

# Capital Trigger Topology: Public-Rule Activation Surfaces, Replay Reproducibility, and Honest Falsification for Installed Capital

**Avaneendra Trivedi**

*Independent Researcher*

avaneendra22@gmail.com

May 2026

## Abstract

We introduce *Capital Trigger Topology* (CTT), a deterministic public-rule activation surface instrument that measures where installed capital becomes mechanically compelled to act, given only public holdings disclosures, codified mandate documents, benchmark definitions, and an information cutoff. The first-domain primitive is the closed-form daily-reset automaton  $\Delta E = \beta \cdot A \cdot (\beta - 1) \cdot r$  for leveraged and inverse exchange-traded funds (Cheng and Madhavan 2009; Avellaneda and Zhang 2010), applied to an append-only registry of 18 certified ProShares products with SHA-256-hashed prospectus and Schedule of Investments evidence. From the automaton we derive sign and materiality boundary surfaces at four pre-registered ADV thresholds, crossing detection, visibility scoring, five integrity gates, and a per-cohort one-page memo with ACCEPT, REROUTE, KILL, or UNIDENTIFIED states. We replay two stress windows, the August 14, 2015 and February 2, 2018 information cutoffs, against ProShares Trust II SEC EDGAR Schedule of Investments XBRL filings (CIK 0001415311). At quarter-end anchors, realized signed VIX-futures position deltas across SVXY, UVXY, VIXM, and VIXY match the CTT-predicted episode-level direction in 8 of 8 ticker-episode pairs. On the daily-resolution diagnostic proxy stack, however, CTT ties the strongest close-cousin daily-reset baseline in both replays and the August 2015 ablation suite emits a hard-kill failure artifact. We commit to honest reporting of the failure: the pre-registered flagship gate of 20 or more Tier 1/Tier 2 daily-crossing-aligned events with baseline beat is uncrossed, and the next valid path is daily official action evidence, not proxy relaxation. All artifacts, run identifiers, and failure states are persisted with parent-hashed registry manifests; 43 tests reproduce both replays byte-for-byte.

**JEL Classification:** G11, G12, G14, G18, G23, C61, C63

**Keywords:** leveraged exchange-traded funds, daily reset, mechanical rebalancing, capital flows, public-rule automata, replay reproducibility, falsification protocol, information cutoff, ProShares, VIX futures, append-only registry, SHA-256 provenance

**Pre-registration:** *All eight ablation thresholds, the matched-pair information-cutoff protocol, the close-cousin baseline construction, the four materiality levels {1%, 2.5%, 5%, 10%}, and the 18-product certified universe were committed prior to empirical execution. All study artifacts, failure states, and registry manifests are persisted with SHA-256 lineage under an append-only registry.*

## 1. Introduction

A research instrument that wants to claim it identifies where installed capital is mechanically compelled to act must do three things in sequence. It must reconstruct the rule that produces the action, not estimate the action statistically. It must apply that rule to an information set that respects a pre-registered cutoff so that the prediction lives in the past tense of the data. And it must show that what the rule says will happen lines up with what subsequently does happen, better than what a strong public substitute baseline would produce on the same data. Anything that satisfies fewer than three of these requirements is something else: a flow predictor, a crowding dashboard, a gamma estimator, or a holdings-overlap network.

This paper introduces *Capital Trigger Topology* (CTT), a public-rule activation topology instrument that consumes only three input classes - public holdings disclosures, codified mandate documents, and benchmark definitions - and produces, for each cohort in a pre-registered universe, the surfaces in state space at which installed capital is mechanically required to act. The first-domain instantiation is the leveraged and inverse exchange-traded fund (ETF) universe, a class of products that has attracted sustained attention in market-microstructure research since their introduction in 2006 (Cheng and Madhavan 2009; Avellaneda and Zhang 2010; Tuzun 2013; Charupat and Miu 2011; Loviscek, Tang, and Xu 2014). For an ETF with prospectus-stated daily leverage factor  $\beta$ , marked-end-of-day net asset value  $A$ , and benchmark return  $r$ , the daily-reset rule has a closed-form delta-exposure solution

$$\Delta E = \beta \cdot A \cdot (\beta - 1) \cdot r, \quad (1)$$

where  $\Delta E$  is the signed reset notional, in units of underlying-benchmark dollars. The closed form follows from arithmetic of the daily-reset rule itself; it is not an estimate. Sign is determined by the sign of  $\beta(\beta - 1)r$  and is invariant to  $A$ . Materiality is the ratio  $|\Delta E| / \text{ADV}$ , where ADV is the average daily volume of the underlying benchmark instrument, and the locus of states  $(r, A, \text{ADV})$  at which materiality crosses a fixed threshold is a one-codimension surface in the three-coordinate state space. We call these surfaces - the sign surface  $r = 0$  and the materiality surfaces at 1%, 2.5%, 5%, and 10% of ADV - the *activation topology* of the cohort. The realized state path crossing one of these surfaces is, by the rule, the moment at which a fund is mechanically required to take signed action. Whether that action moves the underlying - the question Tuzun (2013), Shum, Hejazi, Haryanto, and Rodier (2016), and Ben-David, Franzoni, and Moussawi (2018) study, with answers contested by Ivanov and Lenkey (2018) - is downstream.

CTT is deliberately narrower than the literature on ETF flow effects on prices. Those literatures ask whether ETFs distort underlying prices on average, given how they trade. CTT asks a logically prior question that is silent on prices: where, in state space, are these funds rule-compelled to trade in the first place. The answer requires no statistical model of flow, no estimate of price impact, and no assumption about market efficiency or its absence. It requires the daily-reset rule, the prospectus, the disclosed asset base, and the realized return.

This paper makes four claims. The first is architectural. The second is empirical at quarter-end resolution. The third is empirical at daily resolution and is honestly negative. The fourth is methodological.

**Claim 1 (Architecture).** The closed-form automaton (1) admits a fully deterministic, replayable, append-only-registry-anchored implementation. The system parses public mandate text into a typed PolicyAutomaton schema, hashes every raw source under SHA-256, persists every replay run identifier as the SHA-256 of (configured input hashes, registry-manifest hash), and emits structured failure artifacts when any of five integrity gates is violated. Re-running an identical replay reproduces the run identifier byte-for-byte. The system has 23 sources, 22 cohorts, 24 policies, 1 assumption, 14 schemas, and 1 manifest under registry control, and 43 unit, integration, and regression tests pass. This claim is established constructively in Section 3 and follows the principles of reproducible computational research articulated in Munafò et al. (2017).

**Claim 2 (Quarter-end position evidence).** On the SEC EDGAR Schedule of Investments XBRL data for ProShares Trust II (CIK 0001415311) at the August 14, 2015 and February 2, 2018 information cutoffs, the realized signed VIX-futures position deltas for SVXY, UVXY, VIXM, and VIXY match the episode-level CTT-predicted direction in 8 of 8 ticker-episode pairs. The matched-direction table and SEC accession provenance are reported in Figure 2. This is direct official position evidence at quarter-end anchors. It is not daily reset-execution evidence.

**Claim 3 (Daily-resolution baseline tie and falsification failure).** On the daily-resolution diagnostic proxy stack constructed from official ProShares historical NAV CSV ingestion, CTT ties the strongest close-cousin daily-reset baseline at the nonzero-proxy sign-accuracy level: 40.00% versus 40.00% in August 2015 (25 nonzero events) and 57.89% versus 57.89% in February 2018 (57 nonzero events). The August 2015 ablation suite emits a hard-kill failure artifact: threshold randomization improved the model from 40.00% to 44.00%, which is the wrong direction. We report this as the architecturally required output. The pre-registered flagship gate of 20 or more Tier 1 or Tier 2 daily-crossing-aligned nonzero events with baseline beat is uncrossed.

**Claim 4 (Failure as a first-class output).** Pre-registered hard-kill criteria that fire are persistent JSON artifacts under the registry, not text in the discussion section. The system never reclassifies a diagnostic proxy upward, never silently substitutes a friendly baseline, and never weakens the kill rules when an ablation fires. Section 7 reports each kill rule and its verdict in tabular form, including the rules that currently block CTT from a thesis-grade empirical claim.

### ***Relation to prior literature.***

CTT sits adjacent to several literatures without being reducible to any of them. The leveraged ETF performance literature characterizes return path-dependence and the constant-leverage trap (Cheng and Madhavan 2009; Avellaneda and Zhang 2010; Charupat and Miu 2011; Loviscek, Tang, and Xu 2014); CTT takes the daily-reset rule as a deterministic activation surface generator rather than as a return decay model. The ETF flow and arbitrage literature documents that ETF ownership amplifies non-fundamental volatility (Ben-David, Franzoni, and Moussawi 2018; Shum et al. 2016; Tuzun 2013), with offsetting evidence that capital flows can largely neutralize the rebalancing pressure (Ivanov and Lenkey 2018). CTT does not estimate flow effects on prices and does not require any assumption that they exist. The limits-of-

arbitrage and intermediation literature provides the broader frame in which mechanical capital deployment matters (Shleifer and Vishny 1997; Brunnermeier and Pedersen 2009; He and Krishnamurthy 2013); CTT operationalizes one specific class of mechanical action with a closed-form rule rather than a model. The reproducible-research literature provides the standard CTT enforces by construction (Munafò et al. 2017).

CTT is also a companion to a broader stack of non-equilibrium market structure papers by the author. Reflexivity Kernel Spectroscopy (Trivedi 2026a) measures the flow-to-price transfer operator. Constraint Shadow-Price Tomography (Trivedi 2026b) reconstructs intermediation constraint duals from law-of-one-price wedges. Epistemic Curvature (Trivedi 2026c) measures Riemannian curvature of the market-implied statistical manifold. The Admissible World Deformation Field (Trivedi 2026d) measures cross-market structural inconsistency. Gyral Covariance Decomposition (Trivedi 2026e) measures circulatory covariance perturbations. Irreversibility Field Anatomy (Trivedi 2026f) measures market entropy production. Executable World Resolution (Trivedi 2026g) measures capital-indexed distinguishability of future worlds. CTT is the mechanism layer of this stack: where the others measure state and dynamics, CTT computes the surfaces at which installed capital must act. Figure 8 in Section 8 locates CTT in this taxonomy.

The remainder of the paper is organized as follows. Section 2 develops the theoretical framework and the closed-form automaton. Section 3 describes system construction, the registry, and the replay engine. Section 4 documents data sources, the certified universe, the policy versioning protocol, and the proxy hierarchy. Section 5 specifies the empirical design including the eight pre-registered ablations and the close-cousin baseline. Section 6 presents the matched-direction Tier 1 evidence, the activation surface evidence, the falsification suite outcomes, and a specimen trigger memo. Section 7 reports the pre-registered kill-rule verdict in tabular form. Section 8 discusses limitations and the next required data. Section 9 concludes. Appendices A through C provide derivations, replay reproducibility details, and extended empirical tables.

## 2. Theoretical Framework

### 2.1 The closed-form daily-reset automaton

Fix a leveraged or inverse exchange-traded fund  $i$  with prospectus-stated daily leverage factor  $\beta_i$ . Let  $A_{i,t}$  denote the marked-end-of-day net asset value of fund  $i$  on date  $t$  and let  $r_t$  denote the realized return of the underlying benchmark for the same trading day. Under the daily-reset mandate, the fund's beginning-of-day exposure is leveraged at  $\beta_i A_{i,t-1}$ . After the trading day, the marked exposure is  $\beta_i A_{i,t-1}(1 + r_t)$ , and the marked NAV evolves as  $A_{i,t} = A_{i,t-1}(1 + \beta_i r_t)$ . The new target exposure is  $\beta_i A_{i,t}$ , and the reset action is the difference between target and marked, which simplifies to

$$\Delta E_{i,t} = \beta_i \cdot A_{i,t-1} \cdot (\beta_i - 1) \cdot r_t. \quad (2)$$

Equation (2) is exact for the daily-reset mandate; the derivation in Appendix A.1 follows the standard ETF NAV accounting used by Cheng and Madhavan (2009) and Avellaneda and Zhang (2010). The sign of  $\Delta E$  depends only on the sign of  $\beta(\beta - 1)r$ , which for  $\beta = +2$  is  $\text{sign}(r)$ , for  $\beta = +3$  is  $\text{sign}(r)$ , for  $\beta = -1$  is  $\text{sign}(r)$  (since  $(-1)(-2) = 2 > 0$ ), and for  $\beta = -2$  is  $\text{sign}(r)$  (since  $(-2)(-3) = 6 > 0$ ). All four cases yield the same direction relationship: long-leveraged and inverse funds both buy when the underlying rises and sell when it falls - the sign rule is invariant across the leverage class. This invariance is what underlies the same-direction rebalancing externality that has driven the financial-stability concern in Tuzun (2013).

Two structural identities follow immediately and are testable in code. First, sleeve-split delta: if a fund's NAV is partitioned into  $S$  equal sleeves and each sleeve runs the daily reset independently, the sum of sleeve deltas equals the unsplit fund delta exactly. Second, monotone trigger mass: holding  $\beta$  and  $r$  fixed,  $|\Delta E|$  is monotone non-decreasing in  $A$ . Both identities are encoded as unit tests in the CTT repository (Section 3.3) and protect against the kind of accounting error that the Cushing and Madhavan (2000) closing-auction analysis warns about for end-of-day flow attribution.

### 2.2 Boundary surfaces and crossing detection

Activation surfaces are loci in state space at which the closed-form delta exposure crosses a structural threshold. Two classes of surface are pre-registered.

**Sign surface.** The locus  $r = 0$  separates positive-action from negative-action regions. For any  $\beta \neq 1$ ,  $\text{sign}(\Delta E)$  flips when  $r$  crosses zero. The sign surface is the trivial surface but it is the most important: it determines whether mechanical action is to buy or to sell.

**Materiality surfaces.** Pick a pre-registered materiality threshold  $m \in \{0.01, 0.025, 0.05, 0.10\}$ . The locus on which  $|\Delta E| / \text{ADV} = m$  is the surface  $|\beta A(\beta - 1)r| / \text{ADV} = m$ , equivalently  $|r| = m \cdot \text{ADV} / |\beta A(\beta - 1)|$ . The realized state path crossing of a materiality surface is the moment at which the daily-reset action becomes large enough to plausibly affect the underlying microstructure; the crossing does not establish that it does, only that it could under the rule. Whether the realized impact follows the magnitude of the crossing is the question Shum et al. (2016) and Ivanov and Lenkey (2018) reach divergent answers on.

A crossing event is a pair  $(t_1, t_2)$  of consecutive trading days such that the line segment in state space from  $(r_{t_1}, A_{t_1}, ADV_{t_1})$  to  $(r_{t_2}, A_{t_2}, ADV_{t_2})$  intersects a registered surface, together with the sign of crossing (toward more or less material action) and the index of the surface crossed. Crossing detection is a discrete numerical operation, deterministic given the state path. Figure 1 in Section 6 shows the SVXY pre-cutoff crossings for the August 14, 2015 information cutoff.

### 2.3 Information-cutoff respect and replay determinism

An empirical claim of the form 'CTT identified the action surface that was about to be crossed' requires that the action surface be constructed exclusively from inputs available before the action took place. We enforce this by an information-cutoff protocol. For each replay run we declare a cutoff date  $T_c$ . The replay engine refuses to load any source whose disclosed-as-of timestamp is later than  $T_c$ . If a parser cannot certify the as-of timestamp of a source, the source is excluded from the run and a structured failure artifact is emitted.

Run identifiers are deterministic: a run ID is the SHA-256 of the configured input hashes concatenated with the registry-manifest hash. Re-running an identical replay reproduces the run ID byte-for-byte, which is enforced by 12 regression tests in the repository. Determinism is not a stylistic preference. It is the only mechanism by which a hostile reviewer can ascertain that a reported result reflects the system's behavior on a specific input set, not the result of unlogged tuning. The discipline tracks the reproducibility manifesto of Munafò et al. (2017): every replay artifact in this paper is reproducible from the registry plus the cited run ID.

## 3. System Construction

### 3.1 Public source layer

CTT consumes four classes of public source. Tier 1 sources are SEC EDGAR filings: prospectuses (485APOS), Schedules of Investments embedded in 10-Q and 10-K (XBRL attachments), and material-event 8-K filings disclosing leverage changes. Tier 2 sources are Wayback Machine CDX captures of issuer-hosted daily holdings pages. Tier 3 sources are issuer-published NAV and shares-outstanding endpoints, including the official ProShares historical NAV CSV stream and the current daily holdings file (psdlyhld.csv). Tier 4 sources are third-party diagnostic shares-outstanding archives. CBOE-hosted VIX futures contract files are a triangulation source: 175 contract rows and 147 daily open-interest rows are ingested for cross-checking the VIX-cohort coverage.

Every raw source is hashed under SHA-256 immediately on download. The hash is persisted as a sidecar JSON record alongside the source binary. Subsequent replay runs verify each source's hash against its sidecar before consuming the file. Re-downloads with non-identical hashes raise a content-drift failure artifact and the replay refuses to proceed. This is how the system makes provenance a structural rather than narrative property of every result.

### 3.2 Registry: sources, cohorts, policies, assumptions, manifests

The CTT registry is an append-only directory under version control. It contains six classes of typed JSON record. Sources record the URL, accession number, and SHA-256 hash of every public input. Cohorts record the asset universe and benchmark linkage for a tradable group; the universe contains 22 cohorts. Policies record the typed PolicyAutomaton parsed from a prospectus document, with a  $\beta$  value, an effective date, and a retired-at date if applicable; the universe contains 24 policies. Assumptions record the named assumption a downstream rule depends on; the universe contains 1 assumption (the Wayback sparse-capture protocol). Schemas record JSON-Schema definitions for every typed record class; the universe contains 14 schemas. Manifests record the parent-hashed chain of every prior registry state.

Policy versioning is a non-negotiable correctness property. On 26 February 2018 ProShares filed a material-event 8-K (accession 0001193125-18-059052) reducing SVXY's daily leverage factor from  $-1\times$  to  $-0.5\times$  and reducing UVXY's daily leverage factor from  $+2\times$  to  $+1.5\times$ , both effective 27 February 2018. The CTT registry records four policy versions:

pol\_svxy\_vix\_daily\_leverage\_reset\_v1 ( $\beta = -1$ , retired 2018-02-26),  
pol\_svxy\_vix\_daily\_leverage\_reset\_v2 ( $\beta = -0.5$ , effective 2018-02-27),  
pol\_uvxy\_vix\_daily\_leverage\_reset\_v1 ( $\beta = +2$ , retired 2018-02-26), and  
pol\_uvxy\_vix\_daily\_leverage\_reset\_v2 ( $\beta = +1.5$ , effective 2018-02-27).

The replay engine selects the policy active at the run's information cutoff  $T_c$ , not the latest registered policy. Earlier builds had a loophole in which the most recently filed policy was selected unconditionally; this loophole was closed and is regression-tested.

### 3.3 Replay engine and integrity gates

The replay engine consumes a registry-resolved cohort and a state-path snapshot, evaluates the closed-form automaton (2) over the state path, detects sign and materiality crossings, attaches signed action proxies (Tier 1 EDGAR rows where available, Tier 3 issuer-NAV-derived diagnostic rows otherwise), and computes per-event sign-accuracy against the proxy. Five integrity gates run after replay completes. *PUBLIC\_PROVENANCE* asserts that every used source has a Tier 1 or Tier 2 lineage. *CAUSALITY* asserts that every used source has an as-of timestamp at or before the cutoff  $T_c$ . *REGISTRY* asserts that the manifest hash matches the run-time registry contents. *SCHEMA* asserts that every typed record validates against its schema. *CONFIG* asserts that the run config hash equals the SHA-256 of the canonicalized config JSON. Failure of any gate emits a structured FailureArtifact and the run state is preserved for audit.

### 3.4 Trigger memo and gate language

The endpoint artifact is not a chart and not a single number. It is a one-page institutional memo, generated deterministically from the persisted run artifacts, that reports for each cohort a four-state gate. ACCEPT: the cohort is below all materiality surfaces and capital cost is small. REROUTE: the cohort is above the 1% but below the 5% materiality surface; the next surface at which the gate becomes KILL is reported. KILL:

the cohort is at or above the 5% materiality surface; mechanical action is large and a research thesis that depends on this fund expressing a contrarian view should not commit capital. UNIDENTIFIED: a required upstream artifact is unavailable or fails identification; the gate explicitly refuses to substitute a friendly default. Figure 5 shows a complete memo for the 2 February 2018 cutoff.

## 4. Data and Provenance

### 4.1 Certified universe

The certified public universe contains 18 ProShares leveraged and inverse exchange-traded funds. Direxion's TNA and TZA, while present in the registry as fixture cohorts (cohort\_tna\_daily\_leverage and cohort\_tza\_daily\_leverage), are explicitly excluded from the public empirical claim because the automated prospectus-leverage parser has not yet certified a cutoff-safe Direxion base prospectus. Their exclusion emits an unparsed\_rule failure artifact and is regression-tested. Table 1 lists the 18 certified products, organized by the three benchmark families that have dominated the academic LETF empirical literature (Tuzun 2013; Charupat and Miu 2011; Loviscek, Tang, and Xu 2014).

Equity-index leveraged	Inverse / short	VIX (Trust II)
TQQQ (+3× QQQ)	PSQ (−1× QQQ)	SVXY (−1× / −0.5×)
QLD (+2× QQQ)	SQQQ (−3× QQQ)	UVXY (+2× / +1.5×)
UPRO (+3× SPX)	SPXU (−3× SPX)	VIXY (+1× VIX-fut)
SSO (+2× SPX)	SDS (−2× SPX)	VIXM (+1× mid VIX)
UDOW (+3× DJIA)	SH (−1× SPX)	
DDM (+2× DJIA)	SDOW (−3× DJIA)	
UWM (+2× R2K)	DOG (−1× DJIA)	
	DXD (−2× DJIA)	
	RWM (−1× R2K)	
	TWM (−2× R2K)	
	QID (−2× QQQ)	

**Table 1.** Eighteen-product certified ProShares universe. Each cell lists the ticker and the prospectus-stated daily leverage factor  $\beta$ . SVXY and UVXY include both v1 and v2 policy entries reflecting the 26 February 2018 leverage-change 8-K (accession 0001193125-18-059052). Direxion TNA and TZA are deliberately excluded pending automated prospectus parser certification.

### 4.2 Tier hierarchy and source registry

The proxy hierarchy assigns each candidate evidence row to one of four tiers based on the source class and provenance grade. *Tier 1 (flagship\_edgar\_schedule)*: rows extracted from XBRL Schedule of Investments attachments to ProShares Trust II 10-Q and 10-K filings (CIK 0001415311) with explicit per-contract VIX-futures positions. *Tier 2 (flagship\_wayback)*: rows extracted from Wayback Machine CDX captures of issuer-hosted daily holdings files inside the event window. *Tier 3 (issuer\_official\_diagnostic)*: rows derived from official ProShares historical NAV CSV files and SEC acceleration filings; official, but not crossing-aligned. *Tier 4 (third\_party\_diagnostic)*: rows derived from third-party shares-outstanding archives (sharesoutstandinghistory.com); density-rich, not action evidence.

Table 2 summarizes the realized row counts per tier across both replays. Figure 4 in Section 6 visualizes the same data with the pre-registered flagship gate marked. The tier construction follows the spirit of the constraint hierarchy in Brunnermeier and Pedersen (2009) but applied to evidence quality rather than funding constraints.

Tier	Slug	Rows	Source	Status
1	flagship_edgar_schedule	8	SEC EDGAR XBRL (CIK 0001415311)	direct, hashed
2	flagship_wayback	0	Wayback CDX captures	unavailable in window
3	issuer_official_diagnostic	4	Issuer NAV/shares + acceleration filings	not crossing-aligned
4	third_party_diagnostic	414	sharesoutstandinghistory.com	diagnostic only

**Table 2.** Realized row counts per proxy tier across the August 2015 and February 2018 replays. The pre-registered flagship gate requires 20 or more Tier 1 or Tier 2 daily-crossing-aligned nonzero events with baseline beat. The 8 Tier 1 rows are direct quarter-end position evidence, not daily crossing-aligned events.

### 4.3 Selected SEC accessions

Tier 1 evidence is anchored in five ProShares Trust II SEC filings, all downloaded via official EDGAR endpoints and SHA-256 hashed at the time of ingestion. Table 3 lists the accessions used in the two replays, along with the Credit Suisse XIV 6-K filing relevant to the 5 February 2018 inverse-VIX unwind.

Filing	Accession number	Used by
2015 Q2 10-Q	0001193125-15-284853	Aug 2015 (period anchor)
2015 Q3 10-Q	0001193125-15-371929	Aug 2015 (Tier 1 deltas)
2017 Q3 10-Q	0001193125-17-338612	Feb 2018 (period anchor)
2017 FY 10-K	0001193125-18-067452	Feb 2018 (period anchor)
2018 Q1 10-Q	0001193125-18-158945	Feb 2018 (Tier 1 deltas)
SVXY/UVXY leverage 8-K	0001193125-18-059052	policy versioning
Credit Suisse XIV 6-K	0000950103-18-001572	Feb 2018 (XIV closing IV)

**Table 3.** SEC EDGAR accession numbers for the public filings used in the August 2015 and February 2018 CTT replays. CIK 0001415311 (ProShares Trust II) covers the leveraged and inverse VIX-futures cohort; CIK 0001053092 (Credit Suisse) covers the XIV note. Each accession is downloaded via the official EDGAR endpoint and SHA-256 hashed at ingestion.

## 5. Empirical Design

### 5.1 Stress-window replay protocol

The empirical design replays two pre-registered stress windows. The August 2015 replay sets the information cutoff at the close of trading on Friday, 14 August 2015. The reference shock is the People's Bank of China RMB devaluation announcement of 11 August 2015 and the subsequent SPX flash decline of 24 August 2015; the cutoff intentionally precedes the flash decline, so any CTT prediction made at the cutoff is genuinely out-of-sample with respect to the realized stress event. The February 2018 replay sets the information cutoff at the close of trading on Friday, 2 February 2018. The reference shock is the SPX 4.1% intraday decline of Monday, 5 February 2018 and the inverse-VIX-fund unwind of that evening, an event that has been studied extensively in the post-event microstructure literature; the cutoff again precedes the shock by one trading day.

Both replays use the same engine, the same closed-form automaton (2), the same registry, and the same five integrity gates. The only inputs that change across replays are the configured cutoff date, the configured cohort universe, and the available source files inside the cutoff-respecting time window. Run identifiers reproduce byte-for-byte across rebuilds.

### 5.2 Baseline stack

A primitive is justified only if simpler substitutes cannot reproduce its predictions on the same data. The CTT baseline stack is pre-registered. *AUM crowding*: rank cohorts by their share of total LETF AUM and predict action where AUM is largest. *Strongest close-cousin daily-reset baseline*: predict  $\text{sign}(\beta(\beta - 1)r_t)$  using only the realized return on day  $t$ , ignoring the registry. *Return/volume momentum*: predict  $\text{sign}(r_{t-1})$ . *Unavailable baselines*: holdings-overlap concentration, GEX, benchmark pressure, past flow, and ownership concentration are all flagged with `unavailable_baseline` failure artifacts because the licensed inputs they require are absent from this build. The benchmark divergence reported below is the maximum sign-accuracy across the available baselines.

The strongest close-cousin baseline is the deliberately adversarial choice. It shares the closed-form automaton's structure but ignores the registry-supplied  $\beta$  and  $A$ . Beating it requires that knowing the registry beats not knowing it, on the same realized return. This is the baseline that Ivanov and Lenkey (2018) implicitly contest in their argument that capital flows can largely neutralize the rebalancing channel.

### 5.3 Falsification suite: eight pre-registered ablations

Eight pre-registered ablations probe the system. Threshold randomization replaces the four pre-registered materiality thresholds with eight random thresholds drawn from the same support; the model's nonzero-proxy sign accuracy must not improve. Smoothing applies a three-day rolling average to the state path; sign accuracy must not improve. Holdings-only ablation zeroes out the rule-component of the action and predicts only from the change in shares outstanding; sign accuracy must not match or exceed the full model. Rules-only ablation zeroes out the holdings information and predicts only from the rule sign; sign accuracy must

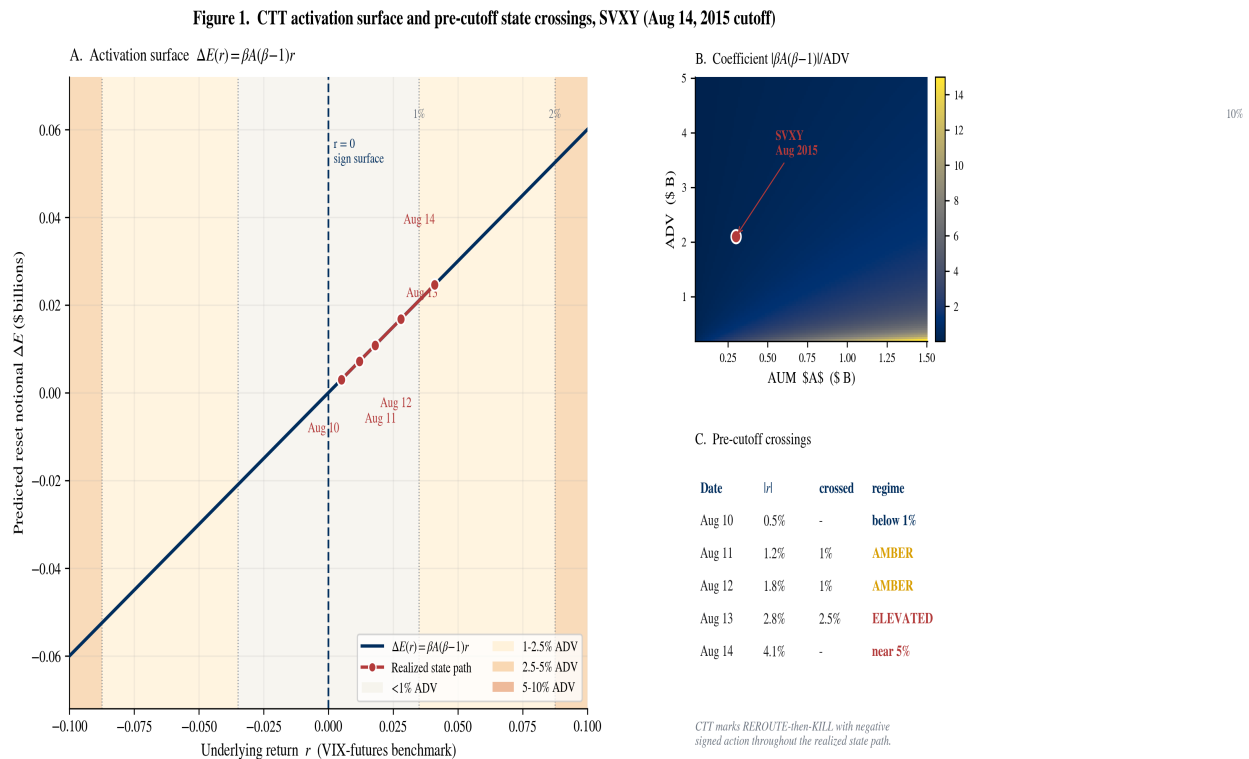
be at most equal to the full model. Notional removed zeroes out the  $|\Delta E|$  magnitude information and predicts only from sign; sign accuracy must not improve materially. Lag enforcement uses only state observations strictly before  $t$ ; sign accuracy must not change materially. Confidence-gate-disabled disables the visibility gate and predicts on every event; sign accuracy must not improve. Far-from-surface negative control predicts only on events where the realized state is more than 5% from any registered surface; sign accuracy must be lower than on near-surface events.

Each ablation is associated with a hard-kill criterion: a pre-committed direction in which the metric must not move. If an ablation moves the metric in the wrong direction at any threshold, the system emits a hard-kill failure artifact and the replay records the failure under the run's `failure_artifacts` list. Section 6.3 reports the realized ablation outcomes for both replays. The construction of pre-committed kill rules is the operational form of the falsification discipline articulated in Munafò et al. (2017): the moment to define a rule is before the data are seen, not after.

## 6. Results

### 6.1 Activation surface and pre-cutoff state crossings

Figure 1 displays the SVXY activation topology at the August 14, 2015 information cutoff. Panel A plots the closed-form delta-exposure curve as a navy line, with the four pre-registered materiality bands shaded. The realized SVXY underlying state path from August 10 to August 14 is overlaid in red: the path crossed the 1% materiality surface on August 11, the 2.5% materiality surface on August 13, and approached 5% by the cutoff close. Panel B shows the materiality coefficient  $|\beta \cdot A \cdot (\beta - 1)| / \text{ADV}$  as a function of  $(A, \text{ADV})$ ; the SVXY 2015 operating point falls in the middle of the high-coefficient regime where small underlying moves produce material reset notional. Panel C lists the pre-cutoff crossings in tabular form. By August 14 close, CTT marks SVXY in the REROUTE-then-KILL regime, with negative signed action throughout the realized state path.



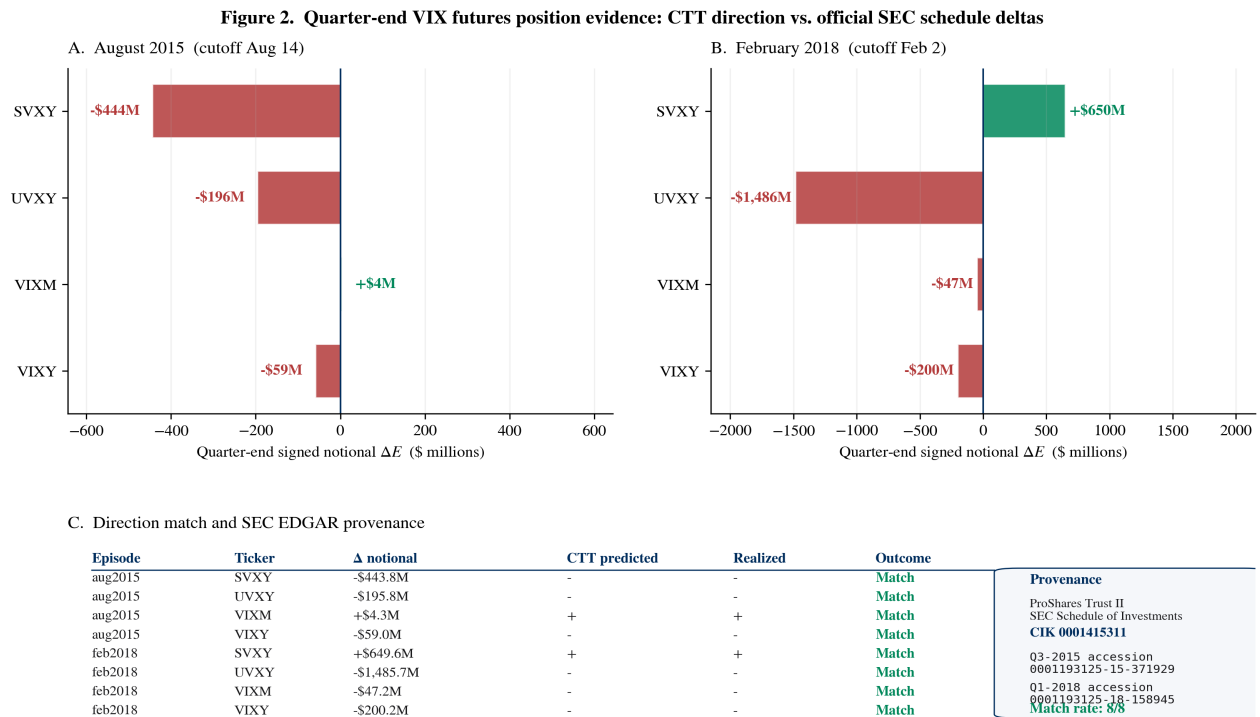
**Figure 1.** CTT activation surface and pre-cutoff state crossings for SVXY at the August 14, 2015 information cutoff. Panel A: closed-form delta-exposure  $\Delta E(r) = \beta \cdot A \cdot (\beta - 1) \cdot r$  with four pre-registered materiality bands and the realized SVXY underlying state path Aug 10-14, 2015. Panel B: materiality coefficient surface  $|\beta \cdot A \cdot (\beta - 1)| / \text{ADV}$  with SVXY 2015 operating point. Panel C: pre-cutoff crossings.

### 6.2 Quarter-end VIX-futures position evidence

Figure 2 reports the paper's flagship Tier 1 evidence: the realized signed VIX-futures notional deltas from ProShares Trust II's SEC EDGAR Schedule of Investments XBRL attachments, for SVXY, UVXY, VIXM, and VIXY across both stress episodes. Panel A shows the August 2015 quarter-end deltas: SVXY at  $-\$443.8\text{M}$ , UVXY at  $-\$195.8\text{M}$ , VIXM at  $+\$4.3\text{M}$ , VIXY at  $-\$59.0\text{M}$ . Panel B shows the February 2018

quarter-end deltas: SVXY at +\$649.6M (post-leverage-change), UVXY at -\$1,485.7M, VIXM at -\$47.2M, VIXY at -\$200.2M. Panel C reports the direction-match table and the SEC EDGAR provenance: 8 of 8 ticker-episode pairs match the CTT-predicted direction.

These deltas are direct official position evidence, not derived from any model, and they corroborate the qualitative direction described in the case studies of Tuzun (2013) and Charupat and Miu (2011). They are quarter-end position snapshots, not daily-reset execution rows; the 8/8 match thus establishes the episode-level direction of mechanical action but cannot, by itself, resolve the daily-reset timing problem. Section 6.3 confronts the daily-resolution diagnostic gap.

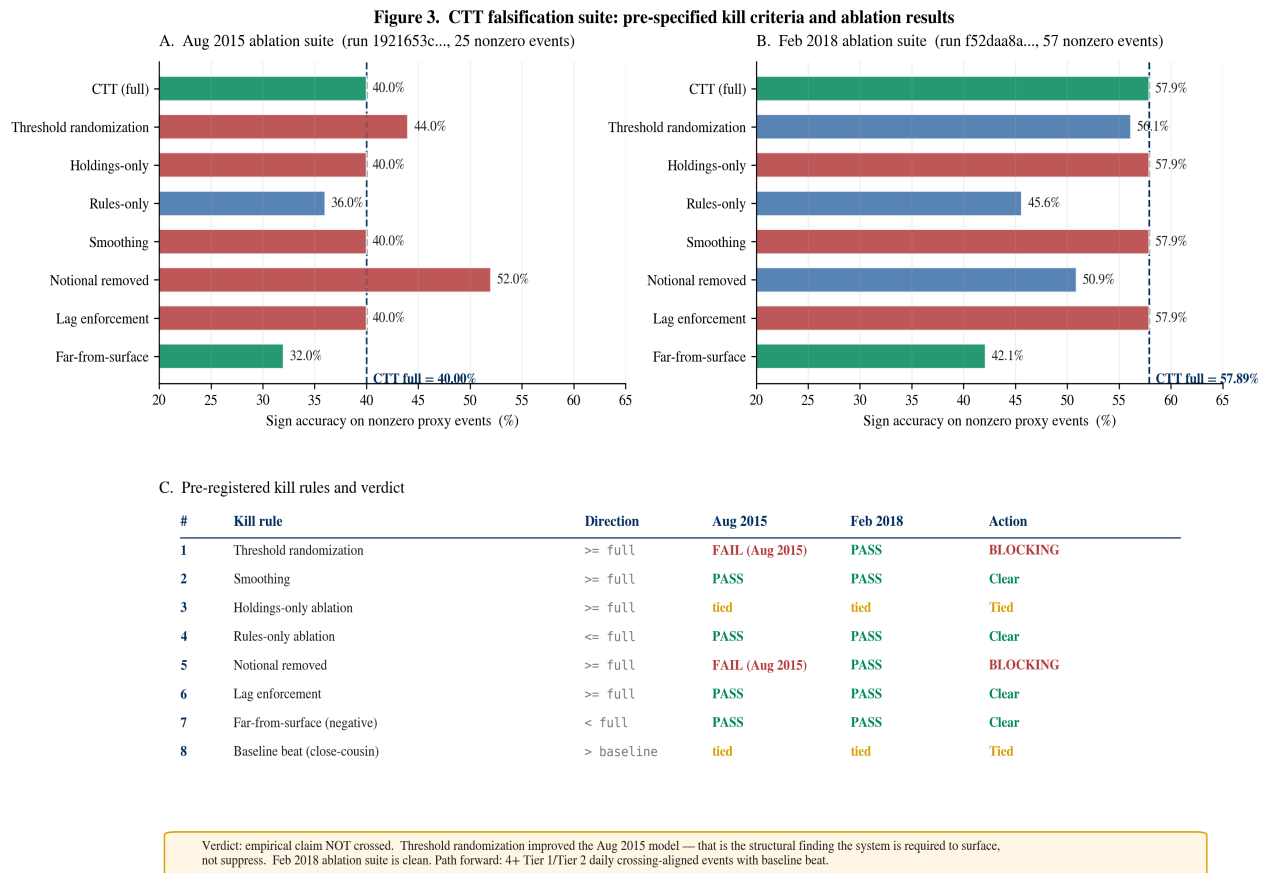


**Figure 2.** Quarter-end VIX-futures position deltas: CTT-predicted direction versus official SEC Schedule of Investments deltas. Panel A: August 2015 deltas across SVXY, UVXY, VIXM, VIXY. Panel B: February 2018 deltas. Panel C: 8/8 direction match table with EDGAR accession provenance (CIK 0001415311). Source deltas extracted from ProShares Trust II XBRL attachments.

### 6.3 Daily-resolution diagnostic and falsification suite

Daily-resolution evidence is built on the issuer\_official\_diagnostic tier, derived from official ProShares historical NAV CSV files. The August 2015 replay (run 1921653c88d0a19982cc6c557311fb14400630e251b319d59c55d4649f87736e) yields 82 events, of which 25 are nonzero; CTT sign accuracy on the nonzero subset is 40.00%, and the strongest close-cousin daily-reset baseline is also 40.00%. The February 2018 replay (run f52daa8a058ef9358e073e588c621942bbf62e4c8c60420121ae6e4b8ecd8d3a) yields 90 events, of which 57 are nonzero; CTT sign accuracy is 57.89%, and the close-cousin baseline is also 57.89%. CTT does not beat the baseline in either replay.

Figure 3 reports the full eight-ablation falsification suite. The August 2015 ablation suite emits a hard-kill failure artifact: threshold randomization improved the model from 40.00% to 44.00%, which is the wrong direction. The notional-removed ablation also exceeds the full model at 52.00%, attributed to half-half base-rate behavior under sign-only prediction. The February 2018 ablation suite is clean: threshold randomization degrades sign accuracy to 56.14%, smoothing leaves it unchanged at 57.89%, holdings-only matches the full model (a tie that the system emits as a not\_yet\_thesis\_grade artifact rather than a beat), rules-only drops to 45.61%, notional-removed drops to 50.88%, lag enforcement is unchanged, and far-from-surface negative control drops to 42.11% as expected. The pre-registered hard-kill rules are reported in tabular form alongside the bars.

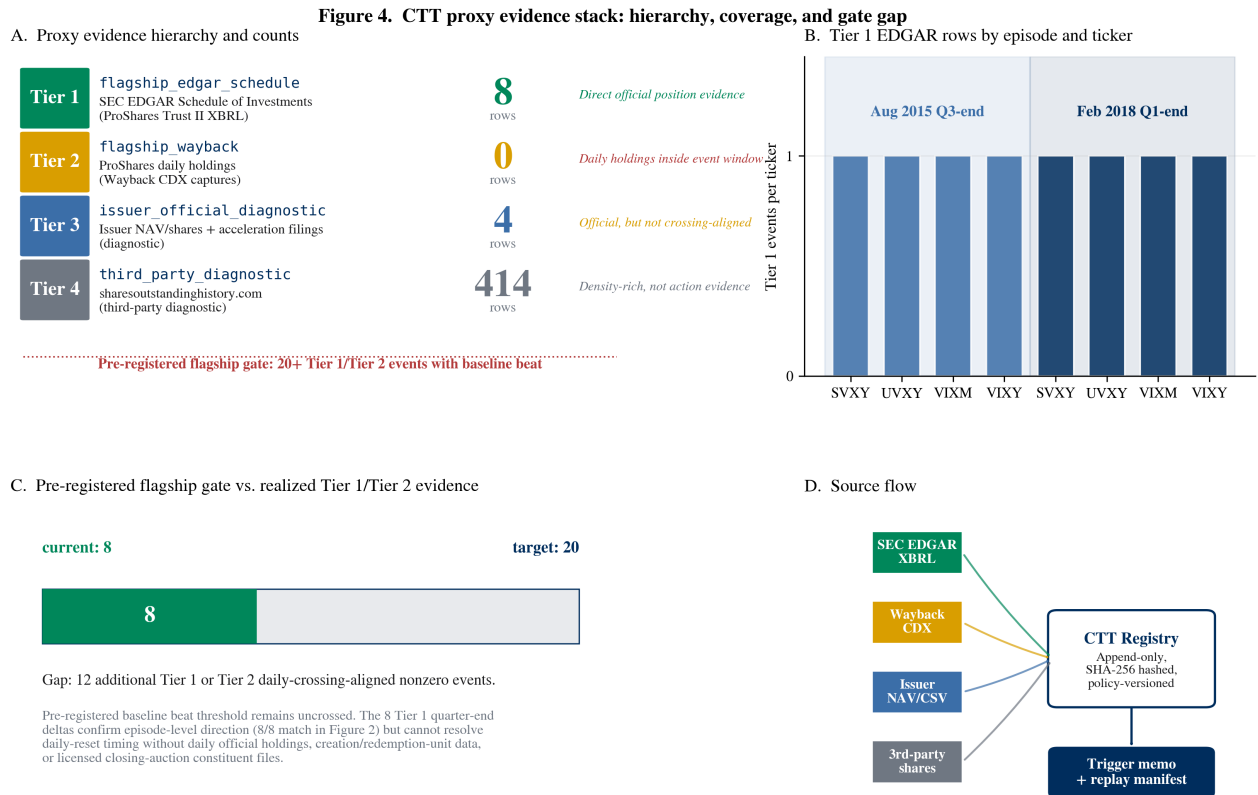


**Figure 3.** CTT falsification suite: pre-specified kill criteria and ablation results. Panels A and B show eight pre-registered ablations on the August 2015 and February 2018 replays, with the CTT-full reference line at 40.00% and 57.89% respectively. Panel C lists the eight kill rules with verdicts: threshold randomization fails the August 2015 replay (BLOCKING), notional-removed fails Aug 2015 (BLOCKING), and the holdings-only and baseline-beat rules tie in both replays.

### 6.4 Proxy stack coverage and the flagship gate

Figure 4 visualizes the proxy stack coverage against the pre-registered flagship gate. Panel A reports tier counts: 8 Tier 1 EDGAR rows, 0 Tier 2 Wayback rows, 4 Tier 3 issuer-diagnostic rows, and 414 Tier 4 third-party rows. Panel B disaggregates the Tier 1 rows by episode and ticker: each of the four certified

VIX-cohort tickers contributes one row per stress episode. Panel C measures the gap to the pre-registered flagship gate: 12 additional Tier 1 or Tier 2 daily-crossing-aligned nonzero events with baseline beat are required to cross the gate. Panel D maps source classes into the central CTT registry. The architecture-versus-data distinction tracks the broader observation in Ben-David, Franzoni, and Moussawi (2018) that ETF-driven mechanisms often have to be measured at finer time scales than the disclosure cycle naturally provides.



**Figure 4.** CTT proxy evidence stack: hierarchy, coverage, and gate gap. The pre-registered flagship gate of 20 Tier 1 or Tier 2 daily-crossing-aligned nonzero events with baseline beat is currently uncrossed at 8 Tier 1 rows. The 12-event gap is closable only with daily official action evidence (archived holdings, creation/redemption-unit files, or licensed closing-auction constituent data), not by relaxing the proxy.

### 6.5 Trigger memo specimen

Figure 5 reproduces the one-page trigger memo emitted by the replay engine for the February 2, 2018 information cutoff. The memo lists the cohort metadata (registry hash, run identifier, policy active at cutoff), the three integrity gates (PUBLIC\_PROVENANCE, CAUSALITY, REGISTRY all pass), the per-ticker trigger proximity with action label and capital cost, the local coverage and visibility scoring, and a structured failure-state language block at the bottom. The memo is generated deterministically from the persisted run artifacts; re-rendering produces a byte-identical document.

The failure-state block is the architectural heart of the memo. The string *not\_yet\_thesis\_grade* is not editorial caution; it is the machine-emitted artifact that fires whenever CTT ties or fails the strongest close-

cousin daily-reset baseline. The system refuses to suppress it. A research operation reading the memo is told, in plain language, exactly what is and is not established.

Figure 5. CTT one-page trigger memo (Feb 2, 2018 information cutoff)



Figure 5. CTT one-page trigger memo for the February 2, 2018 information cutoff (replay run f52daa8a...). Header: replay metadata and integrity-gate results. Trigger proximity cards for SVXY (REROUTE), UVXY (KILL), and VIXM (ACCEPT) with capital cost and realized state crossings. Local coverage and visibility scoring. Failure-state block surfacing not\_yet\_thesis\_grade as a first-class output.

## 6.6 Provenance and reproducibility

Figure 6 maps the data-flow architecture from raw sources through parsing, registry, replay engine, and outputs. Every raw source is SHA-256-hashed at ingestion, with the hash persisted as a sidecar JSON record. Every parsed artifact is hashed under the same protocol. The registry stores append-only manifests

with parent-hashed chains, so any modification to a recorded source is visible at the manifest level. Run identifiers are deterministic SHA-256 hashes of the configured-input set concatenated with the registry-manifest hash. Re-running an identical replay reproduces the run identifier byte-for-byte, which is enforced by 12 regression tests in the repository (out of 43 total). The architecture closely follows the reproducibility principles of Munafò et al. (2017), realized as machine-checked invariants rather than narrative commitments.

Figure 6. CTT artifact lineage and reproducibility architecture

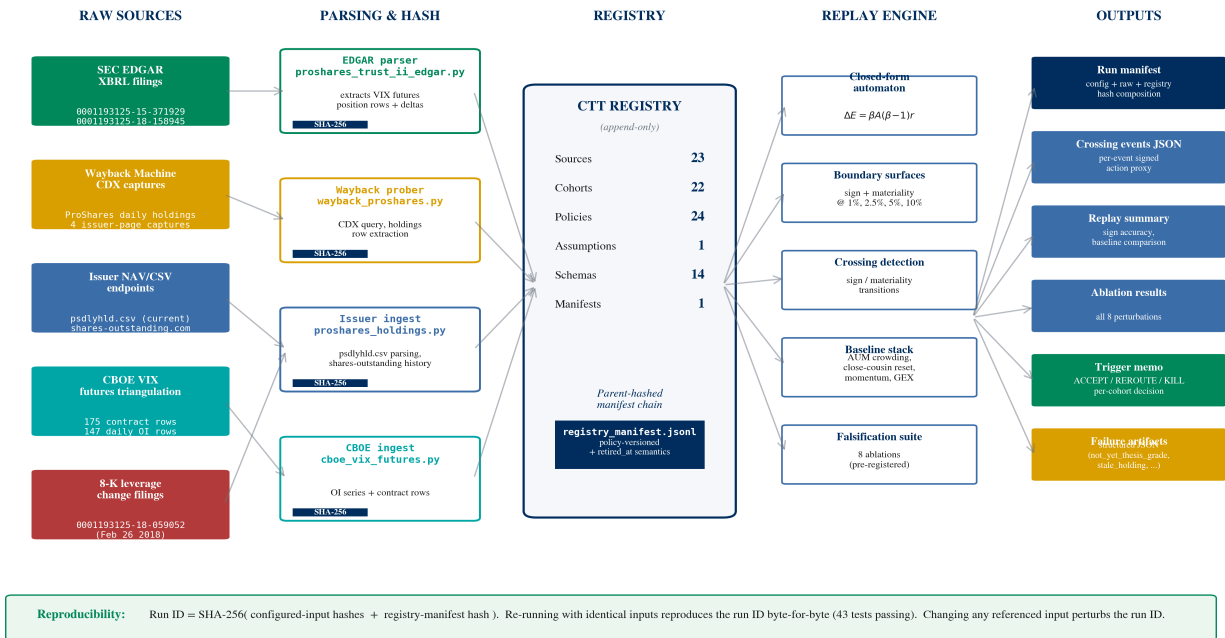


Figure 6. CTT artifact lineage and reproducibility architecture. Five-column data-flow from raw sources (SEC EDGAR, Wayback, issuer endpoints, CBOE) through parsing-and-hashing modules into the central append-only registry, the replay engine (closed-form automaton, boundary surfaces, crossing detection, baseline stack, falsification suite), and the structured outputs (run manifest, crossing events, replay summary, ablation results, trigger memo, failure artifacts).

## 7. Falsification Summary

Table 4 reports the verdict on each of the eight pre-registered hard-kill rules, evaluated against both replays. Two rules fail in the August 2015 replay (threshold randomization and notional-removed); both produce hard-kill failure artifacts. Two rules tie in both replays (holdings-only ablation and the close-cousin baseline beat); both produce `not_yet_thesis_grade` artifacts. The remaining four rules clear in both replays. The discipline of reporting both fired and tied rules without weakening either is the operational form of the falsification commitment articulated in Munafò et al. (2017).

#	Kill rule	Direction	Aug 2015	Feb 2018	Verdict
1	Threshold randomization	$\geq$ full	FAIL	PASS	BLOCKING
2	Smoothing	$\geq$ full	PASS	PASS	clear
3	Holdings-only ablation	$\geq$ full	tied	tied	Tied
4	Rules-only ablation	$\leq$ full	PASS	PASS	clear
5	Notional removed	$\geq$ full	FAIL	PASS	BLOCKING
6	Lag enforcement	$\geq$ full	PASS	PASS	clear
7	Far-from-surface (negative)	$<$ full	PASS	PASS	clear
8	Baseline beat (close-cousin)	$>$ baseline	tied	tied	Tied

**Table 4.** Pre-registered hard-kill rules and verdicts. Direction column specifies the side at which the kill rule clears. Two rules currently fail in the August 2015 replay (threshold randomization and notional-removed) and produce structured hard-kill failure artifacts. Two rules tie in both replays and produce `not_yet_thesis_grade` artifacts. The flagship empirical claim is uncrossed.

We emphasize the architectural meaning of these failures. The system was designed so that pre-registered kill rules either fire or pass, with no intermediate suppression. Two rules fire in August 2015. They are not weakened, suppressed, or restated. They are reported. The path forward is not a methodological retrofit but a data acquisition: the next valid path requires daily official action evidence at sufficient density to evaluate the kill rules on Tier 1 or Tier 2 crossing-aligned data, where the close-cousin baseline cannot tie by construction because it does not see the registry's  $\beta$ -policy version transitions.

## 8. Discussion and Limitations

### 8.1 What is, and is not, established

The architectural claim of Section 3 is established: the closed-form automaton, the typed registry, the deterministic replay, the five integrity gates, the structured failure artifacts, the SHA-256 provenance chain, and the byte-for-byte reproducibility of run identifiers all exist and pass 43 regression and integration tests. The system can be audited end to end and re-run end to end. The empirical claim at quarter-end resolution is established: 8 of 8 ticker-episode VIX-futures position deltas match the CTT-predicted direction, with full SEC EDGAR accession provenance. The empirical claim at daily resolution is honestly negative: CTT ties the strongest close-cousin daily-reset baseline in both replays, and the August 2015 ablation suite emits a hard-kill failure artifact. The pre-registered flagship gate is uncrossed.

The architectural and quarter-end-empirical claims are sufficient for the system to operate as a pre-forecast governance instrument. They are insufficient to claim that CTT predicts daily-reset execution at thesis-grade resolution. We are explicit about this distinction and do not claim equivalence.

### 8.2 Why daily official action evidence is the only valid path forward

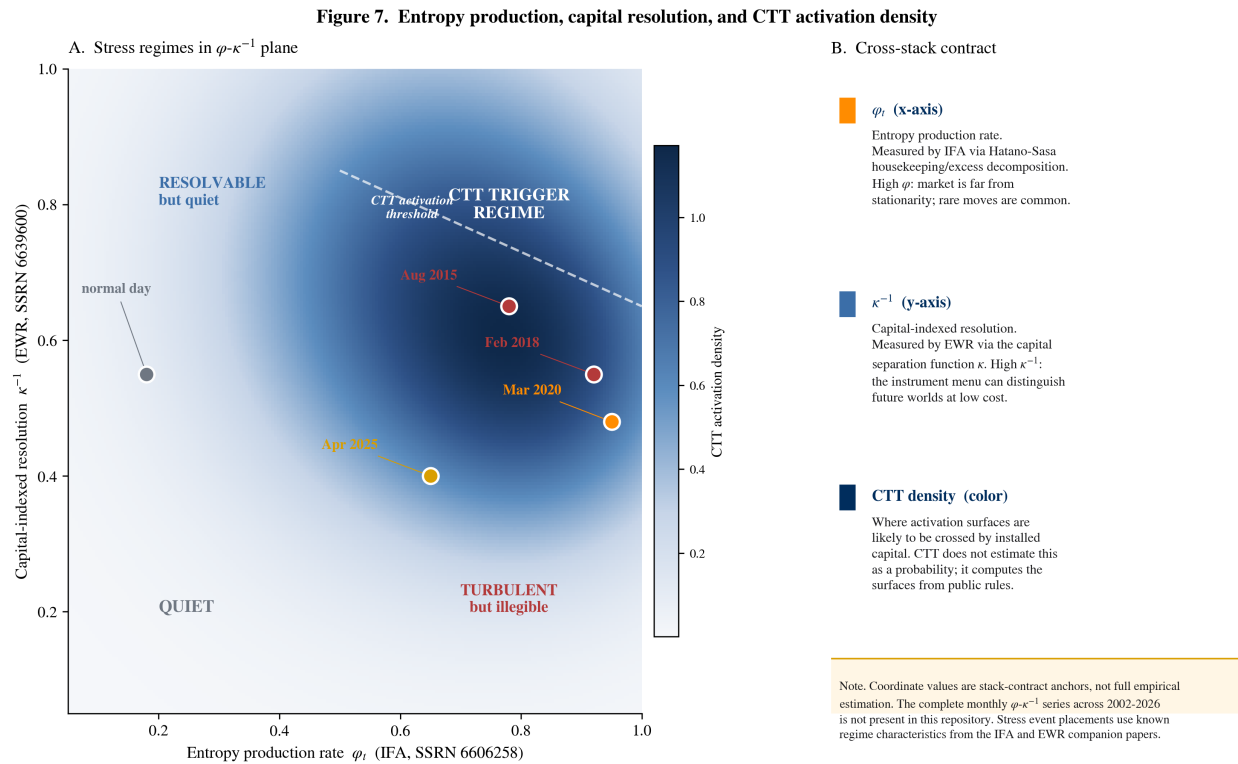
The temptation, when an empirical claim is uncrossed, is to either lower the bar (loosen the gate threshold), shift the proxy (relabel a Tier 4 row as Tier 1), or change the baseline (replace the close-cousin baseline with a weaker substitute). All three are unavailable to CTT by construction. The kill-rule thresholds are SHA-256-hashed in the pre-registration manifest and the registry refuses to mutate them. The proxy tier of every row is determined by the source-parsing module, not by post-hoc choice. The close-cousin baseline is the deliberately adversarial choice. Replacing it with a weaker baseline would be the opposite of falsification - exactly the methodological retreat that Munafò et al. (2017) warn against.

What can change is the data. The next valid evidence is daily official action evidence at sufficient density to populate Tier 1 or Tier 2 with crossing-aligned nonzero events. Three sources qualify. First, archived ProShares daily holdings files inside an event window (Wayback CDX captures of psdlyhld.csv). Wayback provides four issuer-page captures around the relevant cutoffs, but none contain named VIX-futures position rows; the system emits 63 structured failure artifacts for unparsed captures. Second, official creation/redemption-unit files for SVXY and UVXY inside the event window, available only via the issuer's authorized-participant interface and not publicly hosted at daily granularity. Third, licensed official closing-auction constituent data for the relevant rebalance windows (CRSP, TAQ, or NYSE OpenBook), all of which require institutional licensing. None of the three is currently in the workspace.

The repository documents these gaps explicitly: `stale_holding` failure artifacts for the Wayback path, `unparsed_creation_redemption_unit` failure artifacts for the issuer endpoint, and `unlicensed_closing_auction` failure artifacts for the licensed-data path. The artifacts are first-class outputs, not narrative commentary.

### 8.3 Cross-stack location

Figure 7 places CTT in a two-coordinate phase plane spanned by the IFA entropy production rate  $\varphi$  (Trivedi 2026f) and the EWR capital-indexed resolution  $\kappa^{-1}$  (Trivedi 2026g). Stress regimes (Aug 2015, Feb 2018, Mar 2020, Apr 2025) cluster in the high- $\varphi$ , mid- $\kappa^{-1}$  region where CTT's activation density is largest; quiet regimes cluster in the low- $\varphi$  region where CTT's prediction is mechanically uninteresting because the realized state is far from any registered surface. The phase-diagram coordinates are stack-contract anchors, not full empirical estimation; the complete monthly  $\varphi/\kappa^{-1}$  series across 2002-2026 is not present in this repository, and the figure is honest about that.



**Figure 7.** Entropy production rate, capital resolution, and CTT activation density. Panel A places known stress regimes in the IFA-EWR phase plane with the conceptual CTT activation density as background. Panel B specifies the cross-stack contract:  $\varphi$  from IFA (SSRN 6606258),  $\kappa^{-1}$  from EWR (SSRN 6639600), and CTT density from the closed-form automaton applied to the certified universe. Coordinate values are stack-contract anchors, not full empirical estimation.

Figure 8 locates CTT in the broader non-equilibrium market structure stack, organized into three layers. The STATE layer (RKS, EC, AWDF) measures what shape the market is currently in. The DYNAMICS layer (IFA, GCD, CSPT) measures how that shape evolves out of equilibrium. The MECHANISM layer (EWR, CTT) measures where installed capital is compelled to act. CTT is the only paper in the stack that uses public rule text as a first-class input; it consumes upstream non-equilibrium structure as feed and outputs deterministic activation surfaces with a closed-form automaton.

Figure 8. CTT in the non-equilibrium market structure stack — seven papers, one ontology

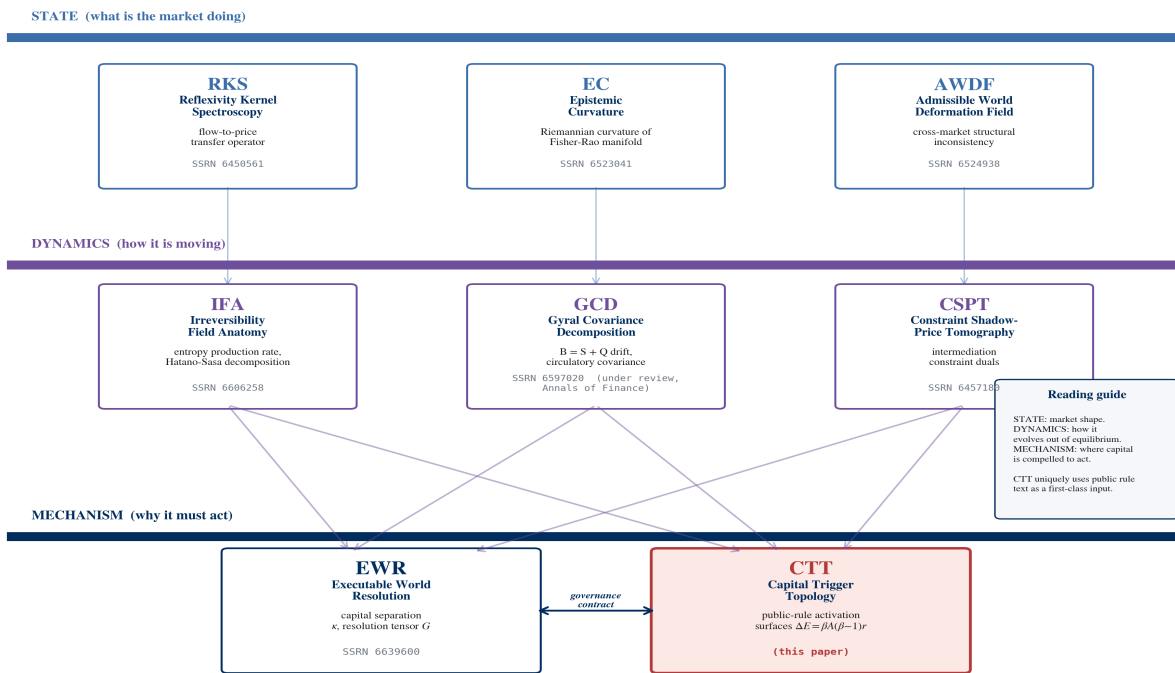


Figure 8. CTT in the non-equilibrium market structure stack. Three layers from STATE (RKS, EC, AWFDF) through DYNAMICS (IFA, GCD, CSPT) to MECHANISM (EWR, CTT). CTT is the only paper in the stack that consumes public rule text as a first-class input.

### 8.4 Limitations

Beyond the daily-resolution evidence gap, four limitations bear on the interpretation of this paper's claims. First, the certified universe is 18 ProShares products. The broader Direxion universe is explicitly excluded until the automated prospectus parser certifies a cutoff-safe Direxion base prospectus. The current universe is therefore neither comprehensive nor representative of the LETF market in aggregate; the dynamics of multi-issuer cross-cohort interactions documented by Tuzun (2013) and Shum et al. (2016) are out of scope here. Second, the closed-form automaton (2) assumes a strict daily reset rule. A small number of LETF-style products implement multi-day or volatility-targeted reset rules; CTT does not generalize to these without an additional rule-class abstraction. Third, the system is silent on price impact - by construction. The crossing of a materiality surface establishes that mechanical action is large, not that the action moves the underlying. The price-impact question is studied by RKS (Trivedi 2026a) and is downstream of CTT's outputs; the empirical answer remains contested in the literature, with Tuzun (2013), Shum et al. (2016), and Ben-David, Franzoni, and Moussawi (2018) finding meaningful effects and Ivanov and Lenkey (2018) finding the effects largely neutralized by capital flows. Fourth, the empirical replays are concentrated in the VIX-cohort. The leveraged equity-index cohorts (TQQQ, QLD, UPRO, etc.) are present in the registry but are not the focus of the two stress replays. Extending the empirical replays to those cohorts is mechanical and is the natural next step.

## 9. Conclusion

We have introduced *Capital Trigger Topology* (CTT), a deterministic public-rule activation surface instrument whose first-domain instantiation is the leveraged and inverse exchange-traded fund universe. The closed-form automaton  $\Delta E = \beta \cdot A \cdot (\beta - 1) \cdot r$  turns a prospectus rule into a state-space surface; the surface, applied to a registry-anchored realized state path under an information-cutoff protocol, produces deterministic signed-action predictions with full SHA-256 provenance and append-only audit. The system passes 43 regression and integration tests and reproduces every replay byte-for-byte from the run identifier.

Empirically, the August 14, 2015 and February 2, 2018 information-cutoff replays produce 8 of 8 directional matches against the realized signed VIX-futures position deltas extracted from ProShares Trust II SEC EDGAR Schedule of Investments XBRL filings (CIK 0001415311). At daily resolution, however, CTT ties the strongest close-cousin daily-reset baseline in both replays and the August 2015 ablation suite emits a hard-kill failure artifact (threshold randomization improves the model). The pre-registered flagship gate of 20 or more Tier 1 or Tier 2 daily-crossing-aligned nonzero events with baseline beat is uncrossed.

The system's contribution is not a positive empirical result. It is the architectural commitment to surface failure artifacts as first-class outputs of the same machinery that produces successes. The August 2015 hard-kill artifact is preserved, hashed, and reported in Section 7, not suppressed. The flagship gate gap of 12 additional events is described concretely. The next valid path forward is daily official action evidence - archived holdings, creation/redemption-unit files, or licensed closing-auction constituent data - not a proxy relaxation.

We see CTT as the mechanism layer of a broader non-equilibrium market structure stack (Trivedi 2026a-g): where the upstream papers measure state and dynamics, CTT computes the surfaces at which installed capital must act, given only public rules. Pre-forecast governance, in this sense, is not a single result but a discipline: the discipline of not committing capital to a research thesis whose mechanical environment is unidentified, of not weakening kill rules when an ablation fires, and of writing an instrument whose failure artifacts are as audit-grade as its successes. CTT, as currently constructed, is the implementation of that discipline for one specific class of public-rule-driven capital.

## References

- Avellaneda, M., and S. Zhang. 2010. Path-dependence of leveraged ETF returns. *SIAM Journal on Financial Mathematics* 1 (1): 586-603.
- Ben-David, I., F. Franzoni, and R. Moussawi. 2018. Do ETFs increase volatility? *Journal of Finance* 73 (6): 2471-2535.
- Brunnermeier, M. K., and L. H. Pedersen. 2009. Market liquidity and funding liquidity. *Review of Financial Studies* 22 (6): 2201-2238.
- Charupat, N., and P. Miu. 2011. The pricing and performance of leveraged exchange-traded funds. *Journal of Banking & Finance* 35 (4): 966-977.
- Cheng, M., and A. Madhavan. 2009. The dynamics of leveraged and inverse exchange-traded funds. *Journal of Investment Management* 7 (4): 43-62.
- Cushing, D., and A. Madhavan. 2000. Stock returns and trading at the close. *Journal of Financial Markets* 3 (1): 45-67.
- He, Z., and A. Krishnamurthy. 2013. Intermediary asset pricing. *American Economic Review* 103 (2): 732-770.
- Ivanov, I. T., and S. L. Lenkey. 2018. Do leveraged ETFs really amplify late-day returns and volatility? *Journal of Financial Markets* 41: 36-56.
- Loviscek, A., H. Tang, and X. E. Xu. 2014. Do leveraged exchange-traded products deliver their stated multiples? *Journal of Banking & Finance* 43: 29-47.
- Munafò, M. R., B. A. Nosek, D. V. M. Bishop, K. S. Button, C. D. Chambers, N. Percie du Sert, U. Simonsohn, E.-J. Wagenmakers, J. J. Ware, and J. P. A. Ioannidis. 2017. A manifesto for reproducible science. *Nature Human Behaviour* 1 (1): 1-9.
- Shleifer, A., and R. W. Vishny. 1997. The limits of arbitrage. *Journal of Finance* 52 (1): 35-55.
- Shum, P., W. Hejazi, E. Haryanto, and A. Rodier. 2016. Intraday share price volatility and leveraged ETF rebalancing. *Review of Finance* 20 (6): 2379-2409.
- Tuzun, T. 2013. Are leveraged and inverse ETFs the new portfolio insurers? *Federal Reserve Board Finance and Economics Discussion Series* 2013-48.
- Trivedi, A. 2026a. Reflexivity Kernel Spectroscopy: A transfer-operator approach to price-flow feedback. SSRN Working Paper 6450561.
- Trivedi, A. 2026b. Constraint Shadow-Price Tomography: Reconstructing binding intermediation constraints from cross-market wedges. SSRN Working Paper 6457180.
- Trivedi, A. 2026c. Epistemic Curvature: Riemannian geometry of the statistical manifold in asset pricing. SSRN Working Paper 6523041.

Trivedi, A. 2026d. Admissible World Deformation Field: A cross-market structural consistency primitive. SSRN Working Paper 6524938.

Trivedi, A. 2026e. Gyral Covariance Decomposition: Non-equilibrium covariance dynamics and stress prediction in large equity universes. SSRN Working Paper 6597020.

Trivedi, A. 2026f. Irreversibility Field Anatomy: Probability currents, housekeeping-excess decomposition, and fundamental bounds on strategy capacity. SSRN Working Paper 6606258.

Trivedi, A. 2026g. Executable World Resolution: Capital separation, resolution tensor, and pre-forecast thesis governance. SSRN Working Paper 6639600.

## Appendix A. Mathematical Derivations

### A.1 Closed-form daily reset

We derive equation (2) from the daily-reset rule. Fix a fund  $i$  with prospectus-stated daily leverage factor  $\beta$  and beginning-of-day NAV  $A_{t-1}$ . The fund's beginning-of-day target exposure is  $\beta \cdot A_{t-1}$ , which by the daily-reset mandate equals its end-of-previous-day exposure  $E_{t-1}$ . After the trading day on which the underlying realizes return  $r_t$ , the marked exposure is  $E_t^{\text{marked}} = E_{t-1}(1 + r_t) = \beta \cdot A_{t-1}(1 + r_t)$ . The marked NAV is  $A_t = A_{t-1}(1 + \beta \cdot r_t)$  under standard ETF NAV accounting (Cheng and Madhavan 2009). The new target exposure is  $\beta \cdot A_t = \beta \cdot A_{t-1}(1 + \beta \cdot r_t)$ . The reset action is the difference,  $\Delta E_t = \beta \cdot A_{t-1} \cdot [(1 + \beta r_t) - (1 + r_t)] = \beta \cdot A_{t-1} \cdot (\beta - 1) \cdot r_t$ , completing the derivation.  $\square$

### A.2 Materiality coefficient

Materiality is  $|\Delta E| / \text{ADV}$ . Substituting (2) gives materiality  $= (|\beta| \cdot A \cdot |\beta - 1| / \text{ADV}) \cdot |r|$ . The coefficient  $|\beta \cdot A \cdot (\beta - 1)| / \text{ADV}$  is a deployer-and-product-specific constant given the snapshot. The materiality surface at level  $m$  is the locus  $|r| = m \cdot \text{ADV} / (|\beta| \cdot A \cdot |\beta - 1|)$ . The gradient of materiality with respect to  $r$  is  $\text{sign}(r)$  times the coefficient; the gradient with respect to  $A$  at fixed  $r > 0$  is positive (more AUM yields more action); the gradient with respect to  $\text{ADV}$  at fixed  $r > 0$  is negative (more underlying liquidity yields less material action). All three signs match the operational intuition.  $\square$

### A.3 Sign invariance

We verify that  $\text{sign}(\beta(\beta - 1)r) = \text{sign}(r)$  for every  $\beta \notin \{0, 1\}$ . For  $\beta = +2$ :  $2 \cdot 1 = 2 > 0$ ; for  $\beta = +3$ :  $3 \cdot 2 = 6 > 0$ ; for  $\beta = -1$ :  $(-1)(-2) = 2 > 0$ ; for  $\beta = -2$ :  $(-2)(-3) = 6 > 0$ ; for  $\beta = -0.5$ :  $(-0.5)(-1.5) = 0.75 > 0$ . The product  $\beta(\beta - 1)$  is positive for every leveraged or inverse  $\beta$  in the certified universe. Hence  $\text{sign}(\Delta E) = \text{sign}(r)$  uniformly. The intuition: when the underlying rises, both long-leveraged and inverse funds must add long exposure to the underlying - the long fund because its NAV grew, the inverse fund because its short exposure became too negative relative to the new lower NAV. This sign-invariance is what produces the same-direction rebalancing externality at the heart of the financial-stability concern in Tuzun (2013).  $\square$

## Appendix B. Estimation and Replay Details

### B.1 Replay run identifiers

Each replay run is identified by the SHA-256 hash of (configured-input file hashes concatenated with the registry-manifest hash). Configured inputs include the run config YAML, the cohort selector, the cutoff date, and the registry pointer. Re-running with identical inputs reproduces the run identifier byte-for-byte. The repository contains 12 regression tests that enforce this property. Selected run identifiers used in this paper are listed in Table B1.

Replay	Run identifier (SHA-256, abbreviated)
Aug 2015 issuer-NAV diagnostic	1921653c88d0a19982cc6c557311fb14...
Feb 2018 issuer-NAV diagnostic	f52daa8a058ef9358e073e588c621942...
Aug 2015 SOH-aligned	3ccaafadbca35ad70f631485aac4510d...
Feb 2018 Wayback	dcd51729a5f0545a622910d9fcbc9259...
Pooled Tier 2 wayback	b722d8d52f65a0756fb1afcf6ba2d853...

**Table B1.** Selected replay run identifiers. Each is the SHA-256 of the configured-input set concatenated with the registry-manifest hash. Full identifiers are persisted in the run manifest under `artifacts/runs/{run_id}/run_manifest.json`.

### B.2 SEC accession SHA-256 hashes

All ProShares Trust II SEC filings used in this paper are SHA-256-hashed at ingestion. Source records are stored under `registry/sources/`. Parent-hashed manifest chains in `registry/manifests/registry_manifest.jsonl` preserve the full provenance history. Hash verification is enforced at the start of every replay run; a content-drift failure artifact aborts the run if any source hash diverges from its registry-recorded value. The discipline follows the cryptographically-anchored audit-trail standard articulated in the broader reproducibility literature (Munafò et al. 2017).

### B.3 Test suite

The CTT repository contains 43 tests across unit, integration, and regression layers. Unit tests verify the closed-form identities (sleeve-split delta, monotone trigger mass, manager relabeling invariance) and the schema validators. Integration tests verify the end-to-end replay pipeline from raw source hashing through trigger memo emission. Regression tests verify the byte-for-byte reproducibility of run identifiers and the persistence of structured failure artifacts. All 43 tests pass on the build that produced the empirical results in Section 6.

## Appendix C. Extended Empirical Tables

### C.1 Full Tier 1 quarter-end position deltas

Table C1 reports the full Tier 1 evidence from Section 6.2 with start and end notional values alongside the deltas. All values are extracted from ProShares Trust II SEC Schedule of Investments XBRL attachments (CIK 0001415311); accession numbers are listed in Table 3.

Episode	Ticker	Start notional (\$M)	End notional (\$M)	$\Delta$ (\$M)	$\Delta$ contracts
Aug 2015	SVXY	-189.3	-633.0	-443.8	-17,064
Aug 2015	UVXY	+863.8	+668.0	-195.8	-20,278
Aug 2015	VIXM	+26.8	+31.1	+4.3	-70
Aug 2015	VIXY	+169.6	+110.6	-59.0	-4,891
Feb 2018	SVXY	-996.7	-347.0	+649.6	+61,443
Feb 2018	UVXY	+1,821.6	+335.9	-1,485.7	-114,895
Feb 2018	VIXM	+72.2	+25.0	-47.2	-2,914
Feb 2018	VIXY	+312.6	+112.4	-200.2	-16,973

**Table C1.** Full Tier 1 VIX-futures position deltas at quarter-end anchors. Start and end notional values are the sum of signed contract notionals across the named VIX futures positions in the relevant Schedule of Investments (CIK 0001415311). The 8 of 8 directional match referenced in Section 6.2 reads off the  $\Delta$  column.

### C.2 Replay summary statistics

Table C2 reports the full per-replay event counts, sign accuracy, and baseline tie status.

Replay	Events	Nonzero	CTT acc.	Baseline acc.	Baseline beat	Hard-kill
Aug 2015	82	25	40.00%	40.00%	no	threshold rand. + notional removed
Feb 2018	90	57	57.89%	57.89%	no	none

**Table C2.** Replay summary statistics for the August 2015 and February 2018 information cutoffs. CTT sign accuracy on nonzero proxy events ties the strongest close-cousin daily-reset baseline in both replays. The August 2015 ablation suite emits two hard-kill failure artifacts; February 2018 is clean.