. In this interior range, both channels contribute quality-tracking information, and the accuracy ceiling is determined by the information content of the combined signal rather than by the specific mixing weight w . The plateau is therefore a saturation phenomenon: beyond $w \approx 0.3$, the marginal benefit of shifting further weight toward the price channel is approximately offset by the marginal cost of losing the EWMA's noise-suppression benefit on the non-corrective component of flows. The plateau ends near $w \approx 0.8$ because, beyond that point, the EWMA weight becomes too dilute to buffer transient flow noise, causing accuracy to decline toward the noisier pure-dTAO level. Mechanism designers choosing a hybrid weight have a wide margin of insensitivity — any $w \in [0.3, 0.8]$ achieves the same allocation quality under short selling.

B.6 Theorem 4: Separating equilibrium under filtered Taoflow

The filtered Taoflow mechanism excludes short-selling flows from the EWMA calculation. This requires the protocol to distinguish short flows from organic unstaking flows, which creates a routing-incentive game. We model this as a simple Bayesian game.

Setup. There are two agent types. Type U (organic unstakers) sell alpha tokens they already hold, motivated by liquidity needs or portfolio rebalancing unrelated to quality assessment. Type S (short sellers) borrow alpha tokens from a protocol lending pool at interest rate $r > 0$ and sell them to correct overvaluation. Both agent types can choose between two routing channels: direct AMM sale (flow enters the EWMA) or borrow-and-sell through the lending contract (flow excluded from the EWMA). Type U does not earn a correction profit and does not pay a borrow rate unless they route through the lending contract. Type S earns gross correction $G(\lambda)$ at slippage cost $S(\lambda)$. The protocol offers a correction subsidy $s \geq 0$ per borrowed-and-sold position, paid only through the lending channel, financed from the emission budget. The protocol cannot distinguish Type U from Type S directly, but it can observe which channel each flow routes through.

Theorem 4 (Separating equilibrium). *Assume (i) the correction subsidy exceeds the borrow rate, $s > r$, and (ii) the short seller is profitable after subsidy, $G(\lambda) - S(\lambda) + s - r > 0$. Then there exists a Bayesian Nash equilibrium in pure strategies in which: (a) all Type U agents sell through the direct AMM channel; (b) all Type S agents sell through the lending-pool channel; (c) the EWMA receives only Type U flows, and the filtered Taoflow mechanism achieves its baseline allocation accuracy ($\rho \approx 0.879$) in equilibrium.*

Proof. Consider the strategy profile in which Type U plays "direct sale" and Type S plays "lending pool." Type U's best response: Type U earns no correction profit and incurs no liquidation externality on own-subnet emissions. Direct sale yields payoff 0. Lending-pool route yields payoff $-r - s'_U$, where s'_U is the subsidy Type U would receive. Because the subsidy s is conditional on the lending contract flagging a position as a short-correction trade and Type U does not have a short position to close, Type U receives $s'_U = 0$ and pays only $-r$. Direct sale is therefore strictly preferred by Type U: $0 > -r$ for $r > 0$.

Type S's best response: if Type S routes through the lending pool, their payoff is $G(\lambda) - S(\lambda) - r + s > 0$ by assumption (ii). If Type S deviates to direct sale, they avoid the borrow rate but also forgo the subsidy, yielding $G(\lambda) - S(\lambda)$. The difference is (direct - lending) = $(G(\lambda) - S(\lambda)) - (G(\lambda) - S(\lambda) - r + s) = r - s < 0$ by assumption (i). Therefore the lending pool is strictly preferred. Both type-conditional best responses strictly prefer the prescribed actions. The strategy profile is a strict Bayesian Nash equilibrium. In equilibrium the EWMA receives only Type U flows, and the filtered mechanism's allocation properties are realized.

The two conditions are observable and governance-set: r is the lending contract's interest rate, and s is a per-trade subsidy funded from the emission budget. The ratio s/r must exceed 1, and the subsidy must be large enough to make short selling profitable after costs. A simple sufficient rule: set $s = 2r$ so $s - r = r$ covers the natural positive profit margin for informed shorts at moderate intensities. The filtered Taoflow mechanism is therefore implementable, not merely theoretically valid, provided the protocol can maintain a subsidy that satisfies $s > r$ and the budget constraint holds.

Appendix C: Parameter Sensitivity Analysis

To address the concern that the 2×2 interaction may be an artifact of specific parameter choices, we sample 100 parameterizations via Latin hypercube from the prior over 12 model parameters: flow-component weights (w_{fund} , w_{mom} , w_{herd} , $w_{\text{noise_persist}}$, $w_{\text{noise_transient}}$), noise degrees of freedom, persistence and momentum autocorrelations, flow scale and base multiplier, quality-pool correlation, and EWMA half-life. Each parameterization is screened against tightened validation targets: price-quality correlation $\in [0.50, 0.90]$, top-10 concentration $\in [0.40, 0.80]$, excess kurtosis $\in [3, 25]$. These ranges are substantially tighter than the main analysis (which uses $[0.25, 0.90]$, $[0.35, 0.85]$, $[3, 35]$).

For each passing configuration, we compute the allocation accuracy change from short selling ($\lambda = 0.20$) under both dTAO and Taoflow at twenty replications per condition with paired seeds. This 20-replication resolution is an order of magnitude higher than an earlier exploratory run at 2 replications per condition, which left three dTAO "exceptions" (deltas of -0.008 , -0.003 , -0.001) within one standard error of zero and therefore statistically indistinguishable from the positive sign-flip result.

Of 100 sampled parameterizations, 58 pass the tightened validation screen. Of these 58 passing configurations, the Taoflow sign is negative in 58 of 58 (100%), the dTAO sign is positive in 58 of 58 (100%), and the full 2×2 sign-flip holds in 58 of 58 (100%). The 95% one-sided Clopper–Pearson binomial lower bound on the sign-flip proportion is 94.97%. No parameterization produces a reversed interaction. The mean dTAO delta across passing configurations is $+0.202$ (range $[+0.085, +0.324]$, std 0.066); the mean Taoflow delta is -0.087 (range $[-0.131, -0.046]$, std 0.019). The minimum positive dTAO delta ($+0.085$) is an order of magnitude larger than the noise floor from the exploratory run, resolving the earlier ambiguity definitively in favor of the sign-flip.

The 100% hit rate at $N = 20$ is consistent with the expectation that the original three exceptions at $N = 2$ were simulation noise, not genuine boundary cases. At $N = 20$, per-configuration standard errors drop by a factor of $\sqrt{10} \approx 3.2$, and the effect sizes in the passing-configuration population remain well-separated from zero. The interaction is therefore not an artifact of a specific parameterization. It emerges from the structural relationship between the emission rule and the information channel, as Theorem 1 predicts: any positive smoothing parameter attenuates transient corrections, and the sign-flip occurs across the full space of validation-passing flow decompositions.