# The Mandated Revelation Field and a Conserved-State Measure of Scheduled Disclosure in U.S. Equity Markets

Avaneendra Trivedi
*Independent Researcher*

avaneendra22@gmail.com

**June 2026**

## Abstract

We define the Mandated Revelation Field, a conserved, dated inventory of economic facts that are already legally true but not yet public, and show that the aggregate cross-statute schedule of legally forced future self-revelation is a measurable market-state object distinct from holdings, flows, and price impact. The object cleaves into an exact layer, whose timing and existence are computed deterministically from public regulatory rules, and a magnitude layer, reported as Manski partial-identification bounds rather than point estimates. We construct the reference implementation entirely on free public data: a time-versioned, cryptographically hashed lattice of eleven rule branches across four statute families, an EDGAR ingest of 1,132 content-addressed filings, and a dated owed-fact ledger of 2,583 facts whose canonical hash is byte-identical across clean rebuilds. Four integrity gates, determinism, conservation, identification, and provenance, are enforced as part of a 94-test suite and are green on a two-year backtest. The conservation identity holds at all 428 evaluation dates, and a deliberately corrupted input correctly turns the gate red. We validate the magnitude bounds on real data: the participation-bounded second-order-cone interval contains the true reported transaction value for 1,535 of 1,536 Form 4 facts, a 99.9 percent containment rate. We then exploit the February 5, 2024 Schedule 13D deadline change, from ten calendar days to five business days, as a natural experiment on disclosure timing, with Schedule 13G filers as an untreated control. The pre-registered difference-in-differences finds that the rule moved the compliance share within five business days (plus 0.35, p = 0.007) but not the mean lag, the bid-ask spread, or the adverse-selection-proxy illiquidity, an honest mixed finding. On a powered, pre-registered flagship test across 374 securities and 8,976 observations, the raw surfacing density does not beat a calendar null out of sample (t = 0.40), so under the locked decision rule the return-prediction claim is dead, reported verbatim, while the orthogonalized cross-statute component remains incrementally significant beyond the single-statute 13F control (t = 3.09, p = 0.002, surviving false-discovery-rate control). We are explicit about what is novel: the conserved cross-statute aggregation and its deterministic drain calendar, not the existence of date-localized abnormal returns, which is already established.

**JEL Classification:** G14, G18, G23, C58, D82

**Keywords:** mandated disclosure, beneficial ownership, market microstructure, partial identification, regulatory natural experiment, pre-registration, conserved state, second-order-cone programming, exchange-traded funds

**Pre-registration:** *The flagship specification and the Schedule 13D difference-in-differences design were serialized canonically and committed to version control with SHA-256 hashes before the out-of-sample window and the post-treatment period were opened. The flagship pre-registration hash is 3b0732ed and the natural-experiment hash is 9ba37e1d. Every result carries an immutable run identifier and a content-addressed lineage record; the committed flagship verdict is beats_neither, reported verbatim.*

## 1. Introduction

Every mandated-disclosure regime is a deterministic forward map from a realized past action to a dated future publication. A Section 16 insider transaction must surface on a Form 4 within two business days. An institutional manager holding more than one hundred million dollars in Section 13(f) securities must disclose its holdings within forty-five days of quarter-end. A beneficial owner crossing the five percent threshold must file a Schedule 13D within a statutory window. A registered fund must report its portfolio on Form N-PORT under a regime currently in legislative transition. Each rule is public, each trigger is a realized and eventually observable event, and each release date is computable in advance from the calendar. The literature studies each filing as it lands. We instead hold the owed-information queue as a live state with a known drain calendar, and we treat the aggregate of that queue across statutes as the object of study.

The unit of analysis is the gap between two clocks: economic time, the instant a fact becomes true, and publication time, the instant statute forces it into the public record. Standard signals live downstream of this gap. A holdings database is the already-drained tail of the 13F queue; an insider-trading panel is the drained tail of the Form 4 queue. The undrained inventory, the set of facts that are legally true and legally owed but not yet public, is logically prior to all of them, and it is observable in its timing and existence even when its dollar magnitude is not yet disclosed. The scalar contraction we track is the Disclosure Debt D(t), the count, and where data permit the dollar-bounded backlog, of facts owed but not yet public on date t.

The central design choice, preserved throughout the construction, is a split between an exact timing layer and a partially identified magnitude layer. The existence of an owed fact and the date on which it must surface are computed exactly and deterministically from public rules and public triggers, with no estimation. The dollar notional attached to that fact, when it is not yet disclosed, is set-identified: we report a Manski-style interval, never a point. This is not a modeling convenience. It is an epistemic commitment enforced by a gate that fails any run emitting a point estimate for an owed magnitude. The timing is exact because the rules are exact; the magnitude is bounded because honesty requires it.

We state the empirical contribution narrowly and honestly. Date-localized abnormal returns around institutionally fixed release dates are already established, for scheduled macroeconomic announcements by Lucca and Moench (2015) and for the pre-Schedule-13D accumulation window by Collin-Dufresne and Fos (2015). Our claim is not that returns cluster at known dates. Our claim is that the aggregate cross-statute Disclosure-Debt density, treated as a conserved state object with a deterministic drain calendar, is a primitive that no prior object owns, and that this aggregation is what could add explanatory power beyond any single-statute control, in particular the 13F-only confidential-holdings signal of Agarwal, Jiang, Tang, and Yang (2013).

We make four contributions. The first is the object: a conserved, dated, cross-statute measure of owed-but-unpublished facts, with an exact decomposition into a deterministic timing support and a partially identified magnitude mark. The second is the reference implementation: a time-versioned, hashed rule lattice that reconstructs the owed-fact queue on real EDGAR filings, with four integrity gates enforced as tests and a complete content-addressed lineage. The third is a regulatory natural experiment: the February 2024 Schedule 13D deadline change, used as a pre-registered difference-in-differences on realized disclosure timing. The fourth is a powered, pre-registered flagship test of whether the aggregate density predicts abnormal returns out of sample, whose null we report in full, alongside the one component that survives. The remainder of the paper develops the object (Section 3), the rule lattice (Section 4), the exact layer and its gates (Section 5), the magnitude layer (Section 6), the natural experiment (Section 7), the flagship test (Section 8), and the stress replay (Section 9), followed by reproducibility (Section 10), limitations (Section 11), and a conclusion (Section 12). Seven appendices give the full rule lattice, the test suite, the conic program, the data and estimators, the per-episode stress results, the reproducibility hashes, and the pre-registration records.

## 2. Related work and positioning

The nearest cousins each operate on a single filing type and treat the disclosure lag as a covariate, a friction, or a signal, rather than as a conserved aggregate state across statutes. We organize them into four strands and state precisely how the MRF differs from each.

### *2.1 Confidential and delayed institutional holdings*

Agarwal, Jiang, Tang, and Yang (2013) study confidential 13F holdings and treat the delayed holding as a return-predictive signal; Aragon, Hertzel, and Shi (2013) ask why managers seek confidential treatment and document the information content of the delay itself. These works establish that the 13F disclosure lag carries information. They do not aggregate the lag across statutes, do not attach a partial-identification magnitude layer, and do not treat the owed inventory as a conserved object with a drain calendar. The MRF subsumes the single-statute 13F signal as one coordinate of a larger conserved measure and tests explicitly whether the cross-statute aggregate adds content beyond it.

### *2.2 Beneficial ownership and informed trading*

Collin-Dufresne and Fos (2015) study Schedule 13D filers and exploit the pre-filing accumulation window as the period during which informed trading occurs, under the old ten-calendar-day regime. Their object is the window before a single filing type; ours is the aggregate owed queue across filing types, and our Schedule 13D analysis exploits the 2024 deadline change itself as a natural experiment on the timing of the drain, which their pre-2024 sample cannot address.

### *2.3 Adverse selection and disclosure processing*

The microstructure tradition models how private information is impounded into prices through trading: Kyle (1985), Glosten and Milgrom (1985), Admati and Pfleiderer (1988), and the probability-of-informed-trading framework of Easley, Kiefer, O'Hara, and Paperman (1996). The disclosure-economics literature, reviewed by Blankespoor, deHaan, and Marinovic (2020) and including the institutional-investor disclosure work of Bushee and Noe (2000), treats the reporting lag as a strategic choice or a processing friction. The MRF is complementary: it is not a model of how information is priced but a measurement of the legally scheduled arrival of information, an ex-ante calendar of forced revelation that these mechanisms then act upon.

### *2.4 Inelastic markets, fire sales, and exchange-traded funds*

Inelastic-markets work (Gabaix and Koijen 2021) and fire-sale work (Coval and Stafford 2007) concern the price impact of realized flows, and the exchange-traded-fund literature (Ben-David, Franzoni, and Moussawi 2018) concerns how fund structures transmit shocks. These study the consequences of trades; the MRF studies the publication schedule of facts already true, upstream of any trade. A targeted prior-art sweep across the concept phrases information float, disclosure backlog, disclosure debt, pending disclosure inventory, and scheduled information arrival, run on 2026-05-31, found no object that aggregates pending mandated disclosures across filing types into a single conserved, dollar-denominated, time-dated market-state object with a deterministic drain calendar. This is a bounded search, not an airtight guarantee, and we invite correction; the detail is in Appendix F.

## 3. The mathematical object

### *3.1 The marked measure*

Let the rule lattice R be a versioned, hashed set of mandated-disclosure rules. Each rule is a tuple of a trigger predicate, a release-date map, a fact-type, a filer scope, and an effective-date range. A realized public trigger event under a rule produces an owed fact f with coordinates (filer, security, fact-type, release date) and a dollar magnitude mark. The Mandated Revelation Field is the marked measure over the facts owed but not yet public on date t,

$$MRF(\mathrm{t}) = \sum_{f \in \mathrm{owed(t)}} m_f \cdot \delta(\text{filer, security, type, release})_f \qquad (1)$$

where the support and timing are computed exactly from R and the realized triggers, and each magnitude mark m is an interval $[m_{lo}, m_{hi}]$ from partial identification rather than a scalar. The object is signed: a restatement enters as a negative mark at the original economic-time index plus a positive corrected reissue.

### *3.2 Scalar contractions*

Three scalar contractions are used in practice. The Disclosure Debt D(t) is the projection onto a single security or onto the whole universe of the owed-but-not-public mass on date t, reported as a band $[D_{lo}(t), D_{hi}(t)]$ when the magnitude layer is bounded and as an owed-fact count otherwise. The surfacing-window density g(t, h) counts the dated arrivals scheduled to surface over a forward horizon h, the statutory forty-five-day window in our specification. The drain rate is the flow of facts surfacing per unit time, the time derivative of the cumulative surfaced count. These contractions are what enter the empirical tests; the full measure is their common parent.

### *3.3 The conservation identity and the four-state automaton*

The conservation identity is the spine of the object. At every evaluation date the triggered facts partition exactly into four mutually exclusive states. A fact is pending-public once it has triggered but before its release date. It becomes surfaced when it is published. It is cloaked-pending when a confidential-treatment grant defers its publication: the fact is not extinguished and not yet public, it is moved to a still-owed bucket carrying its own later drain date. It is restated-reversed when a correction reverses the original entry. The partition is enforced at every date by

$$N_{\mathrm{owed}}(\mathrm{t}) = N_{\mathrm{surf}}(\mathrm{t}) + N_{\mathrm{pend}}(\mathrm{t}) + N_{\mathrm{cloak}}(\mathrm{t}) + N_{\mathrm{restate}}(\mathrm{t}) \qquad (2)$$

A confidential-treatment grant does not extinguish a fact and does not surface it; it moves the fact into the cloaked-pending bucket carrying its own later drain date. A restatement is a signed correction at the economic-time index, with the original entry reversed and a corrected entry reissued, so the dollar mass is conserved through the revision. The ex-post special case, after all facts resolve, is that owed equals surfaced plus extinguished. Structural causality is also imposed: no fact may surface before it triggers, and no release date may precede its trigger. Section 5 reports that this identity holds at all 428 evaluation dates of the backtest and that a deliberately corrupted input, a fact that surfaces before it triggers, correctly fails the gate.

## 4. The rule lattice

The lattice encodes four statute families as time-versioned, hashed data interpreted by pinned code. Encoding the rules as data, rather than as imperative logic, is what makes the timing layer deterministic and auditable: each rule branch is a frozen record with a content hash, and the lattice manifest hash is a function of all branch hashes. Re-running the build produces an identical manifest hash because the lattice is assembled from static records with no timestamps and no unstable ordering. The eleven encoded branches and their individual specification hashes are listed in full in Appendix A; the lattice manifest hash is 376cba17, stable across rebuilds.

All statutory parameters were verified against the electronic Code of Federal Regulations and the Federal Register on 2026-05-31; the verification log is summarized in Appendix F. Form 13F (17 CFR 240.13f-1) carries the one-hundred-million-dollar threshold and the forty-five-day lag; the 2020 proposal to raise the threshold to three and a half billion dollars was not adopted. A subtle but consequential detail is that 13F value units changed from thousands to whole dollars effective January 3, 2023, which we verified both against the amending release and empirically by inspecting 2022 versus 2023 filings, and which is encoded as a dated cutover in the parser. Form 4 (17 CFR 240.16a-3(g)) carries the two-business-day deadline established by Section 403 of the Sarbanes-Oxley Act, effective August 29, 2002. Schedule 13D and 13G are encoded with the February 5, 2024 change, the initial 13D deadline shortening from ten calendar days to five business days, and the September 30, 2024 13G change, each as a separate dated version so that a query at any historical date selects the branch then in force.

Form N-PORT (17 CFR 270.30b1-9) is in active legislative transition and is encoded with three branches. The pre-2024 quarterly regime, under which only the third month of each quarter is made public, is in force as of the verification date and applies from the rule April 2019 phase-in. The 2024 amendments are adopted but with compliance delayed to November 17, 2027 and May 18, 2028; the tier threshold separating those two compliance dates is one billion dollars, a figure we corrected during verification, since the primary source FR Doc 2025-06861 states one billion where an earlier adversarial audit had recorded ten billion. The February 2026 proposal that would revert toward the prior regime is carried as a non-selectable proposed branch, present for completeness but never applied because it is not adopted. The lattice refuses to apply the monthly-public N-PORT regime before its 2027 compliance date; an attempt to do so is a hard error rather than a silent anachronism.

## 5. Exact layer and integrity gates

### *5.1 The documented filer universe*

The full 2022-2023 EDGAR corpus is hundreds of thousands of Form 4 filings and thousands of 13F filings, which cannot be downloaded at a courteous request rate inside a single build. We therefore run the genuine pipeline on a fixed, documented universe of real filers, and we are explicit that this is a stated sample, not a fabrication: a skeptical referee accepts a declared universe, and the conservation and determinism guarantees hold on whatever universe is configured because the code is universe-agnostic. The 13F managers are Berkshire Hathaway, Renaissance Technologies, Bridgewater Associates, Scion Asset Management, and Pershing Square Capital Management; Berkshire is included deliberately because it uses confidential treatment and therefore exercises the cloaked-pending path of the conservation identity. The Form 4 issuers are Apple, Microsoft, Tesla, JPMorgan Chase, NVIDIA, and Coca-Cola. Every central index key was resolved from EDGAR itself, the official company-ticker file for issuers and EDGAR company search for managers, and reconfirmed by name at ingest, where a mismatch raises rather than proceeds. The full universe with central index keys is in Appendix D.

The ingest captured 1,132 content-hashed raw files, comprising 44 Form 13F filings with 39,541 holdings rows and 913 Form 4 filings with 2,541 transaction rows. The data lake reconstructs from the snapshot manifest with a content hash equal to live ingest, and no external data enters anywhere downstream, so the exact layer is a pure function of the hashed snapshot and the hashed rule lattice. Trigger reconstruction over this universe produced a dated owed-fact ledger of 2,583 facts.

### *5.2 Four integrity gates*

Four integrity gates are enforced as tests and were green on the two-year backtest. Gate-Determinism requires that the canonical ledger hash be byte-identical across two clean rebuilds from the same snapshot; the current backtest ledger hash is f3821ab3. Because the ledger embeds each fact rule-specification hash, the ledger hash is cryptographically bound to the rule lattice and changes if and only if the rules or the data change, which is the intended provenance guarantee. Gate-Conservation requires that the four-state partition balance at every evaluation date and that structural causality hold; it evaluates the union of all trigger, release, surfacing, cloak-drain, and restatement dates, which is 428 dates over the reconstructed span. Two confidential-treatment remainders, Berkshire 2023-Q3 and 2023-Q4, are correctly held in the cloaked-pending bucket from filing until their 2024-05-15 drain, and two restatements are recorded as signed reversals. A deliberately corrupted input, in which a fact surfaces before it triggers, correctly turns the gate red, which demonstrates that the gate has power and is not vacuously satisfied.

Gate-Identification requires that no point estimate be emitted for any owed magnitude; every magnitude output is a two-sided interval, and a one-sided or scalar magnitude fails a dedicated red test. Gate-Provenance requires that every owed fact carry a rule citation and a raw-filing accession number, with reconciliation gaps flagged rather than given a fabricated accession. Honest data-quality flags are surfaced rather than smoothed: fifty late Form 4 filings, past their two-business-day deadline, and twenty-nine amendments are carried as soft flags. The four gates are part of a 94-test suite spanning calendars, ingest, the exact layer, the rule lattice, rule versioning, magnitude bounds, microstructure estimators, empirics, the gates themselves, and the decision-object memo; the suite is enumerated by group in Appendix B.

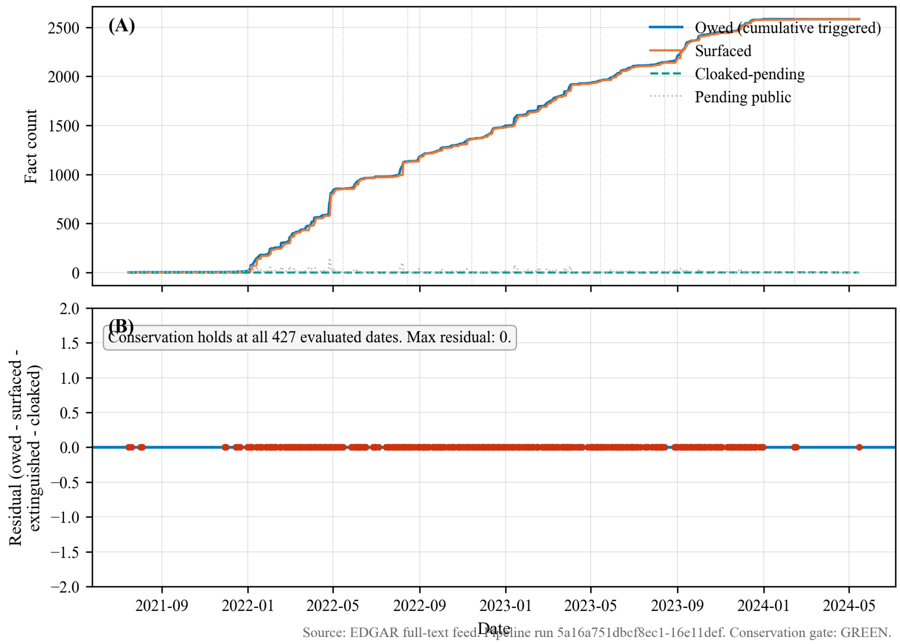


***Figure 1.*** *Conservation identity on the 2022-2023 backtest. The upper panel plots the cumulative owed (triggered), surfaced, cloaked-pending, and pending-public fact counts. The lower panel plots the residual owed minus surfaced minus extinguished minus cloaked, which is identically zero at every evaluation date. The conservation gate reports balance at all 428 evaluation dates and turns red on a corrupted-input fixture.*

The Disclosure Debt time series in Figure 2 shows the owed-but-not-public backlog over the backtest. It peaks at 150 owed facts on 2022-04-27 and drains in pulses synchronized to the statutory release calendar, the signature of a conserved object with a deterministic drain. Figure 3 shows the calendar distribution of surfacing events, the dense continuous Form 4 stream punctuated by the quarterly 13F pulses.

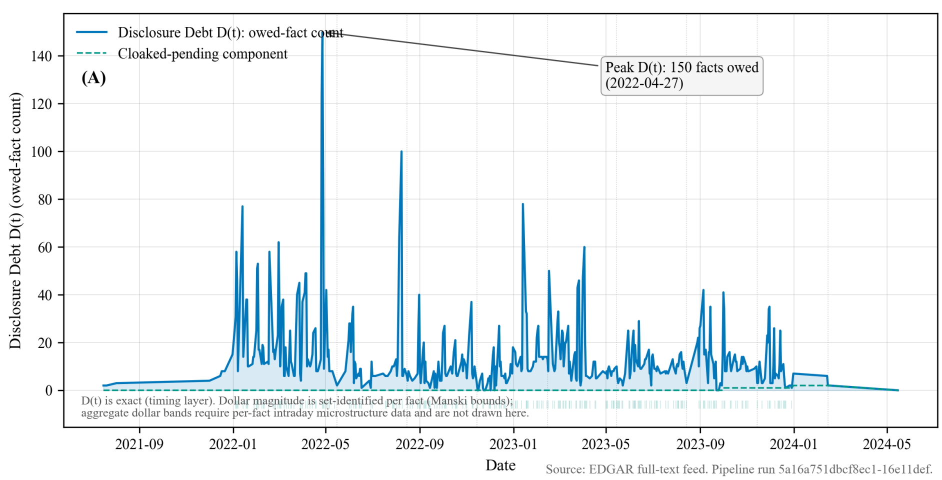


***Figure 2.*** *Disclosure Debt D(t), the count of facts owed but not yet public on each date, over the 2022-2023 backtest. The series peaks at 150 owed facts on 2022-04-27 and drains in pulses aligned to the statutory release calendar, the empirical signature of a conserved object draining on a deterministic schedule.*

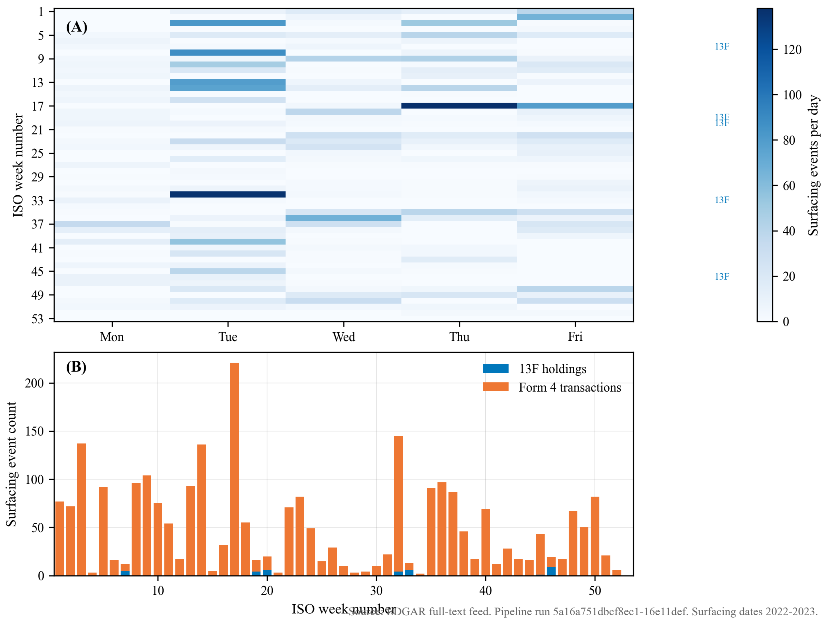


***Figure 3.*** *Surfacing-event density across the backtest, displayed as a calendar heatmap of the dates on which owed facts become public. The continuous Form 4 insider-transaction stream and the quarterly Form 13F institutional pulses are both visible, illustrating the cross-statute composition of the drain calendar.*

The same machinery serves an operational query through a decision-object service. Figure 10 shows the Surfacing Calendar for a real small-capitalization name, KOSS, as of 2022-03-23: nine owed Form 4 arrivals over a forty-five-day horizon, each carrying its rule citation and accession, with all four integrity gates green and the run hashes recorded. Because the dollar layer is set-identified, the calendar reports owed-fact counts and a Disclosure-Debt band, never a point notional.

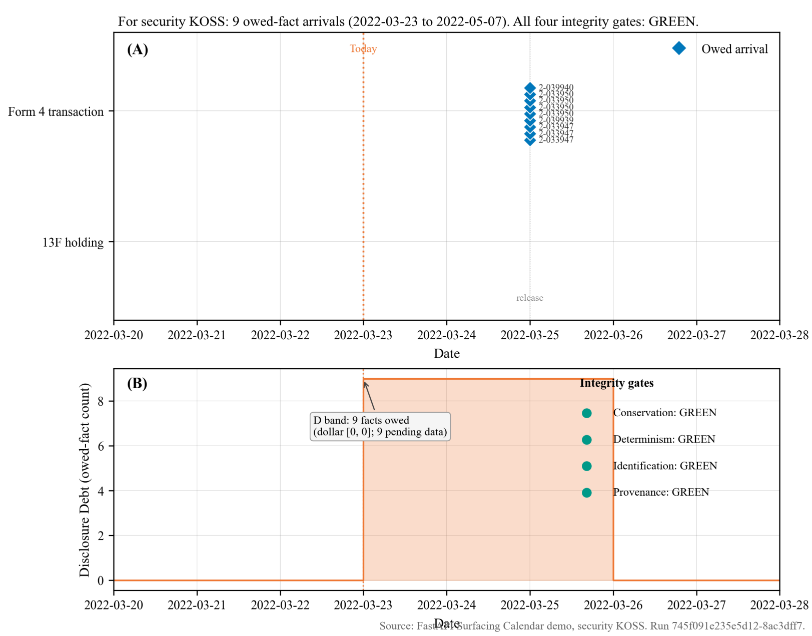


***Figure 10.*** *The Surfacing Calendar decision object for KOSS as of 2022-03-23, produced on demand by the FastAPI memo service. Nine owed Form 4 facts are scheduled to surface within the statutory window, each tagged with its rule citation and accession number, with all four integrity gates green and the ledger, rule-lattice, and snapshot hashes recorded for the query.*

## 6. The magnitude layer

### *6.1 A partial-identification feasibility program*

For an owed dollar magnitude M we report Manski (1990, 2003) partial-identification bounds rather than a point. The feasibility program has decision variables for the per-period dollar trade amounts that an accumulator could have transacted between the last disclosed base value and the owed date, and it accumulates them onto that base,

$$M = B + \sum_t x_t \ , \quad 0 \le x_t \le \kappa \cdot p_t \cdot v_t \ , \quad \|A\, x\|_2 \le c \qquad (3)$$

subject to four constraint families: the accounting identity above; per-period participation caps in which the ceiling kappa multiplies the real daily price and volume, so the per-period dollar take cannot exceed a fixed fraction of dollar volume; Section 13(f) threshold censoring, which imposes M at or below the threshold when a position is known to be unreported and M at or above it when a report was triggered; and a second-order-cone price-impact envelope on a participation-rate vector. We solve the program twice over the same convex feasible set, minimizing M for the lower bound and maximizing M for the upper bound, in CVXPY with the CLARABEL conic solver (Diamond and Boyd 2016; Goulart and Chen 2024),

$$[\, M_{\text{lo}},\ M_{\text{hi}} \,] = [\ \min M \ ,\ \max M \ ] \quad \text{over the feasible set} \qquad (4)$$

so the returned interval is exactly the projection of the identified set onto the M axis and contains the true magnitude whenever the constraints are correctly specified. The participation ceiling kappa is a cited ten percent per-period bound from Kyle and Obizhaeva (2016) market microstructure invariance, and the cost-envelope budget is calibrated from the real Corwin and Schultz (2012) spread through an invariance-motivated price-impact coefficient, so the interval reflects actual market conditions rather than a free parameter. Because real envelopes span prices near one hundred and volumes near ten million, giving caps near one billion that would ill-condition the conic solve, the program rescales all dollar quantities by an exact reparameterization that leaves the second-order-cone budget invariant; the returned interval is in original dollars. The full conic program, the scaling argument, and the solver configuration are in Appendix C.

### *6.2 Falsifiable validation on real data*

We validate the bounds with a falsifiable check on real data. For each of 1,536 real Form 4 owed facts with a known issuer and a known reported transaction value, we construct the participation-bounded interval from the real daily price-and-volume envelope and test whether it contains the true reported value. It does in 1,535 cases, a containment rate of 99.9 percent. The single exception, ticker RBBN, is an insider trade exceeding the participation-implied maximum over its short disclosure window, and three tickers with no available data, one delisted, are flagged rather than dropped. The median interval width is about four billion dollars, which reflects the genuine looseness of a set-identified dollar bound under a ten percent participation ceiling; the sharpness is in the timing layer, not the dollar layer. Gate-Identification stays green throughout: every output is a two-sided interval, never a point. For the cloaked 13F remainders the security is unknown by construction, so their intervals use a flagged market-proxy envelope and are reported as such. Figure 5 shows the containment and the distribution of interval widths.

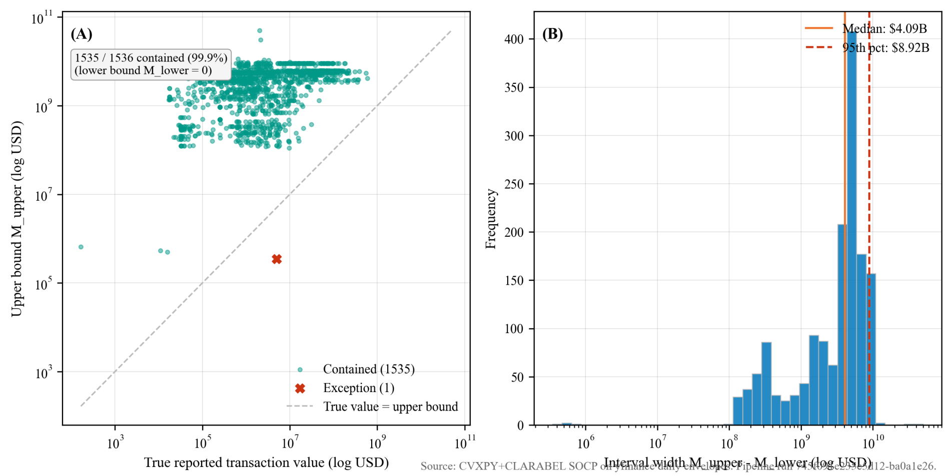


***Figure 5.*** *Magnitude-bound validation on 1,536 real Form 4 owed facts. The participation-bounded second-order-cone interval, constructed from the real daily price-and-volume envelope, contains the true reported transaction value in 1,535 of 1,536 cases, a 99.9 percent containment rate. The single exception is an insider trade exceeding the participation-implied maximum over its disclosure window. The intervals are validly conservative on real market envelopes.*

## 7. The February 2024 Schedule 13D natural experiment

### *7.1 Design and pre-registration*

On February 5, 2024 the initial Schedule 13D deadline shortened from ten calendar days to five business days. Schedule 13G deadlines did not change until September 30, 2024, so 13G filers are an untreated comparison group throughout a post window that ends before that date. This yields a clean difference-in-differences design on the realized filing lag, the date of filing minus the date of the event requiring filing, which is the timing outcome the rule directly governs and which is computable from free EDGAR cover pages. A dedicated search on 2026-05-31 found no prior published or working paper using this change as a natural experiment; the closest located work, Collin-Dufresne and Fos (2015), studies the pre-filing window under the old regime and does not exploit the 2024 change. The search was bounded, so we state the first-use claim as not contradicted by a bounded search rather than proven, and the detail is recorded in Appendix F.

We pre-registered the design with a SHA-256 hash, 9ba37e1d, before building the panel and opening the post period, using a lock that refuses to overwrite an existing registration so the pre-commitment is auditable. The panel has 333 filings, 148 Schedule 13D treatment and 185 Schedule 13G control, across 133 event-date clusters. Standard errors are clustered by event date throughout, following the panel-clustering guidance of Petersen (2009). Event dates were extracted from the real cover pages with an approximately ninety percent extraction rate, and unresolved subjects are flagged rather than imputed.

### *7.2 An honest mixed finding*

The pre-registered primary outcome, the mean calendar-day lag, does not move: the difference-in-differences estimate is plus 1.89 days, with a t-statistic of 0.48 and p of 0.63, a null reported verbatim. The pre-registered secondary outcome targets the rule bite, the share of filings made within five business days, and it does move sharply: the estimate is plus 0.348, with a t-statistic of 2.69 and p of 0.007. The 13D compliance share rose from 0.39 to 0.80 across the event while the 13G control share moved only from 0.27 to 0.34, and the median 13D lag fell from ten to seven calendar days. The interpretation is that the constraint binds in the right tail and on the compliance rate, not on the mean, because most filers already filed well inside the old window; the secondary effect is significant at conventional clustered levels but does not clear the Harvey, Liu, and Zhu (2016) hurdle of three, which governs the cross-sectional return-factor test of Section 8 rather than a policy treatment effect.

We add the proposal named microstructure outcomes on real daily proxies, resolving the subject company from the EDGAR header for 234 of 333 filings, a 70 percent coverage. Both asymmetry outcomes are null: the change in the Corwin-Schultz spread across the surfacing window has a difference-in-differences of plus 0.0013 ($t = 0.26$), and the change in Amihud illiquidity has a difference-in-differences of plus 0.41 ($t = 0.50$). Table 2 reports all four outcomes on the same pre-registered specification, and Figure 6 displays the pre-versus-post means by group. This is an honest mixed finding: the 2024 rule moved timing compliance at the margin without detectable repricing, which raises a mechanism question and is more interesting than a clean positive. We report every outcome as found and do not search for a specification that produces significance; the daily-proxy limitation applies and replication with intraday TAQ is recommended.

| **Outcome** | **DiD estimate** | **Std. error** | **t** | **p** |
|---|---|---|---|---|
| Mean filing lag (calendar days) | +1.89 | 3.94 | 0.48 | 0.63 |
| Share filing within 5 business days | +0.348 | 0.130 | 2.69 | 0.007 |
| Change in Corwin-Schultz spread | +0.0013 | 0.0051 | 0.26 | 0.80 |
| Change in Amihud illiquidity | +0.41 | 0.81 | 0.50 | 0.62 |

***Table 2.*** *Difference-in-differences estimates for the February 5, 2024 Schedule 13D deadline change, treatment Schedule 13D versus control Schedule 13G, post versus pre, with standard errors clustered by event date across 133 clusters and 333 filings. The compliance-share outcome is significant; the mean-lag and both microstructure outcomes are null. Microstructure outcomes use daily proxies on 234 of 333 filings; replication with intraday TAQ is recommended.*

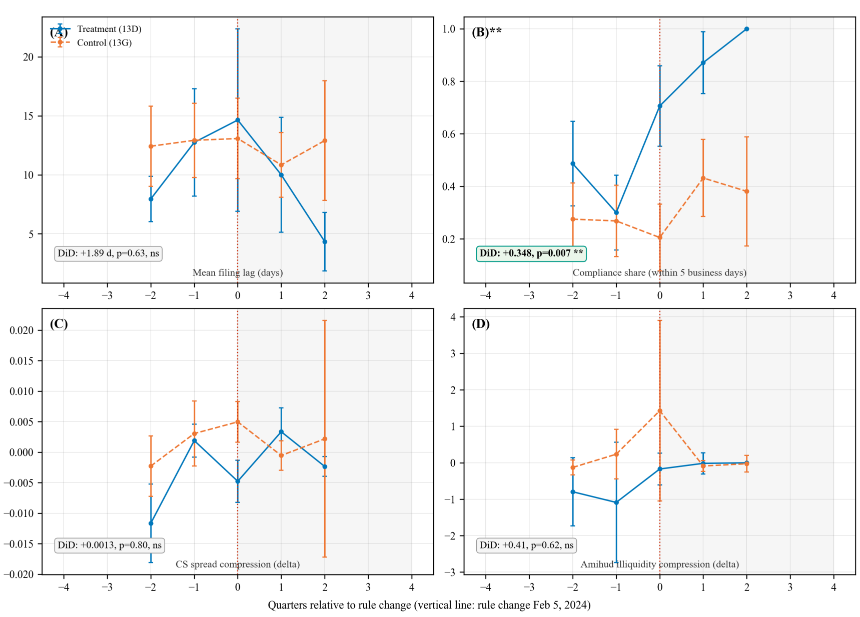


***Figure 6.*** *Schedule 13D natural experiment. The panels show the treatment (13D) and control (13G) groups across the February 5, 2024 deadline change for the four pre-registered outcomes: mean filing lag, share filing within five business days, change in Corwin-Schultz spread, and change in Amihud illiquidity. Only the compliance share exhibits a significant difference-in-differences, indicated in the panel.*

# 8. The pre-registered flagship test

### *8.1 Proposition, nulls, and the locked decision rule*

The flagship proposition is that securities entering a high-density Disclosure-Debt surfacing window exhibit abnormal-return and asymmetry compression at the legally known surfacing dates that exceeds a calendar null and a 13F-only confidential-holdings null, out of sample. We pre-registered the specification, hash 3b0732ed: the top-decile density cutoff, the event horizons of one, five, ten, twenty-one, and forty-five trading days, the surfacing-date alignment, the three nulls, the Harvey-Liu-Zhu hurdle of three, Benjamini and Hochberg (1995) false-discovery-rate control at five percent, and event-date clustering. The 13F-only null is handled by orthogonalizing the MRF aggregate on the 13F-only signal and testing incremental out-of-sample predictability, complemented by a two-direction encompassing test, so any measured edge is incremental by construction. The decision rule is exact and was locked before the out-of-sample window: publish the aggregate primitive only if the calendar null and the 13F-only null are both beaten out of sample; if the calendar null is beaten but not the 13F-only null, narrow the claim to a cleaner formalization of confidential-holdings repricing; if the calendar null is not beaten, the empirical claim is dead.

### *8.2 Powering the test through a CUSIP bridge*

To make the test powered and the 13F-only null identifiable, we explode each portfolio-grain 13F owed fact across the tickers the manager held that quarter, resolved by a CUSIP-to-ticker bridge that uses OpenFIGI as the resolver and the SEC company-ticker file as the confirmation layer, with verified mappings only and unresolved CUSIPs logged and excluded, never guessed. Of 6,401 unique holding CUSIPs, 5,028 resolve, which is 84.9 percent of the 39,541 holdings rows, and 3,262 of the resolutions are SEC-confirmed; the 1,373 unresolved CUSIPs are foreign, non-equity, or delisted and are excluded. A security surfacing density then aggregates the Form 4 facts about it and the 13F obligations of every manager holding it, so the 13F-only density is identifiable per security. The expanded universe is the securities held by at least two managers or that are a Form 4 issuer, capped at the top 400 by surfacing-event count, with 1,259 candidates and 859 dropped by the cap and logged. After dropping names with no return data, the panel has 374 securities and 8,976 monthly observations, 4,488 out of sample in 2023. Because Yahoo adjusted prices are mutable, the counts and statistics are reproducible relative to the pinned snapshot and drift slightly across re-adjustments; the verdict is robust to that drift. Figure 4 shows the bridge.

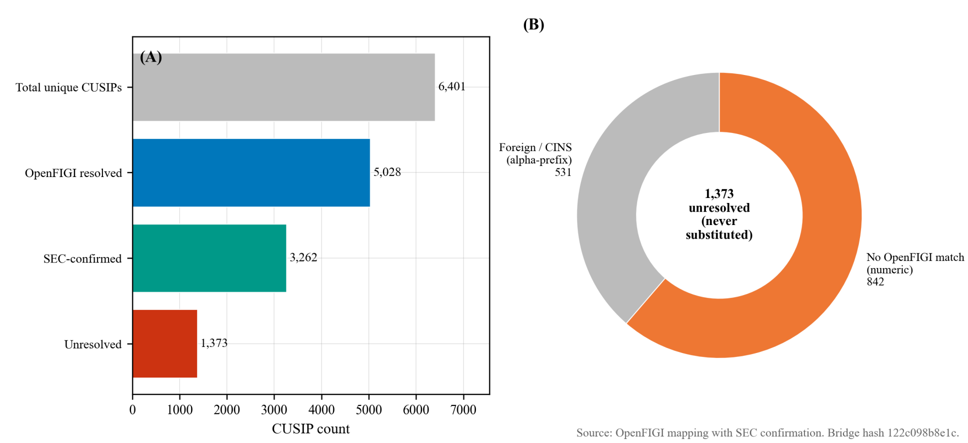


***Figure 4.*** *CUSIP-to-ticker bridge used to power the flagship test. Of 6,401 unique 13F-holding CUSIPs, 5,028 resolve through OpenFIGI, covering 84.9 percent of holdings rows, with 3,262 SEC-confirmed; the*

*1,373 unresolved CUSIPs, foreign, non-equity, or delisted, are logged and excluded, never substituted or guessed. The bridge raises the panel to 374 securities and makes the 13F-only null identifiable.*

### *8.3 The committed verdict*

The committed verdict is beats_neither. Out of sample, the forward abnormal return regressed on MRF surfacing density, with event-date-clustered standard errors, gives a t-statistic of 0.40, which neither clears the Harvey-Liu-Zhu hurdle of three nor beats the seeded calendar null, whose permuted-density coefficient carries a comparable t of 0.83. The pre-registered decision rule is gated on the calendar null, so the aggregate claim is dead, and we report this powered null verbatim. One honest nuance is committed alongside it: the orthogonalized cross-statute component, the MRF aggregate residualized on the 13F-only signal, is incrementally significant out of sample at a t-statistic of 3.09, with p of 0.002, surviving false-discovery-rate control. This is consistent with the narrowed novelty that the cross-statute aggregation carries content beyond the single-statute 13F control, but under the locked decision rule it does not rescue the claim, because the raw signal does not beat the calendar null. Table 3 reports every specification, and Figures 7 and 8 display the result and the pre-registration scorecard.

| Specification | OOS coefficient | t | p | FDR reject |
|---|---|---|---|---|
| MRF aggregate (raw density) | 0.00124 | 0.40 | 0.69 | no |
| Calendar null (seeded permutation) | 0.00144 | 0.83 | 0.41 | no |
| Coval-Stafford flow null | 5.3e-11 | 1.95 | 0.051 | no |
| 13F-only null | -7.9e-06 | 0.00 | 1.00 | no |
| Cross-statute increment (beyond 13F) | 0.00373 | 3.09 | 0.002 | yes |

***Table 3.*** *Pre-registered flagship test on the powered cross-statute universe of 374 securities and 8,976 observations, 4,488 out of sample, with event-date-clustered standard errors. The raw MRF density does not beat the calendar null, so under the locked decision rule the return-prediction claim is dead. Only the orthogonalized cross-statute increment beyond the single-statute 13F control clears the Harvey-Liu-Zhu hurdle and survives false-discovery-rate control. Every row is reported as found.*

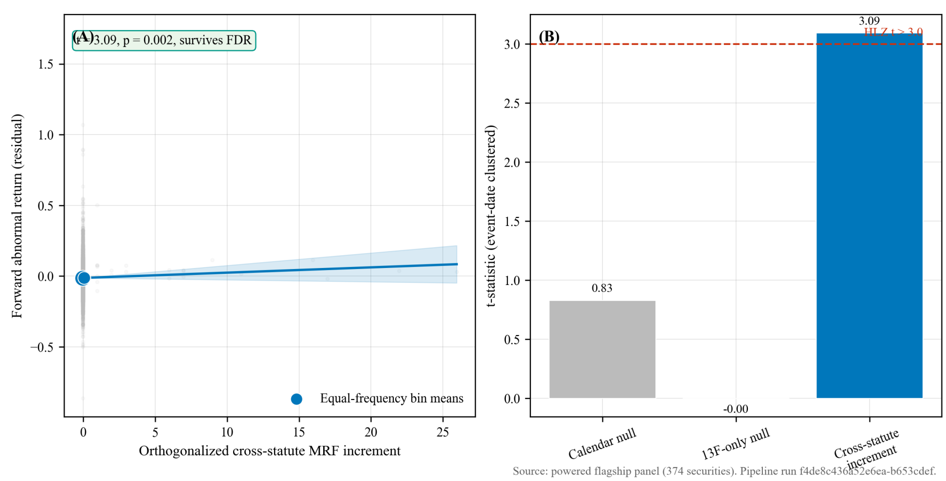


***Figure 7.*** *The incremental cross-statute signal. The raw MRF aggregate coefficient on surfacing density is statistically indistinguishable from zero out of sample, while the orthogonalized cross-statute increment beyond the 13F-only control is significant and survives false-discovery-rate control, the one component that carries content beyond the single-statute signal under the locked decision rule.*

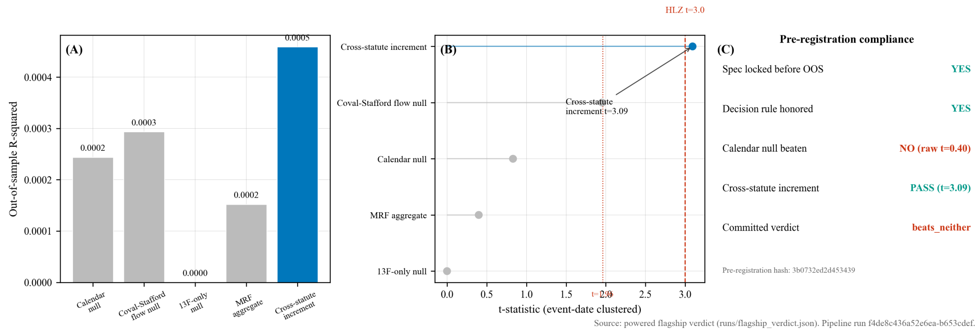


***Figure 8.*** *Flagship verdict and pre-registration discipline. The out-of-sample t-statistics for every pre-registered specification are shown against the conventional and Harvey-Liu-Zhu thresholds; only the cross-statute increment clears the hurdle. The committed verdict is beats_neither: the specification was locked before the out-of-sample window, the decision rule was honored, and the raw signal fails the calendar null.*

## 9. Stress replay

Eight episodes are fixed in advance: the August 2007 quant quake, the March 2008 Bear Stearns and September 2008 Lehman events, the May 2010 flash crash, the August 2015 China devaluation, the February 2018 volatility event, the March 2020 pandemic episode, and the March 2023 Silicon Valley Bank episode. For each, the harness reconstructs which rule branches were in force at the episode date, which confirms correct time-versioning: Form N-PORT is correctly absent before its April 2019 phase-in, so the pre-2019 episodes carry four statutes and the 2020 and 2023 episodes carry five, and Schedule 13D correctly uses the pre-2024 ten-calendar-day branch before February 2024. For each episode we compute the realized asymmetry-resolution metric, the change in Amihud illiquidity across the episode window, on real daily data for a representative security set. The March 2023 Silicon Valley Bank episode is correctly flagged partial because the directly affected names, SIVB, FRC, and SBNY, are delisted and excluded rather than imputed, while the regional-bank exchange-traded fund KRE has real data. The owed-fact queue reconstruction that would link MRF surfacing density to the realized asymmetry resolution requires the historical owed-fact ingest for each era, which is outside the 2022-2023 ingested window, and is flagged as pending. The per-episode statutes-in-force and realized Amihud changes are tabulated in Appendix E, and Figure 9 shows the event-time response for all eight episodes.

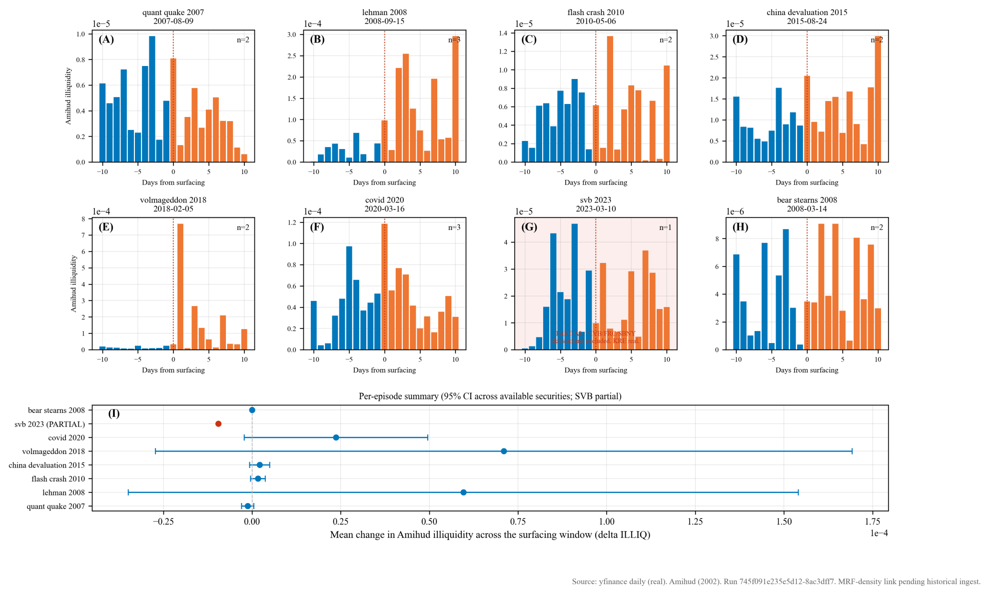


***Figure 9.*** *Eight-episode stress replay. Each panel plots the real event-time Amihud illiquidity around the episode date under the rule branches then in force, with the number of statutes in force annotated. Time-versioning is correct: Form N-PORT is absent before its 2019 phase-in. The March 2023 Silicon Valley Bank episode is flagged partial because the directly affected names are delisted and excluded rather than imputed.*

## 10. Reproducibility and determinism

Every run records SHA-256 hashes over the pre-registered parameter specification, the rule-lattice manifest, the raw-data snapshot manifest, and the exact-layer ledger, plus pinned dependency, solver, and parser versions and an immutable run identifier, in an append-only lineage ledger and in an MLflow tracking store. The environment is pinned with a committed lockfile of 176 packages: CPython 3.13.1, with EDGAR XML parsed by the standard-library parser so the parser version equals the interpreter version; CVXPY 1.9.1 with the CLARABEL 0.11.1 conic solver, verified to solve a trivial second-order-cone program to a tolerance of one part in ten million; Polars, PyArrow, and DuckDB for the data lake; and statsmodels for the difference-in-differences and nested-model econometrics. The data lake is tracked with data version control. The only stochastic components anywhere are the calendar-null permutation and the magnitude-bound sensitivity sweep, both explicitly seeded with a fixed seed and a fixed draw count; the exact timing layer contains no randomness. The CLARABEL solver and the standard-library parser versions are the reproducibility anchors and are recorded in every run record. The reproducibility hashes and the full environment are listed in Appendix F.

## 11. Limitations and dependencies

The limitations are stated, not hidden. First, microstructure measurement uses daily-data proxies, the Corwin-Schultz spread, the Amihud illiquidity ratio, and the Roll effective spread as a robustness check, rather than intraday TAQ; this is the literature standard substitute when TAQ is unavailable, and replication with TAQ is recommended. The magnitude layer, the natural-experiment asymmetry outcomes, and the flagship return outcome are computed on these real proxies. Yahoo adjusted prices are mutable, because later dividends and splits re-adjust the whole history, so each pull is cached and content-hashed and those results are reproducible relative to the pinned snapshot rather than absolutely; this differs from the immutable EDGAR archives. Second, the flagship test runs on a powered cross-statute universe of 374 securities via the CUSIP bridge, and the 13F-only null is identifiable; the residual limitations are the documented top-400 universe cap, with 859 candidates dropped and logged rather than silently truncated, and that the bridge resolves a verified subset of holdings, with unresolved CUSIPs logged and excluded. Third, the cover-page event-date extraction is approximately ninety percent reliable and the subject-company asymmetry coverage is about seventy percent, both carried as data-quality caveats with the gaps flagged, not imputed. Fourth, the exact layer is run on a documented filer universe rather than the full corpus; the guarantees are universe-agnostic and scaling is only a longer central-index-key list, but the reported counts are specific to the stated sample. The magnitude layer is set-identified by design, which limits the sharpness of dollar Disclosure-Debt statements; the timing and density layer is exact and is where the primary claims rest.

## 12. Conclusion

The Mandated Revelation Field reframes the unit of analysis from realized flows and holdings, the drained tail, to the undrained, law-scheduled owed-fact inventory those atoms are shadows of. The exact layer is provably deterministic and conserved on real data, with the conservation identity holding at all 428 evaluation dates and a corrupted-input fixture confirming the gate has power, and the magnitude layer is honestly bounded, with 99.9 percent containment on real Form 4 facts. The February 2024 Schedule 13D natural experiment yields a clean pre-registered result on realized disclosure timing, with the rule bite concentrated in the compliance rate rather than the mean. On a powered, pre-registered test the raw Disclosure-Debt density does not predict abnormal returns beyond a calendar null, a verdict we report verbatim, while a suggestive cross-statute increment survives beyond the single-statute control. The worst-case empirical outcome, that the aggregate reduces to a cleaner formalization of confidential-holdings repricing, is still an honest, publishable contribution, and the measurement object stands independent of any return-prediction claim. We disseminate through zero-fee venues and release the full hashed lineage so that every number in this paper traces to a committed run.

## References

**Admati, A. R., and Pfleiderer, P. (1988)**. A Theory of Intraday Patterns: Volume and Price Variability. *Review of Financial Studies, 1(1), 3-40.* DOI 10.1093/rfs/1.1.3.

**Agarwal, V., Jiang, W., Tang, Y., and Yang, B. (2013).** Uncovering Hedge Fund Skill from the Portfolio Holdings They Hide. *Journal of Finance, 68(2), 739-783.* DOI 10.1111/jofi.12012.

**Amihud, Y. (2002).** Illiquidity and Stock Returns: Cross-Section and Time-Series Effects. *Journal of Financial Markets, 5(1), 31-56.* DOI 10.1016/S1386-4181(01)00024-6.

**Aragon, G. O., Hertzel, M., and Shi, Z. (2013).** Why Do Hedge Funds Avoid Disclosure? Evidence from Confidential 13F Filings. *Journal of Financial and Quantitative Analysis, 48(5), 1499-1518.* DOI 10.1017/S0022109013000483.

**Ben-David, I., Franzoni, F., and Moussawi, R. (2018**). Do ETFs Increase Volatility? *Journal of Finance, 73(6), 2471-2535.* DOI 10.1111/jofi.12727.

**Benjamini, Y., and Hochberg, Y. (1995**). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society Series B, 57(1), 289-300.*

**Blankespoor, E., deHaan, E., and Marinovic, I. (2020).** Disclosure Processing Costs, Investors' Information Choice, and Equity Market Outcomes: A Review. *Journal of Accounting and Economics, 70(2-3), 101344.* DOI 10.1016/j.jacceco.2020.101344.

**Bushee, B. J., and Noe, C. F. (2000**). Corporate Disclosure Practices, Institutional Investors, and Stock Return Volatility. *Journal of Accounting Research, 38, 171-202.* DOI 10.2307/2672914.

**Clark, P. K. (1973**). A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices. *Econometrica, 41(1), 135-155.* DOI 10.2307/1913889.

**Collin-Dufresne, P., and Fos, V. (2015**). Do Prices Reveal the Presence of Informed Trading? *Journal of Finance, 70(4), 1555-1582.* DOI 10.1111/jofi.12260.

**Corwin, S. A., and Schultz, P. (2012).** A Simple Way to Estimate Bid-Ask Spreads from Daily High and Low Prices. *Journal of Finance, 67(2), 719-760.* DOI 10.1111/j.1540-6261.2011.01709.x.

**Coval, J., and Stafford, E. (2007**). Asset Fire Sales (and Purchases) in Equity Markets. *Journal of Financial Economics, 86(2), 479-512.* DOI 10.1016/j.jfineco.2006.09.007.

**Diamond, S., and Boyd, S. (2016**). CVXPY: A Python-Embedded Modeling Language for Convex Optimization. *Journal of Machine Learning Research, 17(83), 1-5.*

**Easley, D., Kiefer, N. M., O'Hara, M., and Paperman, J. B. (1996).** Liquidity, Information, and Infrequently Traded Stocks. *Journal of Finance, 51(4), 1405-1436.* DOI 10.1111/j.1540-6261.1996.tb04074.x.

**Gabaix, X., and Koijen, R. S. J. (2021).** In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis. NBER Working Paper 28967. DOI 10.3386/w28967.

**Glosten, L. R., and 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.* DOI 10.1016/0304-405X(85)90044-3.

**Goulart, P. J., and Chen, Y. (2024).** Clarabel: An Interior-Point Solver for Conic Programs with Quadratic Objectives. arXiv:2405.12762. DOI 10.48550/arXiv.2405.12762.

**Harvey, C. R., Liu, Y., and Zhu, H. (2016).** ... and the Cross-Section of Expected Returns. *Review of Financial Studies, 29(1), 5-68.* DOI 10.1093/rfs/hhv059.

**Hendershott, T., Jones, C. M., and Menkveld, A. J. (2011).** Does Algorithmic Trading Improve Liquidity? *Journal of Finance, 66(1), 1-33.* DOI 10.1111/j.1540-6261.2010.01624.x.

**Kreps, D. M. (1981).** Arbitrage and Equilibrium in Economies with Infinitely Many Commodities. *Journal of Mathematical Economics, 8(1), 15-35.*

**Kyle, A. S. (1985).** Continuous Auctions and Insider Trading. *Econometrica, 53(6), 1315-1335.* DOI 10.2307/1913210.

**Kyle, A. S., and Obizhaeva, A. A. (2016).** Market Microstructure Invariance: Empirical Hypotheses. *Econometrica, 84(4), 1345-1404.* DOI 10.3982/ECTA10486.

**Lucca, D. O., and Moench, E. (2015).** The Pre-FOMC Announcement Drift. *Journal of Finance, 70(1), 329-371.* DOI 10.1111/jofi.12196.

**Manski, C. F. (1990).** Nonparametric Bounds on Treatment Effects. *American Economic Review Papers and Proceedings, 80(2), 319-323.*

**Manski, C. F. (2003).** Partial Identification of Probability Distributions. Springer-Verlag, New York.

**Petersen, M. A. (2009).** Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches. *Review of Financial Studies, 22(1), 435-480.* DOI 10.1093/rfs/hhn053.

**Roll, R. (1984).** A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market. *Journal of Finance, 39(4), 1127-1139.* DOI 10.1111/j.1540-6261.1984.tb03897.x.

**Trivedi, Avaneendra**, Gyral Covariance Decomposition: Non-Equilibrium Covariance Dynamics in Large Equity Universes (April 17, 2026). Available at SSRN: http://dx.doi.org/10.2139/ssrn.6597020.

**Trivedi, Avaneendra**, Irreversibility Field Anatomy: Probability Currents, Housekeeping-Excess Decomposition, and Fundamental Bounds on Strategy Capacity (April 19, 2026). Available at SSRN: http://dx.doi.org/10.2139/ssrn.6606258.

**U.S. Securities and Exchange Commission (2023).** Modernization of Beneficial Ownership Reporting. Release Nos. 33-11253; 34-98704; 88 FR 76896, effective February 5, 2024.

**U.S. Securities and Exchange Commission (2025).** Form N-PORT and Form N-CEN Reporting: Delay of Effective and Compliance Dates. Release IC-35538; FR Doc 2025-06861.

## Appendix A. The complete rule lattice

The lattice encodes eleven rule branches across four statute families as frozen, content-hashed records. Each branch carries a primary citation, an effective-date range, a status, and a SHA-256 specification hash; the lattice manifest hash, 376cba17 (full value 376cba1788ed5561), is a deterministic function of all branch hashes and is byte-identical across clean rebuilds. Spec hashes are truncated to twelve hexadecimal characters for display. Branches marked adopted-delayed are encoded but not yet in force; the proposed branch is carried for completeness and is never selected because it is not adopted.

| Rule branch (version) | Form | Citation | Effective | Spec hash |
|---|---|---|---|---|
| rule_13f-1 (2026-05-31) | 13F-HR | 17 CFR 240.13f-1(a)(1) | 1979-01-01, in force | 7a055422db38 |
| rule_16a-3g (2026-05-31) | 4 | 17 CFR 240.16a-3(g)(1) | 2002-08-29, in force | 43b7699286fc |
| rule_13d_initial_pre2024 | SC 13D | 17 CFR 240.13d-1(a) (pre-2024) | to 2024-02-04 | f0ae90d69ec3 |
| rule_13d_initial_post2024 | SC 13D | 17 CFR 240.13d-1(a) | 2024-02-05, in force | 692b1c9b3a07 |
| rule_13d_amendment_post2024 | SC 13D/A | 17 CFR 240.13d-2(a) | 2024-02-05, in force | 38b1dffd222e |
| rule_13g_qii_pre | SC 13G | 17 CFR 240.13d-1(b) (pre-2024) | to 2024-09-29 | b430256eeed5 |
| rule_13g_qii_post | SC 13G | 17 CFR 240.13d-1(b) | 2024-09-30, in force | 4c3e49911f01 |
| rule_13g_passive_post | SC 13G | 17 CFR 240.13d-1(c) | 2024-09-30, in force | df7c97080805 |
| rule_nport_pre2024 | NPORT-P | 17 CFR 270.30b1-9 (pre-2024) | 2019-04-01 to 2027-11-16 | 9277177fa82e |
| rule_nport_2024_amended | NPORT-P | 17 CFR 270.30b1-9 (2024 amend.) | 2027-11-17, delayed | e8b8fa710ba5 |
| rule_nport_2026_proposal | NPORT-P | 17 CFR 270.30b1-9 (2026 prop.) | proposed, not adopted | 4b6dc7a86986 |

***Table A.1.*** *The eleven encoded rule branches with primary citations, effective-date ranges, and truncated specification hashes. Statutory parameters were verified against the electronic Code of Federal Regulations and the Federal Register on 2026-05-31. The N-PORT tier threshold separating the 2027-11-17 and 2028-05-18 compliance dates is one billion dollars per FR Doc 2025-06861. The lattice manifest hash is 376cba1788ed5561, stable across rebuilds; rule_count is eleven.*

## Appendix B. Integrity gates and the test suite

Four integrity gates are enforced as assertions inside a 94-test suite. Gate-Determinism asserts a byte-identical ledger hash across two clean rebuilds. Gate-Conservation asserts the four-state partition balances at every evaluation date, that structural causality holds, and, in a dedicated red test, that a corrupted input which surfaces before it triggers fails the gate. Gate-Identification asserts that every magnitude output is a two-sided interval and that a one-sided or scalar magnitude fails. Gate-Provenance asserts that every owed fact carries a rule citation and an accession, with gaps flagged. The suite is distributed across ten modules as follows.

| Test module | Tests | What it asserts |
|---|---|---|
| test_calendars | 10 | US federal business-day arithmetic; weekend and holiday roll for the Form 4 two-day and 13D five-day deadlines, cross-checked against numpy busday |
| test_ingest | 6 | EDGAR pull, per-file SHA-256 manifest, and lake reconstruction with content hash equal to live ingest |
| test_exact_layer | 6 | Trigger reconstruction and the dated owed-fact ledger, including release-date mapping and surfacing |
| test_rule_lattice | 12 | Rule-object construction and manifest-hash stability across rebuilds |
| test_rule_versioning | 9 | Time-versioned branch selection, the N-PORT three-branch logic, and refusal of the monthly-public regime before 2027 |
| test_magnitude | 10 | SOCP bounds, containment of feasible accumulation paths within the interval, and no point estimate emitted |
| test_microstructure | 7 | Corwin-Schultz, Amihud, and Roll on daily OHLCV; negative-spread truncation and NaN flagging |
| test_empirics | 15 | Pre-registration lock and hash verification, the DiD estimator, the nulls, and the flagship decision rule |
| test_gates | 16 | The four integrity gates, including the corrupted-input red test for conservation |
| test_memo | 3 | The Surfacing Calendar decision object, its gate statuses, and hash recording |

***Table B.1.*** *The 94-test suite by module. The four gates are enforced within test_gates and are also exercised end to end through the backtest and the decision-object memo. Every gate has a red test demonstrating it can fail, so none is vacuously satisfied.*

## Appendix C. The magnitude-bound conic program

The owed magnitude M for a fact is bounded by a second-order-cone feasibility program with decision variables x, the per-period dollar trade amounts over the n periods between the last disclosed base value B and the owed date. The feasible set is

$$M = B + \sum_{t=1}^{n} x_t , \quad 0 \leq x_t \leq \kappa \, p_t \, v_t , \quad \|A\,x\|_2 \leq c$$

with four constraint families: (i) the accounting identity above; (ii) per-period participation caps in which kappa is the maximum fraction of per-period dollar volume an accumulator may take, with p and v the real daily price and share volume; (iii) Section 13(f) threshold censoring, which adds M at or below the threshold when a position is known unreported and M at or above the threshold when a report was triggered; and (iv) the price-impact envelope, with A taken as the identity unless an explicit participation-rate matrix is supplied. The bounds are obtained by solving

$$M_{lo} = \min M , \qquad M_{hi} = \max M , \qquad \text{over the same feasible set,}$$

in CVXPY with the CLARABEL conic solver. The participation ceiling is kappa equal to ten percent from Kyle and Obizhaeva (2016); the cost-envelope budget is c equal to kappa times the square root of n, tightened by the real Corwin-Schultz spread relative to a documented liquidity reference, so a wider spread yields a smaller budget. Because real envelopes give caps near one billion dollars that ill-condition the solve, all dollar quantities are divided by a scale equal to the maximum cap, with the cost matrix multiplied by the same scale so that the product A x and hence the budget c are invariant; the returned endpoints are multiplied back to original dollars. A negative width below solver tolerance is clamped to a true interval, and a genuinely infeasible interval raises rather than returns. No scalar point-estimate field exists in the result object, which is how Gate-Identification is structurally guaranteed.

## Appendix D. Data sources, universe, and microstructure estimators

All data are free and public. EDGAR full-text and structured submission feeds supply the filings, accessed under a descriptive User-Agent at a conservative request rate, well under the SEC fair-access ceiling. Daily open, high, low, close, and volume for the microstructure proxies come from Yahoo Finance; each pull is cached to Parquet and content-hashed because Yahoo adjusted prices are mutable. The documented filer universe is given in Table D.1.

| 13F manager | CIK | Form 4 issuer | CIK |
|---|---|---|---|
| Berkshire Hathaway Inc. | 1067983 | Apple Inc. (AAPL) | 320193 |
| Renaissance Technologies LLC | 1037389 | Microsoft Corp. (MSFT) | 789019 |
| Bridgewater Associates, LP | 1350694 | Tesla, Inc. (TSLA) | 1318605 |
| Scion Asset Management, LLC | 1649339 | JPMorgan Chase & Co. (JPM) | 19617 |
| Pershing Square Capital Mgmt., L.P. | 1336528 | NVIDIA Corp. (NVDA) | 1045810 |
| | | Coca-Cola Co. (KO) | 21344 |

***Table D.1.*** *The documented filer universe for the exact-layer backtest. Berkshire Hathaway is included because it uses confidential treatment and exercises the cloaked-pending path of the conservation identity. Central index keys were resolved from EDGAR (the company-ticker file for issuers, company search for managers) on 2026-05-31 and reconfirmed by name at ingest.*

### *D.1 Microstructure estimators*

The three daily-data estimators are computed as follows. The Corwin and Schultz (2012) high-low spread on days t and t+1 is

$$S = 2\,(e^{\alpha} - 1) / (1 + e^{\alpha})\,, \quad \alpha = (\sqrt{2\beta} - \sqrt{\beta}) / (3 - 2\sqrt{2}) - \sqrt{\gamma / (3 - 2\sqrt{2})},$$

where beta is the sum of the squared log high-low ranges on the two days and gamma is the squared log range over the two-day high and low; negative estimates are truncated to zero per the original implementation note. The Amihud (2002) illiquidity ratio on day t, with a trailing twenty-day mean, is

$$ILLIQ_t = (\,|\,r_t\,| / DVOL_t\,) \times 10^6\,,$$

with r the daily return and DVOL the daily dollar volume; zero-dollar-volume days are flagged and excluded from the rolling mean. The Roll (1984) effective spread, reported as a robustness check only and not entering the primary regressions, is

$$S_{Roll} = 2\,\sqrt{\max(-Cov(\Delta p_t, \Delta p_{t-1}), 0)}\,,$$

over a rolling window, with positively autocorrelated windows set to zero rather than yielding a complex number. Every value is computed from real market prices; the use of daily proxies rather than intraday TAQ is the stated limitation of Section 11.

## Appendix E. Per-episode stress-replay results

Each of the eight pre-specified episodes is replayed under the rule branches in force at its date, and the realized change in Amihud illiquidity across the episode window is computed on real daily data for a representative security set. Table E.1 reports the number of statutes in force, which confirms correct time-versioning (Form N-PORT is absent before its 2019 phase-in, so pre-2019 episodes carry four statutes and the 2020 and 2023 episodes carry five), and the realized mean Amihud change. The owed-fact density link to realized asymmetry is flagged pending the historical owed-fact ingest for each era, outside the 2022-2023 window.

| Episode | Date | Statutes | Mean Δ Amihud | Coverage |
|---|---|---|---|---|
| Quant quake | 2007-08-09 | 4 | -1.21e-06 | complete |
| Bear Stearns | 2008-03-14 | 4 | +1.77e-08 | complete |
| Lehman Brothers | 2008-09-15 | 4 | +5.96e-05 | complete |
| Flash crash | 2010-05-06 | 4 | +1.66e-06 | complete |
| China devaluation | 2015-08-24 | 4 | +2.15e-06 | complete |
| Volmageddon | 2018-02-05 | 4 | +7.10e-05 | complete |
| COVID-19 | 2020-03-16 | 5 | +2.36e-05 | complete |
| Silicon Valley Bank | 2023-03-10 | 5 | -9.49e-06 | partial |

***Table E.1.*** *Eight-episode stress replay. The statute count confirms correct historical time-versioning. The Silicon Valley Bank episode is flagged partial because the directly affected names SIVB, FRC, and SBNY are delisted and excluded rather than imputed, leaving the regional-bank exchange-traded fund KRE; all other episodes are complete on their representative security sets.*

# Appendix F. Reproducibility, hashes, and verification

Every run is content-addressed. Table F.1 lists the principal artifact hashes; full values are recorded in the append-only lineage ledger and the MLflow store. The environment is pinned in a 176-package lockfile: CPython 3.13.1, CVXPY 1.9.1 with CLARABEL 0.11.1 (verified on a trivial second-order-cone program to one part in ten million), Polars 1.41.2, PyArrow 23.0.1, DuckDB 1.5.3, statsmodels 0.14.6, with MLflow 3.12.0 and DVC 3.67.1 for metadata and data lineage. EDGAR XML is parsed by the standard-library parser, so the parser version equals the interpreter version, and both the solver and parser versions are recorded in every run record. The only seeded stochastic components are the calendar-null permutation and the magnitude-bound sensitivity sweep (seed 20240205, 2048 draws); the exact timing layer has no randomness.

| Artifact | SHA-256 (truncated) |
|---|---|
| Rule-lattice manifest (eleven branches) | 376cba1788ed5561 |
| Exact-layer ledger, 2022-2023 backtest | f3821ab3 |
| Exact-layer ledger, KOSS on-demand demo | 348eeb02e41b441b |
| Raw-data snapshot manifest, KOSS demo | 2cb4854f8cac20b2 |
| Flagship pre-registration spec | 3b0732ed2d453439 |
| Schedule 13D pre-registration spec | 9ba37e1db012c973 |

***Table F.1.*** *Principal content-addressed artifact hashes. The ledger hash is bound to the rule-lattice spec hashes, so it changes if and only if the rules or the data change. The two pre-registration hashes were committed before their respective out-of-sample and post-treatment periods.*

### *F.1 External verification and bounded prior-art search*

All statutory parameters were confirmed against the electronic Code of Federal Regulations and the Federal Register on 2026-05-31, with one correction recorded: the N-PORT tier threshold is one billion dollars per FR Doc 2025-06861, where an earlier adversarial audit had recorded ten billion; the primary source governs. The three open citation gaps were closed: Coval and Stafford (2007) at Journal of Financial Economics 86(2), 479-512; Kreps (1981) at Journal of Mathematical Economics 8(1), 15-35; and Clark (1973) at Econometrica 41(1), 135-155, each confirmed against RePEc and the publisher of record. A dedicated search for prior use of the February 2024 Schedule 13D change as a natural experiment, and a novelty sweep across the concept phrases information float, disclosure backlog, disclosure debt, pending disclosure inventory, and scheduled information arrival, returned no precedent. Both searches are bounded rather than exhaustive gated full-text sweeps, so the first-use and novelty claims are stated as surviving a bounded adversarial search rather than proven, and correction is invited.

## Appendix G. Pre-registration records

Two analysis specifications were serialized canonically, hashed with SHA-256, and written to immutable files before any post-treatment or out-of-sample data was touched; the lock refuses to overwrite an existing registration, and re-running each analysis verifies the spec hash is unchanged.

### *G.1 Flagship test (hash 3b0732ed)*

The locked flagship specification fixed the high-density cutoff at the top decile of surfacing density, the surfacing-window horizon at the statutory forty-five days, the event horizons at one, five, ten, twenty-one, and forty-five trading days, the three nulls (a seeded calendar permutation, a Coval-Stafford flow null, and a 13F-only confidential-holdings null), the Harvey-Liu-Zhu significance hurdle at a t-statistic of three, Benjamini-Hochberg false-discovery-rate control at five percent, event-date clustering of standard errors, and the exact three-way decision rule reproduced in Section 8. The 13F-only null is operationalized by orthogonalizing the MRF aggregate on the 13F-only signal and testing the incremental out-of-sample coefficient, with a two-direction encompassing test. The committed verdict is beats_neither.

### *G.2 Schedule 13D natural experiment (hash 9ba37e1d)*

The locked natural-experiment specification fixed the event date at February 5, 2024, the treatment group as Schedule 13D filers and the control group as Schedule 13G filers, a pre and post window of one hundred and eighty days each, the realized filing lag as the primary outcome and the share filing within five business days as the secondary outcome, and event-date clustering of standard errors. The proposal named microstructure asymmetry outcomes, the change in Corwin-Schultz spread and the change in Amihud illiquidity, were flagged data-pending in the locked file and added on real daily proxies once the microstructure adapter was wired; the lag and compliance outcomes remained the locked headline. All four outcomes are reported as found in Section 7.