PROGRAMMABLE MARKETS

The Measurement Layer for Programmable Settlement

A category document

Settlement · sttl.live · July 2026

This document is neither an academic whitepaper nor a token whitepaper. It is the foundational document of a market-infrastructure company. The research program behind it is the scientific layer and will be public; the platform described here is the product layer and is public; the algorithms that connect them are proprietary and are not described in this or any external document.

Every number in this document is verified against the underlying manuscript or a live validation receipt; the verification ledger accompanies this file.

Contents

1 The Problem: Markets Measure Information, Not Execution

Every mature financial market is instrumented for information. Prices, volumes, positions, and news are captured, normalized, redistributed, and priced within milliseconds, because a century of market infrastructure — from the ticker to the terminal — was built on the premise that information is the scarce input to good decisions.

Execution received no such infrastructure. Yet execution — the transformation of an intention into a settled change of ownership — is itself produced under scarcity, and that scarcity determines transaction costs, settlement risk, liquidity provision, ordering advantages, congestion dynamics, the funding of security, and ultimately which claims on economic activity are worth holding. Existing data products describe execution outcomes (a fee was paid, a trade printed) without measuring the economic mechanism that produced them: how much immediacy the venue can supply, what waiting costs, what being first is worth, who bears a demand surge, and whether the rights attached to an asset will survive the institutions that grant them.

Programmable settlement markets make this measurable for the first time. When the rules of a market are code, the market's own operating system — its capacity, its rationing rule, its priority auction, its security budget, its claim boundary — becomes an observable object rather than an institutional black box. The company exists to measure that object, continuously, and to sell the measurement.

We do not analyze blockchain data. We model programmable settlement economies. Blockchains are the first laboratory because they expose every rule and every transaction; they are not the destination. Any market that allocates a scarce, imperfectly storable resource under programmable rules — blockspace today; compute, bandwidth, and power markets as they financialize — presents the same economic object. The CME–Silicon Data compute-futures announcement of 2026 is an early signal that this class of markets is forming; our framework is built for it, and our evidence discipline means we will claim coverage of those markets only when we have measured them.

2 The Paradigm: Markets Whose Rules Are Data

Three properties distinguish programmable markets from every market that preceded them, and together they define a new measurement problem.

The allocation rule is observable. An exchange's matching engine is a trade secret; a settlement protocol's fee mechanism is source code. For the first time, the economist's "market design" is not an assumption but a data field. This converts questions that were previously untestable — does a capacity rule change where scarcity goes? — into identified empirical questions, because the rule change itself is exogenous, dated, and public.

The full population is observable. Ledgers record every transaction, every failed attempt, every priority payment, every governance vote, and every death. There is no survivorship filter unless the analyst introduces one. This permits an evidence standard unavailable elsewhere in finance: population-scale falsification, where a claimed mechanism must survive being tested on *everything*, including the assets that died.

The institution is revisable in observable time. Claims on cash flow can be switched on, revoked, and reinstated by governance processes whose every step is recorded. Institutional risk — the deepest and least measurable risk in traditional asset pricing — becomes a hazard rate one can estimate.

The consequence is a new layer of financial infrastructure: between the raw ledger (commoditized) and the investment decision (the client's), there is a missing layer that converts observable market operating systems into economic state — what is scarce, where scarcity clears, who receives the value, and how durable the claim on that value is. That layer is our product.

3 The Science: Fourteen Discoveries

The company does not begin with a hypothesis about programmable markets. It begins with a completed first-generation research program — fourteen papers, referred to here as A through N — containing formal primitives, machine-verified theory, causally identified mechanisms, population-scale falsifiers, and production-ready measurements. The papers are public (the *why*); the platform operationalizes them (the *what*); the estimation and forecasting machinery in between is proprietary (the *how*).

This section states each discovery, its verified headline evidence, and the platform capability it powers. Every number below is quoted from the underlying manuscript. Where a paper's identification stops, we say so: the claim boundary is part of the product.

3.1 Execution and Time (Papers A, B)

Discovery 1 (Paper A). Immediacy has a supply technology, and market design selects it. Immediate execution is a produced service, and different settlement designs accommodate the same demand impulse in structurally different ways. Measured through mechanism restrictions (the overidentifying restrictions of Ethereum's zero-free-parameter fee controller against Bitcoin's fixed-cap first-price design), the absorption of demand into quantity is approximately 0.00 on Bitcoin and 0.95 on Ethereum, with no overlap across seven studied congestion windows — and within Ethereum itself, the EIP-1559 switch moves absorption from 0.02 to 0.95 on the same chain with the same users: the controller, not the crowd, opens the quantity channel. Bitcoin rations by spot price; Ethereum replenishes through a stored price state. One economic object, two opposite supply technologies. *Capability: the platform classifies every supported venue's execution regime in real time and prices the difference.*

Claim boundary. Mechanism identification, not an instrumented demand curve; magnitudes are specific to the studied windows and designs.

Discovery 2 (Paper B). Waiting has a causal price. Using externally timed variation in block arrival — capacity timing that transaction demand cannot cause — an externally late

Bitcoin block raises the mean clearing fee by 0.307 log points, which under the paper's capacity-rationing reading implies a local demand-for-immediacy elasticity near three. The margin that clears is a cutoff: high-value transactions pay through a capacity shortfall, capacity-heavy transactions wait. The tradable object this defines is the net-immediacy premium $\pi = p_{\text{now}} - \mathbb{E}[p_{\text{next}}]$: what executing now costs relative to the validated expectation of waiting. *Capability: the platform publishes π per venue — "waiting 24h saves x%" — as a decision object, not a chart.*

Claim boundary. A local effect identified at the congested margin under the stated timing design; it is not a universal fee model.

3.2 Rank and Design (Papers C, D)

Discovery 3 (Paper C). Being first has a price, and its incidence is measurable. In serialized settlement, ordering rank is a distinct scarce good. The paper separates the payment for precedence M from the gross first-mover value E it purchases, and measures the paid-away share $\theta = M/E$: the fraction of the prize surrendered to obtain it. In the mature market the rank price tracks the prize with elasticity 0.73 ($t = 29.9$), and θ collapses — from a monthly peak of 0.66 in August 2024 to 0.016 by 2026, roughly forty-fold from the peak and twenty-eight-fold on the annual arc — as the payment channel shifts from competitive tip auctions to posted fees. Market plumbing, not opportunity size, determines who keeps the value of being first. *Capability: priority intelligence — gross prize, rank price, and incidence reported separately, so a client can distinguish value creation from value transfer.*

Claim boundary. Measured on reconstructable opportunities; no claim that all extractable value is observed.

Discovery 4 (Paper D). Design does not remove scarcity; it routes it. Subjecting three live capacity statutes — fixed cap, feedback controller, administered tariff — to a common demand impulse, the response appears on different margins: the fixed cap yields a transient spot-price response of roughly +2% per standard-deviation shock; the controller stores an additional persistent state response of roughly +8%; the tariff routes pressure into rank competition and failed execution. A market that looks cheap on posted fees may simply be paying in a margin the client is not watching. The paper also documents a methodological trap this field routinely falls into: defining shock windows in block time rather than clock time can manufacture the bulk of an apparent response mechanically. *Capability: the pass-through map — for any venue, where a demand surge will surface (price, state, waiting, rank, failure) — and design-consulting mandates built on it.*

Claim boundary. Comparative response to a common shock under predetermined statutes; not a welfare ranking of designs.

3.3 Security as an Economic Flow (Papers E, F)

Discovery 5 (Paper E). Security follows value; it does not lead it. Across claims sharing a security technology, protection allocates in proportion to claim value: the log hashrate ratio and log price ratio of shared-hardware siblings cointegrate at unit elasticity (1.009, s.e. 0.028), and the error correction runs from value to security. Across thirteen shared-hardware coin-splits in five lineages, the focal claim won twelve of thirteen — including against a challenger whose payment rail offered 8–32× the capacity and whose manufactured throughput the market priced at zero. Security spend is derived demand — protection purchased for a claim already selected — and therefore cannot serve as an independent value anchor; and service quality does not select which claim gets protected. *Capability: security intelligence that treats hashrate and stake as protection metrics, with any leading interpretation explicitly tested rather than assumed.*

Discovery 6 (Paper F). Fees have not replaced issuance, and shares can lie about it. Security budgets are financed by a stock charge (issuance dilution) and a flow charge (fees). Around scheduled cuts in the stock charge, the measured replacement ratio — the fee increase per unit of withdrawn issuance finance — is -0.009 for Bitcoin's 2013–2025 transition, and at Ethereum's Merge, which cut issuance-funded security by 86% in one scheduled step, validator fee revenue did not merely fail to compensate: it *fell* 29%. Meanwhile the user-funded *share* of the budget jumped from 4.8% to 23.4% purely through denominator shrinkage — the statistic most often cited as fee-market health can rise mechanically while the budget collapses. The long-run financing of settlement security is an open solvency question the market prices as if it were closed. *Capability: security-budget accounts per venue — issuance, fees, replacement, coverage — as a standing risk product.*

Claim boundary. Replacement is identified around scheduled cuts; the paper does not claim what level of security is adequate.

3.4 Networks, Incidence, and the Platform Boundary (Papers G, H, I, J)

Discovery 7 (Paper G). Networks close before claims matter. On the complete population of 6,190,617 token births — no survivorship filter — two-sided circulation (addresses on both sides of flow) predicts sixty-day economic survival at $z = +5.0$, rising to $z = +6.4$ under flow controls; raw holder counts are null once flow is controlled ($z \approx -0.3$); and the pooled survival rate is 0.77%. Claim strength is uninformative at formation and prices only among mature survivors: a network must close economically before its claim design means anything. *Capability: survival intelligence — the lifecycle state machine (formation → closure → maturity → death) with the circulation hazard live; explicitly not a return signal (see Section 6).*

Discovery 8 (Paper H). Modular production can quintuple activity while starving the upstream claim. When production is modular — rollups producing execution, a base layer selling data availability — the routed share is a measured object: around the scheduled Dencun input-cost reform, L2 fee pass-through fell from 0.69 to 0.34 and the upstream claim's realized incidence $\iota = R/G$ fell from 17.4% to 0.49% — a thirty-five-fold collapse — while gross downstream fees grew roughly fivefold. Growth of an ecosystem is not growth of a claim on it. *Capability: incidence maps — one unit of user activity traced through the* *production stack to whichever claim actually receives it.*

Discovery 9 (Paper I). The platform tax that matters is the one the marginal entrant faces. The marginal participation wedge τ and the realized take ι are different objects — correlated at only 0.38 in the studied panel — and it is the wedge that suppresses open entry ($\beta = -0.049$, $t = -2.08$), while realized take is null-to-positive in the same design. Platforms optimizing (or regulators judging) on realized revenue share are reading the wrong statistic. *Capability: the wedge/take separation as a standing analytic for protocol fee design and for policy work.*

Claim boundary. Theory plus a directional panel illustration; a within-platform wedge experiment remains data-gated.

Discovery 10 (Paper J). Sufficiently monetary platforms should pay their producers. In the committed optimum of the claim-boundary model, the holder-routed share is $a_{\text{commit}} = (1 - \Omega\vartheta)/(1 + \Omega)$, which crosses zero at $\vartheta^* = 1/\Omega$: past a threshold of monetary use, extraction from producers optimally reverses into subsidy of them. The result is machine-verified theory and reframes "value capture" debates: whether a token should tax its ecosystem depends on how monetary the token is. *Capability: a scenario engine for claim-routing design; sold as model-implied analysis, never as a live classification of assets.*

3.5 Claims, Duration, and Fundamentals (Papers K, L, M, N)

Discovery 11 (Paper K). Turning on a claim is priced as coordination, not cash flow. Around governance events that create cash-flow rights, abnormal attention loads on returns ($t = 2.1$; wild-bootstrap $p = 0.05$) while the announced claim magnitude loads approximately zero; verified pre-announced activations — the moment the cash flow actually starts — reprice approximately nothing (mean CAR +0.3%, every $|t| \leq 1.5$); and the pop partially reverses. The flagship exhibit: the 2024 Uniswap fee-switch proposal popped +41% ($t = 6.5$) on 6.9× normal volume, then gave the move back over the following quarter. The market treats claim-openings as focal events first and discounted cash flows a distant second. *Capability: governance-event intelligence that separates the attention channel from the claim channel, so clients do not mistake a coordination pop for durable value.*

Discovery 12 (Paper L). Revocable rights are duration assets. A governance-revisable claim is a two-state process with revocation hazard δ and reopening hazard μ, giving claim duration $D_c = 1/(r - \gamma + \delta + \mu)$: institutional softness enters the denominator as a hazard, not as an ad hoc valuation haircut. Pilot estimates place $\delta = 0.155$ and $\mu = 0.225$ per year — an expected open-route life near 6.4 years — with two independent speed estimates, from two different papers' data, agreeing within 11%. *Capability: claim-duration analytics — the discount for revocability computed from measured transition rates rather than analyst judgment.*

Claim boundary. Pilot-power estimates; per-asset hazard pricing awaits richer transition histories.

Discovery 13 (Paper M). Claims price near par — but only conditional on survival. Among surviving fee-generating assets, value-per-fee loads on the routed share with elasticity 0.83 ($t = 3.1$), within one standard error of the frictionless unit benchmark; pooled across the full population the slope is 0.38, the difference being survival confounding. First claim-openings raise the value-per-fee multiple by +0.43 log points ($t = 4.3$, about +54%) — and the same event class reprices where the token is a pure claim while doing nothing where it is money-like (Ethereum's own EIP-1559 burn, a routed-share change of +0.78, moved the multiple not at all). At the level, open application tokens trade at a median 8× annual fees against a theoretical perpetuity value of 8.5×. Cash-flow rights are real pricing objects — for the assets that live. *Capability: survivor-conditioned comparables; pooled and conditional views never silently merged.*

Discovery 14 (Paper N). Most "fundamentals" are the price in disguise — and this is testable. A candidate fundamental must contain variation outside the price filtration before its cointegration with price means anything. The filtration-escape score formalizes this: Bitcoin's realized cost-basis anchor cointegrates impressively (residual ADF −4.51) yet its escape score collapses from 0.017 to 0.004 as the price basis is enriched — it is price memory — and a five-hundred-random-walk pipeline placebo manufactures both the cointegration and the celebrated factor-of-two valuation band from pure noise, the real statistics sitting at only the 71st–80th percentile of the no-anchor null. Dynamically, the anchor chases the price ($t = 4.6$–5.4) while price shows no detectable reversion. One usage-based anchor passes the test, marginally and with disclosed fragility. *Capability: the anchor admission gate — no metric enters our valuation or risk products labeled "fundamental" without passing escape and placebo tests. This discovery is also aimed at our competitors' dashboards.*

4 The Platform: What Exists Today

Credibility in this category is destroyed by describing aspirations as products. This section is a strict inventory.

4.1 Live

- **Settlement-state feed, six venues** (Ethereum, Bitcoin, Solana, Arbitrum, Base, Optimism): daily state estimation in paper-native objects — fee state and percentile, capacity semantics correct per venue design (a controller's deviation-from-target is a state; a fixed cap's fullness is a design constant), regime classification against trailing-year thresholds, and the net-immediacy premium π where a validated forecast exists.
- **Validated forecasts, with receipts.** Next-24h fee-level forecasts ship only with an out-of-sample validation receipt against the best naïve baseline: Ethereum +13.2% vs. AR5 (Diebold–Mariano $t = 16.4$, 3,078 out-of-sample days); Bitcoin +6.7% ($t = 13.3$, 4,824 days). Solana's forecast was *refused* by the same gate (−60.8% vs. persistence, $t = -3.1$) and the refusal is published on the product surface. The gate is the product.
- **An append-only benchmark ledger.** Every forecast is written, hash-chained, before resolution and scored after. The ledger is the accumulating public track record; it cannot be edited, only extended.

- **An institutional application** on the feed: a unified state-graph canvas (each venue as a measured two-state Markov graph with empirical hazards and dwell times; token survival as an absorbing-state graph), pre-trade settlement TCA across venues, and a non-custodial read-only wallet monitor with post-trade gas TCA against the state benchmark. The app proposes; the client's wallet signs; we never touch keys.

4.2 Near roadmap (build order, not promises)

Solana rank-price decomposition (Discovery 3 applied to live leader revenue); the L2 dual price (execution fee and blob input cost as separate objects, per Discovery 8); failure-rate state for tariff venues (Discovery 4); chain-health tier from the security-budget accounts (Discoveries 5–6); public hosting of feed, app, and ledger.

4.3 Deliberately absent

Custody, execution, order routing as agent, return prediction of any kind, and any forecast without a validation receipt.

5 The Moat: Five Layers

Publishing the papers does not enable replication of the platform, for the same reason MSCI's factor papers do not enable replication of MSCI. The moat is layered, and only the first layer is public.

1. **Published theory.** Papers A–N: the primitives, the identification designs, the falsifiers. Public by design — publication is what creates the category and the credibility. The papers explain *why*.
2. **Ontology.** The operational definitions that make fourteen papers one system: which fee statistic is the cutoff, what counts as a claim route, when a network is economically dead, how a venue's design maps to its valid state variables. Summarized publicly; maintained and versioned internally.
3. **Data engineering.** Historical reconstruction that cannot be rebuilt quickly: population-scale token histories, curve-account conservation logic, entity and channel resolution, source-reconciled panels spanning 2016–2026, and the accumulated catalogue of source pathologies (every provider lies differently).
4. **Models.** The estimators, features, and forecasting machinery. Never disclosed, never published, never in a client deliverable. The papers are sufficient to understand the primitives and insufficient to reproduce the models.
5. **Decision products.** Feeds, TCA, state graphs, incidence maps, duration analytics, the anchor admission gate, and the append-only ledger whose age itself is a moat — a track record cannot be backfilled.

6 What We Will Not Claim

This section is unusual in a commercial document. It is our strongest sales asset.

Before building the platform we ran the opposite experiment: an exhaustive, pre-registered search for trading strategies in the thesis's information set — roughly fourteen strategy families under roughly nineteen frozen verdict rules, spanning survival signals, event drift, factor portfolios, maker-side timing, directional state-riding, volatility forecasting, at every grain from 15 minutes to monthly and widths up to 361 assets. The result: no deployable information edge. The one genuine anomaly found (2019–2021 fee-regime riding) was real, replicated, and already arbitraged away. Conditions *forecasting* shows real, growing skill; *return* prediction does not.

The consequences are structural, not rhetorical:

- **We publish no return predictions, ever.** The nulls are our compliance document: the vendor provably trades no information edge against its clients.
- **Forecasts ship only with receipts.** A validation gate stands between every model and the product surface; refusals are published (Solana is live proof).
- **Survival intelligence is not investment advice.** Our own population evidence shows the survival signal has no tradable price content — we say so on the product.
- **Claim boundaries travel with numbers.** Where identification stops, the product says so; associations are never dressed as causes.

7 Generalization: The Structural-Equivalence Test

The company does not expand into industries; it expands into new realizations of the same structure. The framework applies wherever (i) a scarce, imperfectly storable resource flow is capacity-constrained, $x_t \leq K_t$; (ii) allocation follows programmable, observable rules; (iii) settlement updates persistent claims; and (iv) participants compete on time, price, or rank. Under those conditions the objects of Papers A–N — absorption, waiting price, rank incidence, pass-through, security funding, claim duration — admit direct reinterpretation.

Blockspace satisfies all four today and is where the evidence lives. Compute, bandwidth, cloud capacity, and electricity are approaching them as they financialize — the 2026 CME–Silicon Data compute-futures listing is exactly such a market forming. We cite these as motivation, not coverage: each new market enters the platform when we have measured it under the same evidence discipline, and not before.

8 Sequencing

The order of operations is deliberate, and each step gates the next.

1. **Publish the science.** The A–N wave goes to public deposit; Layer 1 of the moat activates on publication, not before.

2. **Host the measurement.** The feed, the application, and the append-only ledger go public; the track record begins accruing in a form that cannot be backdated.
3. **One design partner.** The anchor engagement is a protocol fee-design or execution-cost mandate built on the pass-through and incidence machinery (Discoveries 4, 8, 9) — a deployment, in the Palantir sense, not a dashboard sale.
4. **Standards.** Only after the ledger has age and the papers have citations do we formalize benchmarks and indices. A benchmark no one can audit is marketing; ours will be an append-only file anyone can verify.

The one-sentence version of this document: *programmable markets made the operating system of markets observable; we built the science that reads it, the discipline that refuses to overclaim it, and the platform that sells it.*