LIVE: Priority-Ordered Signal Distribution under Market Impact

The native access token of LiveAgents: a mechanism for allocating execution precedence via time-weighted on-chain stake, with a formal model of alpha decay in sequentially executed agent strategies

LiveAgents Research · liveagents.org ·
Draft v1.1 - September 2026

Abstract. Autonomous trading agents that publish investable signals face an unavoidable microstructure constraint: executing a signal consumes it. The root cause is microstructural: exchange liquidity is finite, order books are shallow relative to strategy flow, and matching engines fill in strict price-time priority - so the first orders to arrive consume the best-priced depth and every later order walks a depleted book. Under the empirically robust square-root law of market impact that emerges from this structure, the first tranche of execution permanently moves prices toward the strategy's target, so identical signals delivered to later recipients carry strictly less realizable alpha. Any signal-distribution platform therefore embeds a sequencing decision, whether disclosed or not. This paper formalizes that decision as an explicit, priced, and auditable mechanism. We survey the matching-engine architectures the platform trades across (price-time CLOBs, pro-rata engines, constant-function AMMs, and brokered routing) and show that in each of them earlier arrival buys better execution. We model per-wave implementation shortfall and fill-ratio decay for a five-wave release schedule, derive the net alpha available to each wave as the composition of exponential information decay and cumulative permanent impact, and specify a Sybil- and manipulation-resistant qualification rule based on the time-weighted average balance (TWAB) of staked LIVE valued at a 30-day TWAP. We prove that the cost of spoofing wave standing scales linearly with the averaging window, derive the causal chain from staked balance through release latency and fill quality to net alpha - including its simulated correlation structure, and the break-even capital at which staking is rational for a risk-averse allocator, and specify the on-chain attestation scheme by which every per-wave fill is published so that the priced priority premium is a measured quantity rather than a marketing claim. All numerical results are model-based simulations under stated parameters; no live track record is claimed. The token confers execution precedence only: it carries no revenue share, profit participation, or claim on managed assets, and every wave is equally reachable through a posted direct subscription payable in BTC, ETH, or major stablecoins.

Keywords: market impact, implementation shortfall, alpha decay, mechanism design, token staking, time-weighted average balance, execution sequencing, transparency.

Nomenclature. LIVE is the token - the on-chain asset whose time-weighted staked balance determines execution-wave standing, and it is the sole subject of the tokenomics in §8. **LiveAgents** (liveagents.org) is the agentic trading platform on which the agents, waves, and measurements described here run. The two are not interchangeable: LIVE is the asset, and LiveAgents is the platform on which holding it buys precedence. Earlier drafts of this paper called the token PAI; LIVE is the correct and final name, and the mechanism is unchanged.

1 Introduction and problem statement

LiveAgents - the agentic trading platform - operates a marketplace of autonomous research and trading agents whose opportunity identification, committee approval, and execution are recorded as an immutable event log. The architecture is strictly non-custodial: client capital never touches the platform. Funds remain in each allocator's own brokerage or exchange account (Alpaca for US equities and spot crypto; Bybit for crypto derivatives), and the platform holds only trade-scoped API permissions, no withdrawal rights, no pooling, no omnibus account. What the platform controls is exactly one thing: the *sequence* in which each account's copy of an agent's buy or sell decision is released to that account's own venue. When an agent's investment committee approves a trade, the resulting signal is simultaneously (i) an instruction to deploy the agent's own book and (ii) an investable disclosure to subscribed allocators. The signal cannot be delivered to all subscribers "simultaneously" in any economically meaningful sense: order-book depth is finite, matching engines serve price-time priority, and each executed tranche consumes the best-priced liquidity and moves the market against every subsequent one (§3.1).

Conventional copy-trading platforms resolve this silently - by latency lottery, by fee tier, or by undisclosed internal ordering - leaving late recipients to absorb an unmeasured and undisclosed transfer. We take the opposite approach: sequencing is made explicit, its cost is measured and published per trade, and precedence is allocated by a transparent rule tied to sustained, on-chain, time-weighted stake in the LIVE token. The design goal is not to eliminate the queue - that is physically impossible - but to make the queue *legible, priced, manipulation-resistant, and continuously verifiable* by the risk-averse allocators it serves.

This paper is written for three audiences. **Quants and strategy developers** who will publish trading agents through LiveAgents: Sections 3–4 give the microstructure model your strategies will execute inside, and §7 the measurement regime your track record will be built from. **Crypto traders** evaluating the token: Sections 5–6 derive exactly what staking buys, in basis points, and what it costs. **Fund managers and allocators** subscribing to agent signals: Sections 3.6 and 6 quantify the relationship between wave standing, fill quality, and the net alpha your capital actually captures - the numbers that belong in your subscription decision. Three questions structure the paper. First, *what exactly does a queue position cost?* Section 3 derives per-wave implementation shortfall from the square-root impact law and information decay. Second, *how is precedence assigned without being gameable?* Sections 4–5 specify the wave mechanism and the TWAB qualification rule, with a manipulation-cost bound. Third, *why would a rational, risk-averse allocator stake at all?* Section 6 derives the break-even capital frontier, and Section 7 specifies the measurement and attestation scheme that lets every claim in this paper be checked against realized fills.

2 Notation and model primitives

Consider a single approved trade in an asset with daily volatility σ (in basis points) and daily traded volume V. The platform executes total quantity Q, partitioned across waves $k \in \{1, 2, 3\}$ with participation rates $q_k = Q_k/V$ released at scheduled offsets $t_1 < t_2 < t_3$ after committee approval. The signal's gross alpha at approval is α_0 (bps), decaying at rate λ as the anticipated move materializes independently of the platform's own trading. Allocator i holds staked balance $B_i(t)$ of LIVE with reference price P^{TWAP} (30-day time-weighted average price). All monetary thresholds θ_w are denominated in USD.

3 Liquidity, order matching, and alpha decay

3.1 The root cause: finite depth under price-time priority

Slippage and alpha decay are not platform artifacts; they originate in the structure of the venues themselves. A limit order book displays finite resting size at each price level, and the depth profile is steep: available size decays rapidly away from the touch, with displayed liquidity at any instant typically a small fraction of one percent of daily volume (Bouchaud, Farmer & Lillo 2009). Matching engines fill incoming marketable orders in strict *price-time priority*: the best-priced resting orders are consumed first, and within a price level, the earliest-arriving orders match first. Two consequences are mechanical. First, any order larger than top-of-book *walks the book* - its marginal fill price rises level by level, so slippage is literally the area between the fill path and the arrival price (Figure 10, right). Second, and decisive for signal distribution: because the engine serves arrival order, the first participants to act consume the cheap depth and leave every later participant a strictly worse book. Liquidity does replenish - market makers refill the ladder - but resilience operates on a timescale of seconds to minutes and refills at *new, higher* price levels once the permanent information content of the flow is impounded. The square-root impact law of §3.2 is the emergent, empirically measured description of exactly this process: latent liquidity theory derives it directly from the shape of the order book that market participants are willing to reveal (Tóth et al. 2011; Donier & Bonart 2015).

Figure 10 — Root cause: finite order-book depth under price-time priority matching

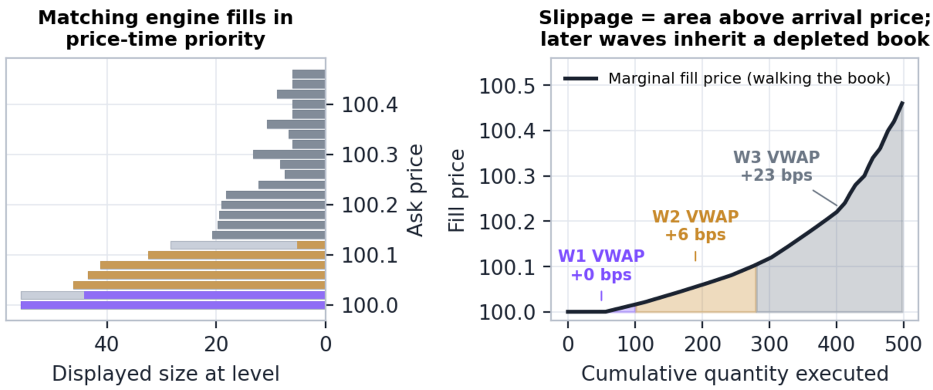

Figure 10. Left: a stylized ask ladder; the matching engine allocates fills in price-time priority, so earlier waves consume the best-priced depth and later waves the remainder (stylized three-cohort illustration; production runs five waves). Right: the marginal fill price walks the book; per-wave VWAP degrades mechanically with queue position.

This is the opportunity the LIVE token exists to price. Given finite depth and priority matching, *some* ordering of signal recipients is imposed by the venue whether the platform acknowledges it or not; the only design freedom is whether that ordering is random, latency-determined, hidden - or explicit, measured, and purchasable. The LIVE mechanism takes the last option: it converts a physical property of exchange microstructure into a transparent priority market.

3.2 A taxonomy of matching engines - and where slippage comes from in each

The venues our agents trade span four distinct matching architectures, and slippage has a different mechanical origin in each. **(i) Price-time priority CLOBs** - the dominant design in crypto (Bybit, Binance, OKX, Coinbase) and in equity exchanges. The engine sorts resting orders first by price, then by arrival time within a price level; incoming marketable flow consumes that sorted queue. Slippage here is book-walking (§3.1, Figure 10): the cost of demanding more size than the touch displays. Arrival time is everything - both for takers (earlier arrival meets a fuller book) and for makers (earlier arrival means a better position in the time-priority queue at your level, hence higher fill probability). **(ii) Pro-rata and hybrid engines** - common in short-rate futures and some options markets. Fills at a price level are allocated proportionally to resting size rather than by arrival time. Slippage manifests as *partial fills*: you receive a fraction of your intended quantity and must chase the remainder at worse prices. Speed matters less within a level but the chase dynamic reintroduces it. **(iii) Automated market makers (AMMs)** - the DEX regime (Uniswap-style constant-function pools). There is no queue at all: price is a deterministic function of pool reserves, and slippage is computable in closed form before you trade. For a constant-product pool, buying a fraction f of reserves costs:

$$\text{CPMM } (xy = k): \quad \overline{P}_{\text{exec}} = \frac{P_{\text{spot}}}{1-f}, \quad P_{\text{post}} = \frac{P_{\text{spot}}}{(1-f)^2}, \quad f = \frac{\Delta y}{y} \quad \Rightarrow \quad \text{premium} = \frac{f}{1-f} \approx f$$

The premium f/(1–f) is the AMM analogue of walking the book - and because the post-trade price moves by (1–f)■^{2}, an AMM broadcasts your impact to the *next* trader with perfect fidelity: sequencing consequences are, if anything, harsher on-chain, with the added adversary of MEV searchers who can insert themselves between observation and execution. **(iv) Brokered best-execution routing** - the US retail equity path (Alpaca). Orders route to wholesalers under best-execution obligations; there is no single book to walk, and slippage appears as spread capture and price-improvement variance rather than visible ladder consumption. Across all four architectures one invariant holds and is the foundation of everything that follows: *the earlier an order arrives relative to correlated flow, the better its execution* - by queue position in (i), by pre-chase allocation in (ii), by pool state in (iii), and by pre-impact NBBO in (iv). The wave mechanism prices that invariant; it does not depend on any one engine design.

3.3 The square-root impact law

A large empirical literature - spanning equities, futures, FX, and crypto - finds that the expected price impact of a metaorder of size Q is well described by a concave, approximately square-root function of participation (Torre & Ferrari 1997; Almgren et al. 2005; Tóth et al. 2011; Donier & Bonart 2015 for Bitcoin). We adopt the standard form, with Y a market-specific constant of order unity and a decomposition into permanent and temporary components with permanent share ϕ:

$$\Delta P_{\text{bps}} \;=\; Y\,\sigma_{\text{bps}}\sqrt{\frac{Q}{V}}, \qquad \Delta P = \Delta P^{\text{perm}} + \Delta P^{\text{temp}}, \quad \Delta P^{\text{perm}} = \phi\,\Delta P$$

Empirical estimates place $Y \in [0.4, 1.0]$ and $\phi \in [0.4, 0.7]$; we use $Y = 0.7$, $\phi = 0.55$ as base-case parameters (Table 1). The concavity of impact in Q is what makes wave partitioning meaningful: splitting flow does not eliminate impact, but the *ordering* of the split determines who bears it.

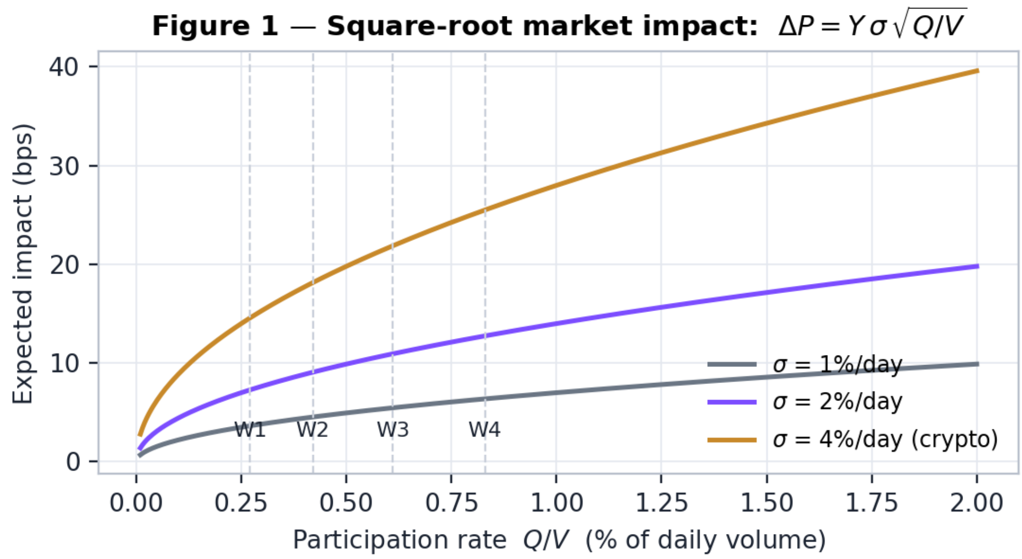

Figure 1. Expected impact as a function of participation rate for three volatility regimes. Dashed verticals mark base-case wave participation rates q1 = 1.08%, q2 = 2.04%, q3 = 3.50% of daily volume.

3.4 Net alpha per wave

Wave k executes after the permanent impact of all earlier waves is impounded in the price, pays its own temporary impact, and captures a signal whose informational content has decayed for t_k seconds. The net alpha available to wave k is therefore:

$$\alpha_k = \alpha_0 e^{-\lambda t_k} - \phi \sum_{j<k} Y\sigma\sqrt{q_j} - (1-\phi)\, Y\sigma\sqrt{q_k}$$

Two properties follow immediately. **Monotonicity:** $\alpha_1 > \alpha_2 > \alpha_3$ whenever $q_j > 0$ - later waves strictly underperform, by construction, not by accident. **Separability:** the degradation decomposes into an information term (governed by λ, dominant for slow signals) and an impact term (governed by Y, σ, ϕ, q, dominant for fast signals in the seconds-to-minutes regime that concerns us). Figure 2 plots the composition under base-case parameters: at t_3 = 180 s the information term has cost only ~2 bps while cumulative permanent impact has cost 19 bps - sequencing risk in this regime is an impact phenomenon.

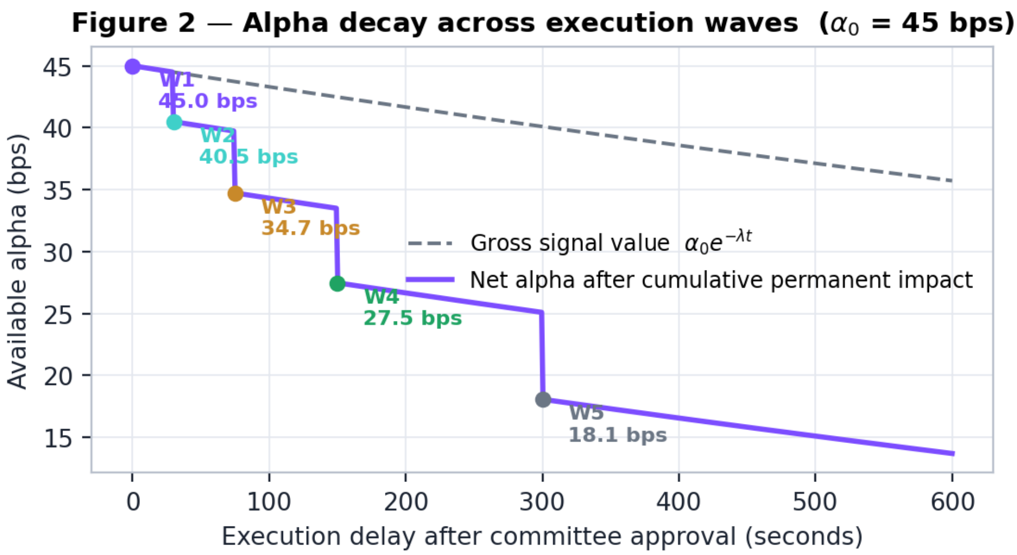

Figure 2. Gross signal value (dashed) and net alpha after cumulative permanent impact (solid) as a function of execution delay. Markers show the five scheduled waves under base-case parameters (α0 = 45 bps, λ = ln2/1800 s, remaining parameters per Table 1).

3.5 Simulated per-wave implementation shortfall

Following Perold (1988), define wave-k implementation shortfall IS_k as the signed difference between the wave's volume-weighted fill and the decision price at committee approval. We simulate 20,000 trades with lognormally distributed participation (σ_{log} = 0.35 around base-case rates) and idiosyncratic execution noise. Results: median IS of 0.0 / 4.1 / 9.1 / 15.2 / 22.3 bps for waves 1–5. These simulated medians are the provenance of the +4 / +9 / +15 / +22 bps figures quoted in platform materials; they are model outputs under Table 1 parameters, to be superseded by measured values (§7) after launch.

Figure 3 — Simulated per-wave shortfall distribution (20,000 trades)

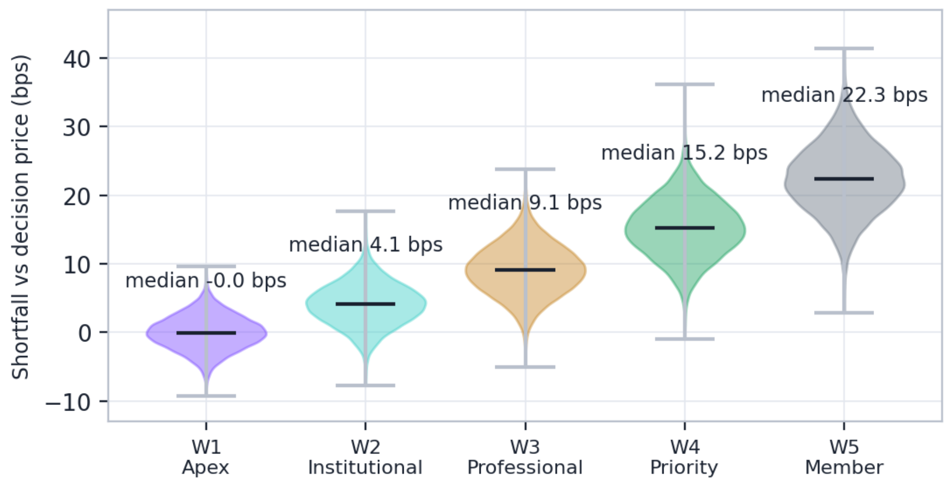

Figure 3. Monte Carlo distribution of per-wave implementation shortfall (20,000 simulated trades, base-case parameters, lognormal participation dispersion). Medians annotated. Illustrative - not a track record.

3.6 Signal speed, fill ratio, and the alpha chain

Sections 3.1–3.5 priced the cost of *arriving late with marketable orders*. This section prices the second, less visible cost of delay: **fill ratio**. Define fill quality FQ as the fraction of intended size executed at or better than the intended level. For marketable orders FQ ≈ 1 and delay is paid in implementation shortfall; for limit orders working at the decision price, FQ itself decays with delay - by the time a late subscriber's order reaches the queue, earlier waves have consumed the levels it hoped to trade at. Worse, the decay is adversely selected: a stale limit order fills *most reliably when the signal was wrong* and the market comes back through it. Figure 11 separates the two: the raw fill ratio declines with delay, and the benign fill ratio, fills that are not adverse selection - declines faster. The gap between the curves is toxic fill: executions you receive precisely because you should not want them (Harris 2003).

Figure 11 — Fill ratio decays with signal delay; surviving fills grow toxic

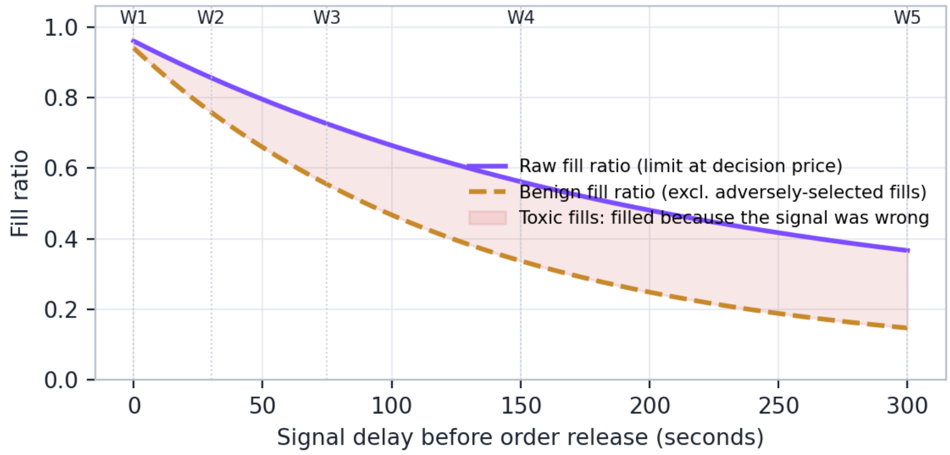

Figure 11. Simulated fill ratio of limit orders working at the decision price, as a function of release delay. Dashed: fills excluding adverse selection. The shaded gap is toxic fill - orders that execute because the thesis failed. Wave release times marked.

Fill quality is not an execution-desk nicety; it is arithmetically the other half of realized alpha. Per trade, $\alpha_{net} = FQ \times \alpha_{avail}(t_w)$ - the alpha still available at your release time, times the fraction of your size that actually captures it. Both factors are decreasing in delay, so realized alpha is doubly convex in queue position. Figure 12 shows the simulated per-trade relationship: net alpha is close to linear in fill quality with ρ ≈ 0.87, and the color gradient shows delay driving both axes simultaneously.

Figure 12 — Net alpha is a linear function of fill quality (20k-trade simulation, 800 shown)

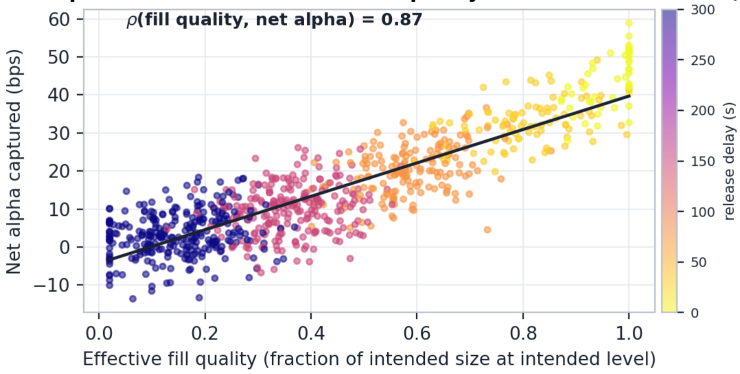

Figure 12. Per-trade net alpha versus effective fill quality (simulation under Table-1 parameters; 800 of 20,000 trades shown, colored by release delay). The correlation is not incidental - it is the identity αnet = FQ · αavail with noise.

The complete causal chain - and the precise sense in which staking LIVE correlates with net alpha - is therefore:

$$\text{stake (PAI TWAB)}\uparrow \;\Rightarrow\; \text{wave}\downarrow \;\Rightarrow\; \text{delay}\downarrow \;\Rightarrow\; \text{FQ}\uparrow,\ \text{IS}\downarrow \;\Rightarrow\; \alpha_{\text{net}} = \text{FQ}\cdot\alpha_{\text{avail}}(t_w)\uparrow$$

Staked TWAB determines wave; wave determines release delay; delay determines both fill quality and implementation shortfall; and those jointly determine the alpha a subscriber's capital nets. Figure 13 shows the end-to-end result under base-case parameters: median net alpha is strictly monotone in stake tier (45.0 / 40.5 / 34.7 / 27.5 / 18.1 bps for W1→W5), a Spearman correlation of 1.0 across tiers by construction, with per-trade dispersion overlapping - a Wave-5 trade can beat a Wave-1 trade on any given day; the tiers separate in expectation, not in every realization. Two disclosures keep this claim honest. First, the correlation is *engineered, disclosed, and priced*, not discovered: the mechanism exists to convert an unavoidable venue-level ordering into a purchasable one. Second, all figures in this section are simulations under Table-1 parameters; the measured versions - realized FQ, IS, and net alpha by wave, per agent, are published continuously under §7, and it is those measured correlations, not these modeled ones, that should drive a subscriber's staking decision.

Figure 13 — Net alpha by staked wave: monotone in stake, Spearman ρ = 1.0 across tiers

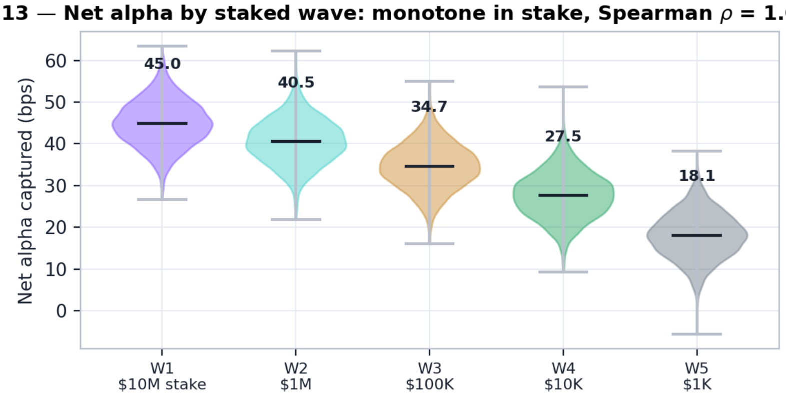

Figure 13. Simulated per-trade net alpha by staked wave. Medians strictly decrease down the ladder; distributions overlap. Model output under Table 1 - measured per-wave values supersede after launch.

Table 1. Base-case model parameters. All downstream numbers in this paper derive from this single parameter set.

Parameter	Symbol	Base case	Source / rationale
Daily volatility	σ	200 bps (2%)	large-cap equity / mid-vol crypto regime
Impact coefficient	Y	0.70	Tóth et al. (2011); Almgren et al. (2005): Y ∊ [0.4, 1.0]
Permanent impact share	ϕ	0.55	Almgren-Chriss decomposition; empirical range 0.4–0.7
Wave participation	q1..q4	0.27% / 0.42% / 0.61% / 0.83% ADV	base-case flow partition (wave 5 = remainder)
Wave release offsets	t1..t5	0 / 30 / 75 / 150 / 300 s	scheduled release policy v2 (five waves)
Gross signal alpha	α0	45 bps	representative imminent-move signal
Information half-life	ln2/λ	1,800 s	imminent-move signal class
Wave thresholds	θ1..θ5	$10M / $1M / $100K / $10K / $1K TWAB	windows 180/120/90/30/7d; ≈ 1% of intended account size; θ5 is universal membership; pegged to the posted direct-subscription price (BTC/stablecoin payable) per epoch
LIVE reference price	P^TWAP	US$0.0444 (30d TWAP)	oracle-attested

4 The wave mechanism

Signal release and execution proceed in five scheduled waves per trade. Each wave is a *cohort of client accounts*, not a pooled order: at each wave's release time the platform submits, via each member account's own API keys, that account's order to that account's own venue. Wave 1 (*Apex*) is the first release cohort, submitted simultaneously with committee approval; Waves 2–4 (*Institutional*, *Professional*, *Priority*) release at offsets t_2–t_4; Wave 5 (*Member*) completes at offset t_5 under FIFO ordering within the cohort. Wave 5 is not free: platform membership itself requires θ_5 = $1,000 of staked LIVE on a 7-day window (or the posted direct-subscription equivalent), so every signal recipient on the platform holds the token - universal skin-in-the-game, while the direct rail preserves severability for those who decline token exposure. There is no netting or aggregation across accounts - every fill lands in the owner's account directly. The mechanism's entire lever is arrival time. The platform does not - and cannot - influence how any venue matches orders; it controls only when each account's order is *released*. Because the venue itself matches in price-time priority (first in, first served at each price level), the release schedule is converted into fill priority by the exchange's own rules: earlier cohorts consume the better-priced depth as a mechanical consequence of §3.1, with no privileged market access of any kind. Three scope clarifications follow. First, sequencing operates *among platform clients relative to one another*: between cohort releases, unrelated market flow interleaves freely, so Wave 1's guarantee is precedence over Waves 2–5 - never precedence over the market. The measured inter-wave premia of §3.5 are defined net of this interleaving. Second, venue microstructure differs: on a central-limit-order-book exchange such as Bybit, price-time priority applies literally and fills walk the displayed ladder as in Figure 10; for US equities routed through a brokerage such as Alpaca, orders are routed to executing venues under best-execution obligations, so the fill path is mediated - arrival sequencing still determines relative fill order and the aggregate impact results hold, but the book-walking picture is an idealization for that venue class. Third, within a cohort, per-account API submissions land with unavoidable millisecond-scale dispersion; the randomized intra-wave ordering declared above converts this jitter into explicit fairness rather than an undisclosed latency race. The venue's matching engine is indifferent to whose account each order belongs to: what walks the book is the aggregate sequenced flow, so the impact and alpha-decay results of §3 apply unchanged to this non-custodial routing architecture. The schedule (t_2, t_3) is a published platform constant, identical across agents and trades within a policy version; no discretionary reordering is possible, and every fill is attributed to its wave in the event log (§7). Within a wave, ordering is randomized per trade to eliminate intra-wave latency games.

The mechanism deliberately trades off two costs. Shorter offsets reduce the impact borne by later waves but concentrate flow and raise instantaneous participation (raising everyone's impact); longer offsets smooth flow but transfer more permanent impact onto later waves. The v1 schedule is chosen so that, under base-case parameters, the staked-wave premia are material enough to price (4 / 9 / 15 / 22 bps down the ladder) but small relative to typical gross signal alpha (45 bps) - later waves still receive positive expected net alpha (Figure 2), which is a design constraint, not a coincidence: a mechanism whose bottom tier receives negative expected net alpha cannot retain its bottom tier.

4.1 On-chain execution: AMM pools and vault subscribers

The wave mechanism extends on-chain along two routes with sharply different properties. **Vaults as subscribers.** On-chain vaults are already structured as delegated-trading pools with share-based accounting (ERC-4626 and equivalents; leader-vault primitives on venues such as Hyperliquid additionally support delegating authorized trading agents with fully on-chain accounting). A vault that authorizes the platform's executor as its agent becomes one more account in a wave cohort: capital remains in the vault contract - the platform custodies nothing - and the vault's on-chain TVL gives the ≈1%-of-capital staking rule a publicly verifiable denominator. Vault operators are therefore a natural institutional subscriber class, and their LIVE stake is sized by observable TVL rather than self-reported account size. The platform does not issue pooled vaults of its own: platform-issued per-wave vaults would convert disclosed sequencing among independent accounts into a collective investment vehicle whose returns vary by token stake - precisely the structure the design excludes (§9). **Direct AMM execution.** Signals may execute as swaps from subscribers' own wallets via scoped session keys, with slippage governed by the closed-form pool mathematics of §3.2(iii). One property changes materially on-chain: transaction ordering is not granted by price-time priority but auctioned to block builders. Wave sequencing therefore survives on-chain only as *inter-block* sequencing via private order flow; within a public mempool, the inter-wave gap is captured by searchers rather than by earlier waves, and on batch-auction venues - which clear all orders in a batch at a uniform price by design - intra-batch priority does not exist and the wave premium compresses toward zero. These venue-dependent properties are disclosed rather than engineered around: the measured per-wave premia of §7 are published per venue class, and a venue on which precedence is worthless will show a premium of zero there. A second-order benefit runs the other way: where execution is on-chain, fills are public by construction, so the per-wave attestation of §7 becomes trustlessly verifiable rather than Merkle-anchored after the fact.

4.2 Reference architecture: a sequencer-gated event log

The mechanism's guarantees are only as strong as the system that enforces them, so the architecture is specified here as part of the design. A natural first instinct - a publish/subscribe bus on which agents publish signals and subscriber accounts receive them for execution - fails on two counts: delivery order does not determine venue arrival order (a later-wave subscriber with faster infrastructure would beat an earlier wave to the exchange), and signal content in a subscriber's hands before execution is an uncontrollable leakage channel. The platform therefore enforces a stricter principle: **release is execution**. Signal content never leaves the platform before orders are placed; what a subscriber 'subscribes' to is sequenced execution in their own account, plus a post-release information feed.

The system is four totally-ordered append-only logs, gated by a single deterministic component. The **Research Log** accepts SignalProposed events from any agent - the one genuinely many-writer, bus-like surface. The **Decision Log** records the committee's consensus: each SignalApproved event carries full vote provenance. The **Release Log** is written by exactly one component, the **Sequencer**: on each approval it snapshots the stake table (wave per account, from an indexer of on-chain staking state), draws intra-wave slots from the verifiable randomness of §7.1, constructs the release manifest, anchors its Merkle root on-chain, and only then emits time-stamped ReleaseOrder events at each account's scheduled instant. The **Fill Log** records venue acknowledgments appended by stateless executor workers that submit orders using per-account, HSM-guarded credentials. Measured outcomes flow from the Fill Log into the calibration and re-weighting loop of §1.

The single-writer Sequencer is the load-bearing choice. It yields total ordering, exactly-once release semantics (idempotency keyed on account × signal, deduplicated at the venue via client order identifiers), and - decisively - determinism: the manifest is a pure function of the Decision Log, the epoch stake snapshot, and the published randomness, so any external party can recompute it and compare against the committed root. The architecture is thereby an audit property, not merely an implementation: §7.1's receipts prove what was scheduled, and replay determinism proves the schedule itself was the only one the inputs permit. Three operational rules complete the design. Subscriber kill switches are control events with precedence over every release, honored by Sequencer and executors independently. Degradation is intra-wave only: under venue rate limits or partial outages, the randomized slot order determines who waits within a wave, and no failure mode may invert ordering across waves - a wave that cannot be honored in its window is marked in the manifest and visible in receipts rather than silently reordered. Finally, timing requirements are deliberately modest - offsets are tens to hundreds of seconds, so the system needs provable ±100ms adherence with published clock-skew distributions, not microsecond engineering; the fairness claim rests on commitments and venue timestamps, never on the platform's own clock.

Figure 14 — Reference architecture: four ordered logs gated by a deterministic sequencer

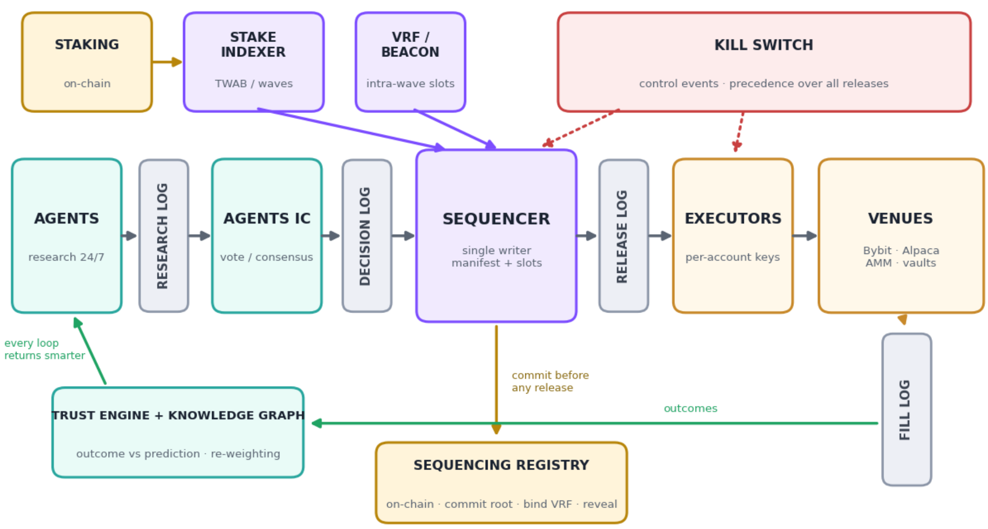

Figure 14. Reference architecture. Agents write to the Research Log; the committee's approvals form the Decision Log; the single-writer Sequencer - fed by the stake indexer and verifiable randomness, and committing to the on-chain registry before any release - emits the Release Log; executors carry orders to venues, whose acknowledgments form the Fill Log; outcomes close the learning loop back to the agents.

5 TWAB staking and manipulation resistance

5.1 Qualification rule

Wave standing is a function of the *time-weighted average balance* of staked LIVE over the qualification window W, valued at the 30-day TWAP of the token:

$$\mathrm{TWAB}_i(t) \;=\; \tfrac{1}{W}\int_{t-W}^{t} B_i(s)\,ds \;\approx\; \tfrac{1}{E}\sum_{e=1}^{E} B_i(t_e), \qquad V_i = \mathrm{TWAB}_i \cdot P_{30d}^{\mathrm{TWAP}}$$

$$w(i) = \min\{k : V_i \geq \theta_k\}, \quad \theta_{1..4} = \$10\text{M}, \$1\text{M}, \$100\text{K}, \$10\text{K}\ (W = 180, 120, 90, 30\text{d});\ , \quad \theta_5 = \$1\text{K}\ (7\text{d})\ \text{for platform membership}$$

Thresholds sit at approximately 1% of each band's intended account size - $10K for ≈$1M accounts up to $10M for ≈$1B accounts - so the stake scales with the capital whose execution it prioritizes - and given the platform's crypto-native audience, stake is expected to rotate in from BTC and other majors, not fiat on-ramps. The discrete form is computed from per-epoch snapshots taken at unpredictable times within each epoch (committee-of-oracles randomness), with W = 7, 30, 90, 120, and 180 days for the Wave-5 through Wave-1 bands respectively - commitment scales with priority. New stake ramps into full weight linearly over the window; unstaking reduces weight immediately and asymmetrically (Figure 4). Balances must be staked in the non-transferable escrow contract - wallet balances, borrowed balances, and LP positions do not qualify.

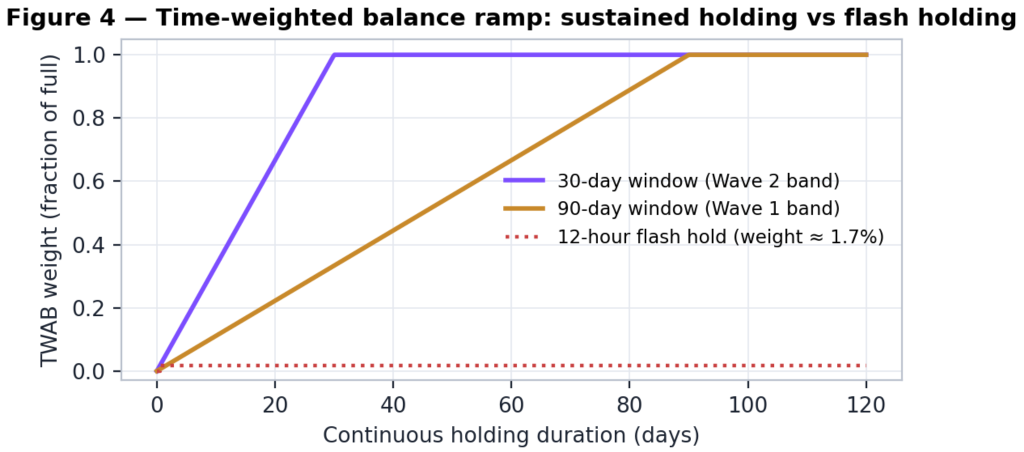

Figure 4. TWAB weight as a function of continuous holding duration for the 30-day and 90-day windows. A 12-hour flash hold (dotted) attains ≈1.7% of full weight - economically useless for band qualification.

5.2 Cost of manipulation

Consider an attacker who wishes to attain full band weight while holding for only h ■ W days per window (e.g., borrowing around snapshots). Because weight is proportional to the time integral of balance, attaining full weight with hold time h requires scaling the balance by W/h, so borrow cost is:

$$C_{\text{attack}}(h) \;=\; \theta_w \cdot \frac{W}{h} \cdot r_b \cdot h \;=\; \theta_w W r_b \;\;\gg\;\; C_{\text{honest}} = \theta_w W r_c, \;\; r_b \gg r_c$$

The attack cost $\theta_w \cdot W \cdot r_b$ is invariant to h - shortening the hold proportionally raises the required balance, and grows linearly in the window W (Figure 5). At θ_1 = $10M, W = 180 days, and a 12% annualized borrow rate, spoofing Wave-1 standing costs ≈ $592,000 per cycle in pure carry on borrowed size 180× threshold, i.e., ≈ $1.8B of borrowed LIVE, over four times the token's fully diluted valuation and therefore physically impossible; even the Wave-4 band (θ_4 = $10K, W = 30 days) requires ≈ $300K borrowed, before liquidity, slippage on acquiring the borrow, and snapshot-timing risk, which dominate in practice given LIVE market depth. Honest qualification at the risk-free-plus-premium carry rate is strictly cheaper for any participant who values the standing at all. Randomized snapshot timing removes the residual strategy of borrowing only around known snapshot instants.

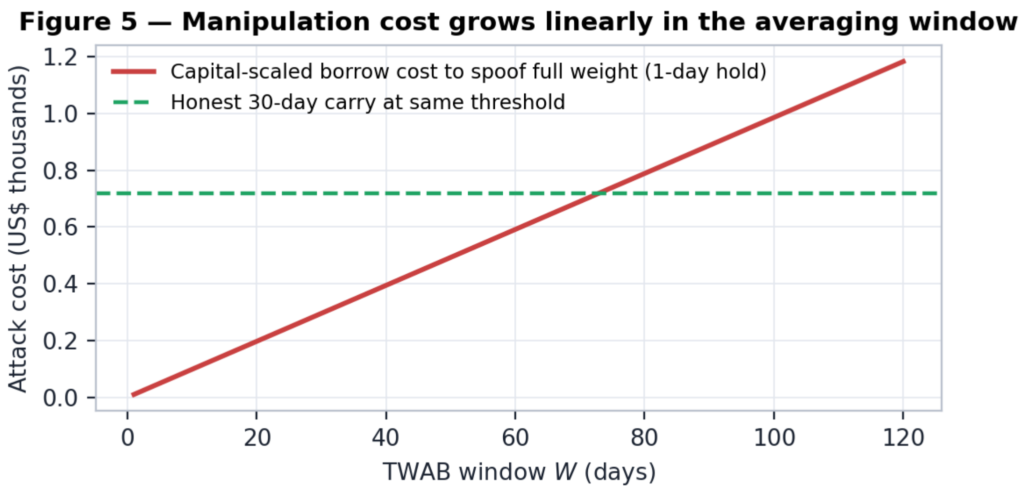

Figure 5. Cost to spoof full band weight via short-duration holding, as a function of the averaging window, versus honest carry at the same threshold. Manipulation cost is linear in W and invariant to hold duration.

5.3 Reflexivity of USD-denominated thresholds

Thresholds are fixed in USD terms and pegged per epoch to the posted direct-subscription price of each wave (payable in BTC, ETH, or stablecoins). Consequently the required token quantity falls as the token price rises:

$$\theta_w \text{ fixed in USD} \Rightarrow N^{\text{req}}_{\text{PAI}} = \frac{\theta_w}{P^{\text{TWAP}}} \quad \text{(token demand falls as price rises)}$$

This is a deliberate negative-feedback choice: token-denominated thresholds would create a reflexive loop (price ↑ → standing value ↑ → demand ↑), which is the signature dynamic of speculative access tokens. USD thresholds make band demand a function of the *service's* value, not the token's momentum, and preserve exact equivalence with the direct rail. The cost is that a falling token price raises the token quantity required - existing stakers are grandfathered within an epoch and receive a top-up grace period of one full window.

6 When is staking rational? The break-even frontier

A risk-averse allocator with capital K delegated to platform agents, trade frequency f (executions per year touching the account), and required annual return $r + \phi_{tok}$ on locked stake (opportunity cost plus a token-volatility risk premium) should upgrade one wave if and only if the expected annual slippage saving exceeds the carry on the required stake:

$$f K \Delta s_w \geq (r + \varphi_{\text{tok}})\,\theta_w \implies K^* = \frac{(r + \varphi_{\text{tok}})\,\theta_w}{f \Delta s_w}$$

Under base-case simulated premia, the four upgrades save Δs ≈ 7, 6, 5, and 4 bps per execution (W5→4, W4→3, W3→2, W2→1) against incremental stakes of $10K, $90K, $900K, and $9M respectively. With f = 30 executions/year and a demanding $r + \phi_{tok}$ = 18% (deliberately conservative to reflect LIVE's volatility), the break-even delegated capital is K* ≈ $86K, $900K, $10.8M, and $135M for the four upgrades - each comfortably below the ≈1%-of-capital account sizes the bands are designed for ($1M, $10M, $100M, $1B), so staking is rational for every band's intended occupant. Upgrades higher in the ladder buy smaller marginal savings on larger marginal stakes, a diminishing-returns property inherited from the concavity of impact (§3.2). Below a band's break-even the direct subscription (BTC/stablecoin) - or simply accepting a later wave - dominates; above it, staking dominates. The frontier is published and recomputed as measured premia update (Figure 6). Note what this framing implies: LIVE staking is an *operating decision about execution costs*, economically analogous to an exchange colocation or direct-feed subscription, and it should be evaluated with the same capital-budgeting arithmetic - not as an investment in the token.

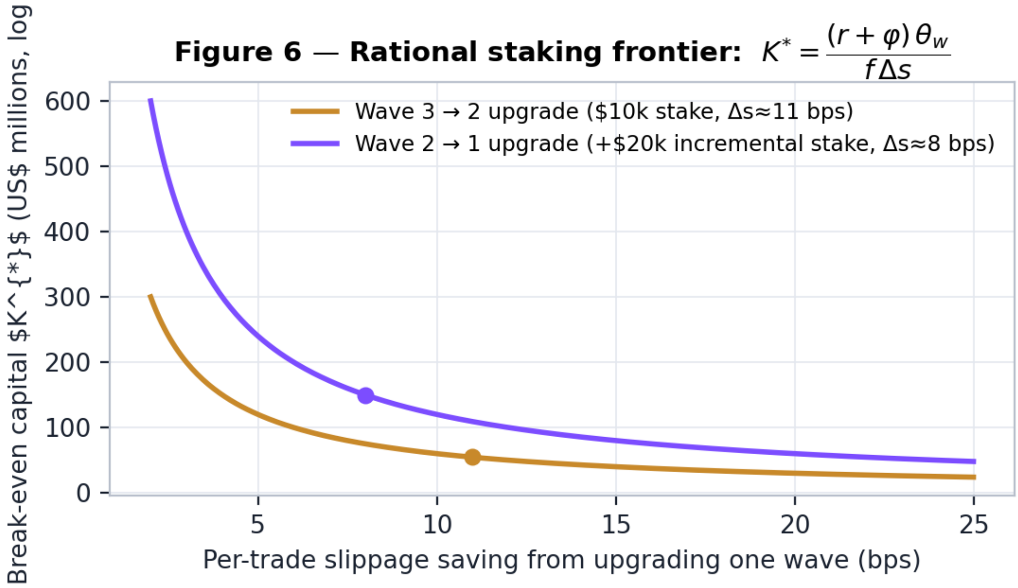

Figure 6. Break-even delegated capital K for each of the four one-wave upgrades (log scale), using the incremental stake each requires (f = 30 executions/yr, r + ϕtok = 18%). Markers show base-case premia; all four sit below their bands' intended account sizes.*

7 Measurement, attestation, and verifiability

Every claim in Sections 3–6 is falsifiable against platform data, by design. For each trade, the platform publishes to its event log and anchors on-chain (Merkle root per epoch): the committee-approval timestamp and decision price P^{dec}; each wave's release timestamp; and each wave's volume-weighted fill. Per-wave implementation shortfall and the inter-wave premium estimator are then:

$$\mathrm{IS}_k = \frac{\bar{P}_k^{\mathrm{fill}} - P^{\mathrm{dec}}}{P^{\mathrm{dec}}} \cdot 10^4 \text{ bps}, \qquad \widehat{\Delta s}_{k\to k-1} = \mathrm{median}(\mathrm{IS}_k) - \mathrm{median}(\mathrm{IS}_{k-1})$$

Rolling 90-day medians of Δs by wave, by agent, and by asset class are published in the transparency feed, and the staking interface displays each user's realized per-trade shortfall against the Wave-1 reference. Three audit properties follow. **Completeness:** fills are reconciled against broker records; a trade absent from the log cannot have occurred through the platform. **Non-repudiation:** epoch Merkle roots are anchored on-chain, so historical premia cannot be retroactively edited. **Falsifiability:** if measured Δs converges to zero, the priced premium is publicly observed to be worthless - the mechanism prices honesty about its own value, and the rational response (documented in advance) is a schedule revision or threshold reduction at the next policy version.

7.1 Queue receipts: proving position without leaking the sequence

Per-wave statistics establish that the schedule is honored on average; subscribers additionally require *individual, online* proof of their own position - without the proof mechanism itself leaking the signal or the queue. The platform uses a commit–execute–reveal protocol. **Commit.** At committee approval, before any order is transmitted, the platform constructs a release manifest: every participating account, its wave, its intra-wave slot (assigned by a verifiable random function, so intra-wave ordering is provably not discretionary), and its scheduled release time. Each entry is hashed with a per-subscriber salt into a Merkle tree, and only the root is anchored to a public chain. The root discloses nothing: not the instrument, not the direction, not cohort membership. **Execute.** Orders release per the manifest; each submission captures the venue's own acknowledgment identifier and timestamp - the ground-truth clock is the exchange's, not the platform's. **Reveal.** After the final wave completes, the manifest, the VRF proof, and the venue acknowledgments are published. Each subscriber receives a **queue receipt**: a Merkle inclusion proof that their (account, wave, slot, scheduled time) tuple was inside the pre-trade root, paired with the venue acknowledgment showing actual release matched schedule within a published clock-skew tolerance. Verification is one comparison: the anchoring block's timestamp precedes the first fill's venue timestamp, therefore the sequence was fixed before anyone traded, and the subscriber's position honored it. A reference verifier is open-source; no trust in the platform's dashboard is required.

On-chain, the protocol is a single minimal registry contract: **commit(root, epoch)** on a fixed cadence (the including block is the tamper-proof 'before' timestamp - the chain is the clock, no oracle required); **bindRandomness(epoch, vrfRequest)** at commit time, so intra-wave slots provably cannot be re-drawn after the fact; and **reveal(epoch, manifestHash, vrfProof, ackDigest)** after the final wave, with the manifest on content-addressed storage. Verification is exposed as an on-chain view function - verify(leaf, proof, epoch) - so any wallet, indexer, or third-party contract can check a subscriber's inclusion against the stored root with no platform infrastructure in the trust path; receipts may optionally be minted as non-transferable tokens, giving allocators a permanent wallet-native record of honored queue position. Per-signal salts are derived from a subscriber master secret, preventing cross-manifest linkage. The two venue regimes then differ in exactly one respect: on the on-chain venues of §4.1 the fills are themselves transactions in blocks after the commitment, so wave monotonicity - including the absence of unmanifested platform-linked flow - is publicly checkable from chain data alone; on CLOB venues the reveal binds venue-signed acknowledgments into the digest, the chain proves the schedule preceded trading, and adherence remains detection under third-party audit as above. Storage cost is one 32-byte root per interval, with all signals in an epoch batched under a second-level tree.

Three leakage vectors are closed by construction. Content leakage: pre-trade, only a hash exists. Cohort leakage: salted leaves prevent any subscriber from inferring who else participates or where they stand. Existence leakage - the subtle one: if roots appeared only when signals exist, the commitment stream would itself be a tradeable signal; the platform therefore publishes commitments on a fixed cadence, inserting indistinguishable dummy roots in idle periods, so the stream carries no information. One limitation is stated rather than obscured: the protocol proves that every manifested order was sequenced correctly; it cannot cryptographically prove the absence of unmanifested flow. Off-chain, operator self-dealing is addressed by detection - drop-copy reconciliation of all platform-associated venue accounts against the manifest, under third-party audit - and by the attested no-proprietary-priority policy; on the on-chain venues of §4.1, execution is public by construction and the residual gap closes.

8 Tokenomics of the LIVE token

8.1 Supply

LIVE has a fixed maximum supply of 10,000,000,000 tokens minted at the token generation event (TGE); the contract exposes no mint function thereafter, so supply is strictly non-increasing. Two mechanisms remove tokens permanently: a fixed burn of 500 LIVE per allocation-slot claim (§4) and the burning of slashed builder bonds (§8.3). At the reference price of US$0.0444 the fully diluted valuation is US$444M; circulating supply at TGE is 2.0B tokens (20%), consisting solely of the liquidity and public tranches.

Figure 7 — LIVE token genesis allocation (max supply 10,000,000,000 · no inflation)

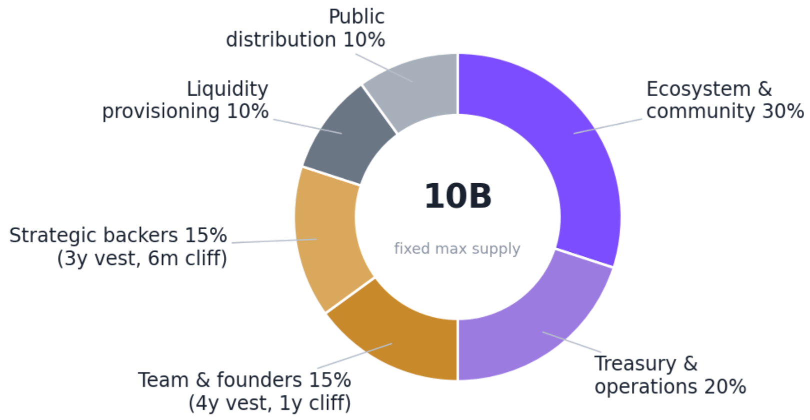

Figure 7. Genesis allocation of the fixed 10B supply. No inflation; no post-TGE minting.

8.2 Allocation and vesting

Vesting is enforced by on-chain timelock contracts published at TGE; no allocation can accelerate its own schedule. Team tokens carry a full 12-month cliff - no team liquidity exists during the platform's first year of measured per-wave premia, aligning insiders with the mechanism's demonstrated (not promised) value.

Table 2. Genesis allocation, vesting, and purpose.

Allocation	Share	Tokens (M)	Vesting	Purpose
Ecosystem & community	30%	3,000	48m linear from TGE	staking incentive programs (non-yield: fee credits, threshold rebates), builder grants, audits, integrations
Treasury & operations	20%	2,000	48m linear from TGE	development, market operations, legal & licensing; quarterly transparency report
Team & founders	15%	1,500	12m cliff + 36m linear	alignment; no unlocked team tokens in year 1
Strategic backers	15%	1,500	6m cliff + 30m linear	capital & distribution partners
Liquidity provisioning	10%	1,000	unlocked at TGE	DEX/CEX liquidity, LP incentives
Public distribution	10%	1,000	unlocked at TGE	broad initial holder base; geo-gated sale

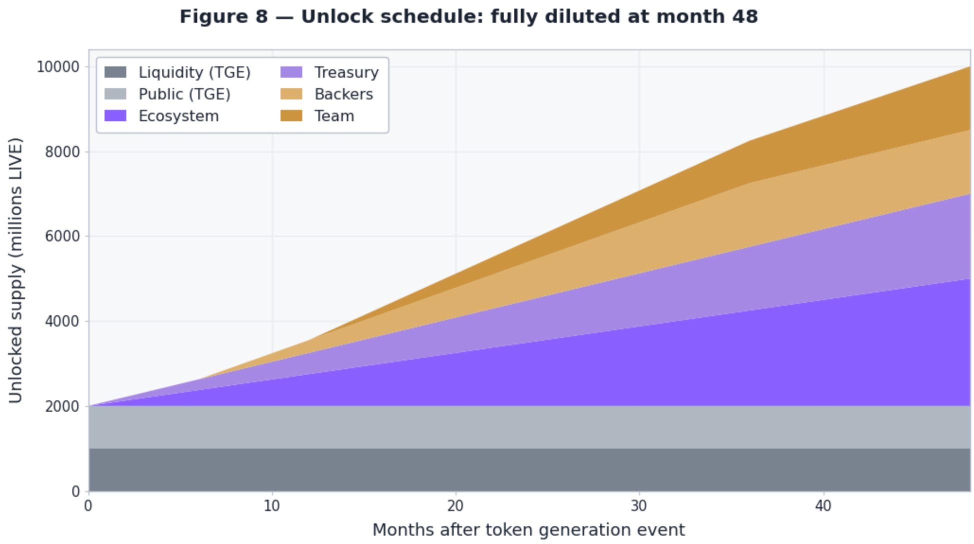

Figure 8. Cumulative unlocked supply by tranche. Checkpoints: 2.0B at TGE, 3.55B at month 12, 5.9B at month 24, fully diluted (10B) at month 48.

8.3 Sinks, and the deliberate absence of value-accrual mechanisms

Net supply change is linked exclusively to platform usage: slot-claim burns (500 LIVE per claim) and slashed builder bonds, which are burned rather than redistributed - redistribution would create a yield expectation. Under low / base / high usage scenarios of 100k / 400k / 1M slot claims per year, cumulative burn over six years is 300M / 1.2B / 3.0B tokens (3% / 12% / 30% of supply; Figure 9). We state explicitly what is **absent**, as binding constraints consistent with §9: no staking yield or emissions to stakers; no buyback funded by platform or subscription revenue; no fee-share, dividend, or treasury distribution to holders. Subscription revenue - in whatever asset it is paid - is never used to purchase LIVE. Demand for the token is designed to derive from exactly one source - the operating value of execution precedence priced in §6 - and supply reduction derives from exactly one source: usage. Any mechanism that would let holders profit from the platform's cash flows is excluded by design, not merely by current policy.

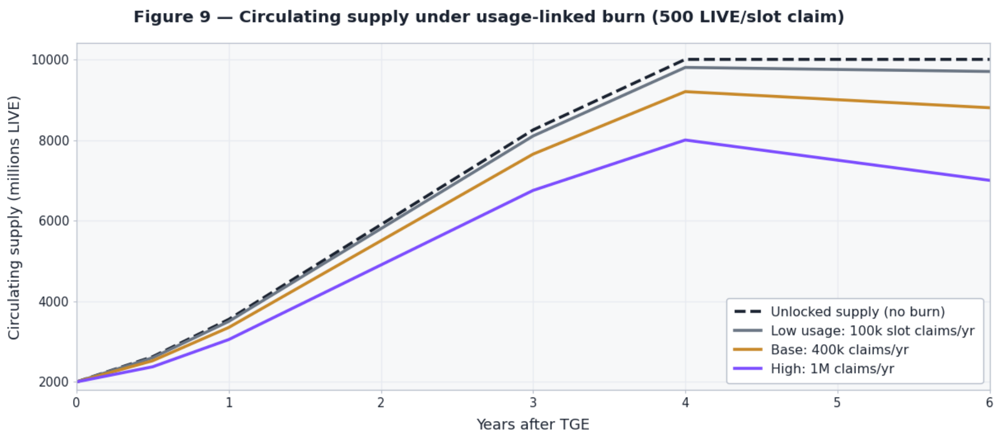

Figure 9. Circulating supply under usage-linked burn scenarios (illustrative; usage rates are assumptions, not forecasts).

8.4 Treasury governance and transparency

The treasury operates under a published mandate (development, market operations, legal and licensing) with quarterly on-chain accounting reports; treasury sales are subject to a disclosed rolling cap (max 1% of circulating supply per 30 days) executed via TWAP to bound market impact - the same impact mathematics of §3 applied to the platform's own conduct. LIVE carries no governance rights over allocation decisions, agent operations, or platform parameters; schedule and threshold revisions are platform policy versions (§7), published with notice periods, not token votes. This is deliberate: governance rights over a trading operation are themselves a step toward the managerial-effort linkage the design excludes.

9 Token design constraints (what LIVE is not)

LIVE confers execution precedence and nothing else. Specifically excluded, as binding design constraints and not merely current policy: no revenue share, dividend, buyback, or fee distribution; no claim on funds under management, agent P&L, or the platform operator; no governance over allocation decisions. Every wave band is equally attainable via a posted direct subscription payable in BTC, ETH, or major stablecoins, and thresholds are pegged to that posted price - the token is one payment rail for a priced service, severable from the service itself. Slot claims burn a fixed LIVE amount, tying net supply change to platform usage rather than speculation. These constraints exist because the mechanism's integrity depends on the token being valued as prepaid execution priority; any cash-flow linkage would convert it into a claim on managerial effort and undermine both its regulatory posture and its economic legibility.

10 Risk factors

10.1 Enumerated paths to zero

Every notice in this document states that LIVE can lose all value. This subsection makes that statement specific: Table 3 enumerates, exhaustively to the authors' knowledge, the scenarios under which the token's value goes to zero or near-zero, the mechanism by which each operates, what the design does about it, and - candidly - the residual risk that remains after mitigation. Scenarios are grouped into five classes: demand collapse, supply-and-market-structure failure, security failure, legal and regulatory action, and integrity failure. A reader evaluating LIVE should price these scenarios, not ignore them; a mechanism whose failure modes cannot be listed is not a mechanism that should hold anyone's capital adjacent.

Table 3. Enumerated zero-value scenarios: mechanism, mitigation, residual risk.

Scenario	Mechanism to zero	Mitigation	Residual
Z1. Premium collapse	Platform flow becomes small vs venue depth; measured inter-wave premia → 0; priority is worthless; demand reduces to membership, which the direct rail also satisfies	§7 publishes the premium estimator continuously; thresholds re-peg down per epoch; schedule revisions can restore separation	HIGH - this is the honest base failure mode of a priority token
Z2. Platform shutdown	LiveAgents ceases operating (funding, key person, strategic exit); the token's sole utility venue disappears	None that preserves token value; timelocked treasury cannot be rug-pulled, but utility ends with the platform	HIGH if platform fails, token value is strictly platform-contingent
Z3. Alpha extinction	Markets adapt (§1); agents stop producing net-positive signals; subscribers leave; no flow, no premium, no demand	Continuous learning loop, open marketplace of third-party agents so the platform is not one strategy; incubation filters	MEDIUM - adaptive markets guarantee individual strategies decay; the bet is on the pipeline, not any agent
Z4. Direct-rail cannibalization	If the direct subscription (BTC/stablecoin) is persistently cheaper than token carry + volatility, rational users pay directly; token demand → speculators only	Epoch re-pegging maintains equivalence, not advantage, for the direct rail; the tension is deliberate, severability requires it	MEDIUM - accepted by design; the direct rail is a legal keel, and it caps token demand
Z5. Ladder unraveling	If Wave-5 expected net alpha turns persistently negative, the bottom tier exits; volume falls; premia shrink for upper tiers; the ladder unravels bottom-up (market-for-lemons dynamic)	Schedule is constrained so W5 expected net alpha stays positive under base parameters (§4); measured W5 alpha is published; policy revision on breach	MEDIUM - the binding design constraint of the whole mechanism

Z6. Reflexive unstake spiral	Token price falls → USD thresholds require more tokens → marginal stakers sell rather than top up → further price fall; grace windows delay but may not stop the cascade	Long TWAB windows slow the loop; top-up grace of one full window; treasury sale cap (1% of circulating / 30d) removes one seller	MEDIUM in a drawdown - USD thresholds are stabilizing for demand but destabilizing for distressed holders
Z7. Unlock overhang	Vesting tranches (backers m6+, team m12+) hit thin liquidity; supply outruns demand growth	On-chain timelocks published at TGE; unlock calendar is fully public (Figure 8); no acceleration possible	MEDIUM - disclosed and datable, but real if adoption lags the calendar
Z8. Liquidity failure	Market makers and LPs withdraw; no venue quotes the token; price is effectively zero regardless of 'value'	10% liquidity tranche at TGE; LP incentives from ecosystem tranche	MEDIUM - liquidity is rented, not owned, in small-cap tokens
Z9. Escrow or contract exploit	Staking escrow drained or bypassed; stolen stake dumped; or band verification corrupted, destroying the mechanism's meaning	No mint function; escrow and timelocks audited pre-TGE; band logic on-chain and verifiable; bug bounty	LOW-MEDIUM, audits reduce, never eliminate, contract risk
Z10. Oracle failure	30d TWAP or snapshot randomness manipulated or broken; band assignment becomes contestable; trust in the mechanism collapses	TWAP over 30 days is expensive to move; committee-of-oracles randomness; §5.2 manipulation-cost bounds	LOW-MEDIUM - thin markets make no oracle invulnerable
Z11. Regulatory action	A regulator (ASIC, SEC, EU under MiCA) characterizes LIVE as a financial product or the platform as an unlicensed adviser/MDA; forced delisting, geo-exit, or unwind	Direct-rail severability (a non-LIVE route to every wave); no cash-flow rights (§8.3, §9); geo-gating incl. AU restriction; jurisdictional structuring; the mechanism is disclosed sequencing among consenting clients, not privileged market access (§4)	HIGH - the single most probable material scenario; mitigations reduce but do not remove characterization risk
Z12. Venue termination	Alpaca or Bybit prohibit sequenced third-party order flow via API, or terminate the platform's access; execution capability collapses	Pluggable broker adapters; multi-venue architecture; venue relationships and ToS review	MEDIUM - venue policy is exogenous and can change without notice
Z13. Integrity scandal	Operator trades ahead of Wave 1, or per-wave statistics are shown to be falsified; trust - the product, evaporates	On-chain Merkle attestation of fills (§7); no-proprietary-priority attestation in the event log; broker-record reconciliation	LOW if attestation holds - and the attestation is precisely what makes falsification detectable
Z14. Systemic crypto event	Exchange, stablecoin, or bridge contagion collapses the venues and rails the token trades on	None within the design's control; multi-venue listing reduces single-point exposure	EXOGENOUS, crypto-systemic risk applies to every token
Z15. MEV extraction (on-chain)	On AMM venues, searchers and block builders capture the inter-wave gap; the premium leaks to MEV rather than accruing to early waves, hollowing out on-chain priority	Private order-flow relays; inter-block (not intra-block) sequencing; per-venue-class premium publication (§7) so on-chain premia are measured, not assumed	MEDIUM on-chain, ordering is auctioned there, not granted; CLOB venues unaffected

Three observations follow from the table. First, the dominant residual risks are Z1 (premium collapse), Z2 (platform failure), and Z11 (regulatory action) - that is, LIVE fails if the service it prices stops being valuable, stops existing, or stops being legal. This is exactly the risk profile the design intends: the token is a claim on the usefulness of execution precedence and on nothing else, so it should die when and only when that usefulness dies. Second, several scenarios are self-announcing: Z1 and Z5 are visible in the published §7 estimators long before they complete, and Z7 is datable from the public unlock calendar - a holder who is surprised by them was not reading the platform's own disclosures. Third, no scenario in the table is mitigated by a promise; every mitigation is a mechanism, a publication, or a structural absence, consistent with the design constraints of §8–§9.

10.2 Ongoing risk factors

Model risk. The square-root law is an empirical regularity with dispersion across assets and regimes; realized premia may differ materially from Table-1 simulations, in either direction. All pre-launch numbers herein are simulations, not measurements.

Premium compression. If platform flow becomes small relative to venue liquidity, inter-wave premia compress toward zero and the staking value proposition weakens; §7's published estimator makes this observable, and thresholds re-peg accordingly.

Token volatility. USD thresholds mean a falling LIVE price increases the token quantity required to maintain standing; stakers bear mark-to-market risk on locked balances that can exceed any execution saving. LIVE can lose all value.

Oracle and snapshot risk. TWAP and snapshot integrity depend on the oracle committee; manipulation of the reference price, while costly, is not impossible in thin markets.

Custody and key risk. The platform never custodies client funds; compromise of the platform cannot move capital out of client accounts, because API permissions are trade-scoped with withdrawal rights disabled at the broker. Residual risks are unauthorized *trading* within an account (mitigated by per-account kill switches and order-size guardrails) and broker/exchange counterparty risk, which clients bear directly at their chosen venue.

Regulatory risk. A token whose holding improves execution outcomes on a discretionary trading platform may be characterized as a financial product in some jurisdictions. Acquisition is geo-gated, availability to Australian residents is restricted pending licensing review, and the direct-rail equivalence is maintained specifically to preserve severability. Prospective participants should obtain independent advice.

Adverse-selection critique. A skeptic may object that Wave 3 is structurally exploited. The design answers are: (i) published per-wave premia let Wave-3 users price their tier accurately; (ii) the schedule is constrained so Wave-3 expected net alpha remains positive under base parameters; (iii) the direct rail (BTC/stablecoin payable) prices the upgrade for anyone whose capital clears the §6 frontier. The mechanism does not remove the queue; it removes the *information asymmetry about the queue*.

11 Conclusion

Sequential execution of shared signals imposes a real, measurable, and unavoidable ordering cost. LIVE converts that cost from an undisclosed transfer into a priced, manipulation-resistant, and continuously audited mechanism: precedence is earned by sustained time-weighted stake or purchased at an equivalent posted price in BTC, ETH, or stablecoins; the value of precedence is measured on every trade and published; and the token is structurally prevented from becoming anything other than prepaid priority. For the risk-averse allocator, the relevant question is not whether the token appreciates but whether the measured inter-wave premium, multiplied by delegated capital and trade frequency, exceeds the carry on the required stake - an arithmetic this paper makes explicit and the platform makes verifiable.

References

Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*, 3(2), 5–39.

Almgren, R., Thum, C., Hauptmann, E., & Li, H. (2005). Direct estimation of equity market impact. *Risk*, 18(7), 58–62.

Bouchaud, J.-P., Farmer, J. D., & Lillo, F. (2009). How markets slowly digest changes in supply and demand. In *Handbook of Financial Markets: Dynamics and Evolution*. Elsevier.

Donier, J., & Bonart, J. (2015). A million metaorder analysis of market impact on the Bitcoin. *Market Microstructure and Liquidity*, 1(2).

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.

Angeris, G., & Chitra, T. (2020). Improved price oracles: Constant function market makers. *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*.

Gatheral, J. (2010). No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7), 749–759.

Harris, L. (2003). *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford University Press.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.

Perold, A. F. (1988). The implementation shortfall: Paper versus reality. *Journal of Portfolio Management*, 14(3), 4–9.

Torre, N., & Ferrari, M. (1997). *The Market Impact Model*. BARRA Research Insights.

Tóth, B., Lempérière, Y., Deremble, C., de Lataillade, J., Kockelkoren, J., & Bouchaud, J.-P. (2011). Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2), 021006.