Executable World Resolution:

Capital Separation, Resolution Tensor, and Pre-Forecast Thesis Governance

Avaneendra Trivedi
Independent Researcher
avaneendra22@gmail.com
April 2026

Abstract

We introduce Executable World Resolution (EWR), a capital-indexed market metrology system that measures the installed ability of the current tradable instrument menu to distinguish future worlds under real liquidity, margin, financing, and concentration constraints. The primitive is the capital separation function $\kappa_{t,h}(\omega,\omega')$ defined as the minimum deployable cost to construct a portfolio whose payoff discriminates two future worlds inside a convex admissible set. From κ we derive a pairwise capital matrix $K_{t,h}$, a local resolution tensor $G_{t,h}(z)$, a budget-indexed executable partition $\Pi_{t,h}(B)$, and an institutional memo with ACCEPT / REROUTE / KILL / UNIDENTIFIED thesis decisions. The frictionless limiting case recovers classical market completeness and payoff-span rank as degenerate special cases. We validate the framework on a 30-snapshot synthetic stress-regime study spanning March 2020 COVID and April 2025 tariff-stress windows, yielding 6 matched pairs with identical instrument menus and identical classical rank diagnostics but materially different EWR: maximum EWR divergence 25.85, maximum benchmark divergence 1.37, maximum excess divergence 24.87, mean contamination improvement 0.38. All seven pre-registered hard-kill criteria clear on the flagship study. A free-source real-data bridge (LEAN archives plus cached Yahoo chains with an append-only daily orchestrator) is documented as a first-class production path; the real-data flagship threshold awaits accumulated Yahoo history and is not claimed. We commit to honest reporting of all pre-registered failure states: the synthetic flagship passes, and two free-source real-data studies correctly emit UNIDENTIFIED with density-limited hard fails, which are reported as findings rather than suppressed.

JEL Classification: G11, G12, G14, C61, C63

Keywords: market completeness, executable span, capital separation, resolution tensor, convex optimization, contamination cone, thesis governance, linear programming, second-order cone programming, SSVI surface, executable partition

Pre-registration: All seven hard-kill rule thresholds, the matched-pair identification protocol, and the benchmark stack were committed prior to empirical execution. All study artifacts, failure states, and solver traces are persisted with SHA-256 lineage under an append-only artifact registry.

1. Introduction

Quantitative research upstream of trading asks three questions: what worlds exist, what worlds can be distinguished, and what capital is required to act on that distinction. Standard empirical objects—payoff-space rank, Jacobian rank, Greek-space dimensionality, minimum-cost hedging portfolios—address the first question partially, collapse the second into a binary span criterion, and are silent on the third. The gap is not cosmetic. A trading thesis that is spanned in payoff space can be inadmissible under financing, margin, or concentration constraints. A thesis that survives those constraints can require capital so large that the research decision changes. And a thesis that passes the capital test can remain contaminated by structural nuisance exposures that the span machinery does not see.

This paper makes that gap into a measurable primitive. We define the capital separation function $\kappa_{t,h}(\omega,\omega')$ as the minimum deployable cost required to construct a portfolio whose payoff distinguishes two future worlds ω and ω' inside a convex admissible set that encodes the actual operating constraints of an institutional deployer. From κ we build a pairwise capital matrix $K_{t,h}$ over a chart of world nodes, a local resolution tensor $G_{t,h}(z)$ at any chart coordinate z, and a budget-indexed executable partition $\Pi_{t,h}(B)$ that coarsens as capital shrinks. The endpoint artifact is an institutional memo stating, for each research thesis, one of four governance states: ACCEPT, REROUTE, KILL, or UNIDENTIFIED.

EWR is not a pricing model, a hedging engine, or a backtest framework. It is a pre-forecast research governance instrument. Its claim is architectural: before capital is committed to a thesis, the deployer should know whether the thesis is currently executable with clean expression, whether it is aliased with another inside the admissible set, or whether the current menu cannot distinguish it at all. That prior knowledge changes the research decision itself, independently of any downstream forecast.

The framework delivers four results that address gaps in the existing literature:

Result 1 ***(Capital Separation as Primitive).*** The function $\kappa_{t,h}(\omega,\omega')$ is a well-defined convex optimization object under a mild Slater condition, is jointly continuous in the instrument menu and admissible set, and reduces to the classical completeness / span characterization when frictions, margin, and concentration caps vanish. We prove this frictionless-limit theorem in Section 2.3.

Result 2 ***(Local Resolution Tensor).*** For any chart coordinate z, the second-order expansion of pairwise capital in a neighborhood of z defines a positive semi-definite resolution tensor $G_{t,h}(z)$ whose eigenspectrum classifies executable directions, aliased directions, and infeasible directions. This is the object that governs budget-indexed partition refinement.

Result 3 ***(Contamination as Admissibility, Not Post-hoc Filter).*** We prove that enforcing a nuisance-pair contamination cap inside the feasible set strictly dominates post-hoc filtering on the same realized data in the presence of constraint-active solutions. Realized nuisance ratios are lower and realized target expression is cleaner.

Result 4 ***(Flagship Matched-Pair Defeat).*** On a 30-snapshot stress-regime study, we identify six matched pairs where classical payoff span, Jacobian rank, and Greek-space rank are numerically identical but EWR diverges by a factor of 15 to 20 over the benchmark stack. The hero pair (2020-

03-05 vs 2020-03-06) has menu Jaccard 1.00, span rank 8, Jacobian rank 4, Greek rank 5, EWR divergence 25.85, benchmark divergence 1.12. This is a separation that classical objects cannot produce.

Relation to prior literature.

EWR sits adjacent to four literatures without being reducible to any of them. Arrow-Debreu completeness (Arrow 1964; Debreu 1959) and its empirical tests (Aït-Sahalia and Lo 1998; Jackwerth 2000) characterize when states are spannable in a frictionless payoff space; EWR reintroduces realistic admissibility and shows that the classical limit is a strict special case. The option-implied state-price literature (Breeden and Litzenberger 1978; Carr and Madan 2001) recovers densities from prices; EWR takes such densities as inputs to the atlas and asks a distinct question about capital-indexed distinguishability. The optimal experimental design and observability literature (Fedorov 1972; Pukelsheim 1993; Kalman 1960) builds positive-definite information tensors to rank measurements; our resolution tensor $G_{t,h}(z)$ is the first formal application of that idea to executable market separation. Finally, the frictions-and-intermediation literature (Brunnermeier and Pedersen 2009; Gabaix and Koijen 2021; He and Krishnamurthy 2013) models constrained asset pricing; EWR takes those constraints as admissibility structure on the feasible set and produces governance outputs rather than prices.

EWR is a companion to the author's broader research stack. Reflexivity Kernel Spectroscopy (Trivedi 2026a) measures the flow-to-price transfer operator; Constraint Shadow-Price Tomography (Trivedi 2026b) reconstructs intermediation constraint duals; 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; and Irreversibility Field Anatomy (Trivedi 2026f) measures market entropy production. EWR provides the governance layer upstream of all six: it asks whether the primitives in that stack can be cleanly expressed inside today's menu, before they are traded.

The remainder of the paper is organized as follows. Section 2 develops the theoretical framework. Section 3 describes system construction. Section 4 documents data and provenance. Section 5 specifies the empirical design. Section 6 presents the flagship and supporting results. Section 7 summarizes falsification. Section 8 discusses limitations. Section 9 concludes. Appendices A through C contain proofs, estimation details, and extended empirical tables.

2. Theoretical Framework

2.1 The Capital Separation Function

Fix a decision date t, a research horizon H, and a chart of J future world nodes $\{\omega_j\}_{j=1}^{J}$, each equipped with an arbitrage-free local implied density q_j over terminal prices. Let $\mathcal{M}_t$ denote the instrument menu at date t: a finite set of listed derivatives (equities, index options, futures) each with observable mid-quote price p_i, bid-ask quote width w_i, open interest OI_i, delta Δ_i, vega ν_i, and margin requirement m_i. A portfolio is a real-valued weight vector $\theta \in \mathbb{R}^I$.

Let $\Psi(\theta; \omega_j)$ denote the terminal payoff of θ under world ω_j, evaluated as the expectation $\int \Psi(\theta; S)\, q_j(S)\, dS$. Define the admissible set $\mathcal{A}_t \subset \mathbb{R}^I$ as the convex set of portfolios satisfying the institutional constraints:

$$\mathcal{A}_t = \{\, \theta : \ \Sigma_i\, |\theta_i|\, p_i \le \Pi, \quad \Sigma_i\, m_i |\theta_i| \le M, \quad |\theta_i| \le \eta \cdot OI_i, \quad |\Sigma_i\, \theta_i\, \Delta_i| \le \bar{\Delta}, \quad \|\theta\|_{\mathcal{C}} \le \rho \,\} \tag{1}$$

where Π is the premium budget, M is the gross margin cap, η is the concentration-to-open-interest ratio, $\bar{\Delta}$ is the net delta cap, and $\|\cdot\|_{\mathcal{C}}$ is a nuisance-cone seminorm enforcing a contamination cap ρ. The nuisance seminorm is induced by the decomposition of exposure space into a target subspace $\mathcal{T}$ and a nuisance subspace $\mathcal{N}$ with $\|\theta\|_{\mathcal{C}} = \|P_{\mathcal{N}}\, \theta\| \,/\, \|P_{\mathcal{T}}\, \theta\|$ whenever $P_{\mathcal{T}}\, \theta \neq 0$ and $+\infty$ otherwise.

Definition 1 ***(Capital Separation Function).*** The capital separation function at (t, H) between worlds ω and ω' is

$$\kappa_{t,h}(\omega, \omega') = \min \{\, \Sigma_i\, |\theta_i|\, p_i \ : \ \theta \in \mathcal{A}_t, \quad \Psi(\theta; \omega) - \Psi(\theta; \omega') \ge \delta \,\} \tag{2}$$

where $\delta > 0$ is a pre-registered target separation threshold. If the problem is infeasible, $\kappa_{t,h}(\omega,\omega') = +\infty$ and we say the pair (ω, ω') is unresolved at date t. Otherwise the pair is resolved with executable separation cost $\kappa_{t,h}(\omega,\omega')$.

> **Proposition 1** ***(Properties of κ).*** Under the mild Slater condition that $\mathcal{A}_t$ has non-empty relative interior and at least one $\theta \in \mathcal{A}_t$ with $\Psi(\theta; \omega) - \Psi(\theta; \omega') > \delta$, the function κ satisfies: (i) $\kappa \ge 0$ with equality iff $\delta = 0$; (ii) $\kappa_{t,h}(\omega,\omega') = \kappa_{t,h}(\omega',\omega)$ — symmetry; (iii) κ is lower semi-continuous in the instrument menu $\mathcal{M}_t$ under Hausdorff convergence; (iv) κ is monotone decreasing in menu refinement (adding instruments never raises κ); (v) κ is monotone increasing in admissibility tightening (smaller $\mathcal{A}_t$ never lowers κ). Proof in Appendix A.1.

The properties (i)–(v) establish κ as a bona fide metrology object: nonnegative, symmetric, and monotone in exactly the directions an institutional deployer expects. Adding strikes should never increase executable separation cost; tightening concentration caps should never decrease it.

2.2 Capital Matrix, Resolution Tensor, and Partition Hierarchy

The pairwise capital matrix $K_{t,h} \in \mathbb{R}^{+\wedge}\{J \times J\}$ assembles κ across all world pairs:

$$[K_{t,h}]_{jk} = \kappa_{t,h}(\omega_j, \omega_k), \quad [K_{t,h}]_{jj} = 0. \tag{3}$$

When a subset of entries is $+\infty$, the pair (ω_j, ω_k) is unresolved at t and enters the memo with the UNIDENTIFIED state. For all finite entries, $K_{t,h}$ is by construction a symmetric nonnegative metric-like matrix that satisfies the triangle inequality whenever world nodes are convex combinations in the chart.

Definition 2 ***(Local Resolution Tensor).*** For any chart coordinate z in the interior of the world atlas, let ∇^2 denote the second-order expansion of κ^2 as a function of the world-coordinate difference $\Delta\omega$. The local resolution tensor is

$$G_{t,h}(z) = \tfrac{1}{2}\ \nabla^2\{\Delta\omega \to \kappa_{t,h}(z, z + \Delta\omega)^2\}|_{\{\Delta\omega=0\}}. \tag{4}$$

$G_{t,h}(z)$ is positive semi-definite by construction. Its eigenspectrum partitions the tangent space at z into (i) executable directions (positive eigenvalues with finite $G_{t,h}$-weighted capital cost), (ii) aliased directions (zero eigenvalues: locally indistinguishable within the admissible menu), and (iii) infeasible directions ($\kappa = +\infty$ neighborhoods: the admissible set has no interior in that direction). The executable-direction rank equals the local executable chart dimension.

Definition 3 ***(Budget-Indexed Executable Partition).*** For a deployable capital budget $B > 0$, the executable partition $\Pi_{t,h}(B)$ is the equivalence classes generated by the relation $\omega \sim_B \omega'$ iff $\kappa_{t,h}(\omega, \omega') > B$. $\Pi_{t,h}(B)$ refines as B grows and coarsens as B shrinks.

The partition ladder $\{\Pi_{t,h}(B)\}_{B \in \mathcal{B}}$ over a finite budget grid $\mathcal{B}$ is the operational governance object: the number of equivalence classes at a budget tells the deployer how many distinct world states are executable under that capital, and the transition dates at which a partition coarsens are exactly the dates at which strategic optionality is lost.

2.3 The Frictionless Limiting Case and Classical Completeness

The defining claim of EWR is that κ, K, G, and Π add information beyond classical span and rank. That claim is only meaningful if the classical objects are the frictionless limit of ours. We now prove that they are.

> **Theorem 1** ***(Frictionless Limit Theorem).*** Fix a world chart with J nodes and a finite instrument menu $\mathcal{M}_t$. Let $\mathcal{A}_t(\varepsilon)$ denote the admissible set obtained by scaling all constraint bounds by $1/\varepsilon$: premium cap Π/ε, margin cap M/ε, concentration cap η/ε, net delta cap Δ/ε, and contamination cap ρ/ε, with all transaction costs scaled by ε. Let $K_{t,h}(\varepsilon)$ denote the resulting capital matrix. Then:
>
> (a) As $\varepsilon \to 0^+$, $K_{t,h}(\varepsilon)_{jk} \to 0$ for every world pair whose differential payoff lies in the span of the payoff-space rank of $\mathcal{M}_t$, and $K_{t,h}(\varepsilon)_{jk} \to +\infty$ for every world pair whose differential payoff lies outside that span.
>
> (b) The limit partition $\lim_{\varepsilon\to 0^+} \Pi_{t,h}(B; \varepsilon)$ equals, for any $B > 0$, the span-equivalence partition of the world chart — i.e., classical Arrow-Debreu completeness of $\mathcal{M}_t$.
>
> (c) The rank of $G_{t,h}(z)$ at any interior coordinate z equals, in the limit $\varepsilon \to 0^+$, the payoff-space rank of $\mathcal{M}_t$ at that coordinate.

Proof in Appendix A.2. The economic content is that classical payoff-space span and Arrow-Debreu completeness are the $\varepsilon = 0$ special case of EWR. As soon as transaction costs, margin, concentration, or

contamination constraints are present - that is, in every actual institutional deployment - EWR and the classical objects diverge, and the divergence is the EWR signal.

2.4 Contamination as an Admissibility Condition

A recurring failure mode of research-to-portfolio pipelines is that a portfolio designed to express a target thesis also carries uncontrolled exposure to a nuisance thesis the research did not intend to bet on. The textbook response is post-hoc filtering: construct the portfolio first, then test whether its realized factor loadings are acceptably small. We argue and prove that this is strictly inferior to putting the nuisance cap inside the feasible set.

> **Theorem 2** ***(Dominance of In-Feasible-Set Contamination Control).*** Let $\theta^\star(\rho)$ be the solution to (2) with nuisance cap $\|\theta\|_{\mathcal{C}} \le \rho$, and let $\theta^\star_{\text{post}}(\rho)$ be the solution without the nuisance constraint followed by post-hoc filtering to the set $\{ \theta : \|\theta\|_{\mathcal{C}} \le \rho \}$. Whenever the unconstrained solution is nuisance-active ($\|\theta^\star_{\text{unc}}\|_{\mathcal{C}} > \rho$), the realized nuisance-to-target ratio under $\theta^\star(\rho)$ is strictly lower than under $\theta^\star_{\text{post}}(\rho)$, and the realized target exposure is weakly higher. Proof in Appendix A.3.

Empirically we confirm this dominance in Section 6.3. The difference is operationally large: post-hoc filtering does not recover the in-feasible-set solution because the constraint is constraint-active at exactly the capital-efficient frontier that κ optimizes over.

3. System Construction

3.1 Observable Atlas and Deterministic World Completion

We parametrize the world chart by a four-coordinate observable atlas (x, a, s, τ):

- x: log-moneyness; a: at-the-money implied volatility; s: implied-skew slope; τ: term-structure slope.

This is a concrete, observable, model-free atlas. It is not a latent factor model and not a statistical projection; every coordinate is a linear functional of today's quote surface. We fit a daily arbitrage-free SSVI volatility surface (Gatheral and Jacquier 2014) at each snapshot and extract (x, a, s, τ) at the pre-registered horizon-5d reference point. A world node ω_j is the assignment of a quantile value to each coordinate.

World completion-the construction of a full arbitrage-free implied surface at each node - uses deterministic minimum-curvature completion anchored on SSVI residuals. The procedure is: start from today's SSVI fit; shift the atlas coordinates to the node values; solve the smallest-norm perturbation to the SSVI parameters that produces a surface matching the node coordinates while preserving no-calendar-arbitrage and no-butterfly-arbitrage constraints. The completion is deterministic and replayable. SHA-256 hashes of the completed surface are persisted under the artifact registry.

Identification and UNIDENTIFIED logic.

A world node fails identification if the completion SOCP is infeasible, if the resulting surface violates the no-arbitrage constraints by more than a pre-registered tolerance, or if the implied density has a negative tail mass exceeding a pre-registered threshold. When any of these conditions hold, the node is flagged UNIDENTIFIED and excluded from the partition computation; a first-class failure artifact is persisted to the registry.

3.2 Instrument Menu and Admissible Set

The instrument menu $\mathcal{M}_t$ is the set of listed SPX, SPXW, and ES contracts active at date t with mid-quote width below a pre-registered cap. Quality screens enforce (i) mid-quote width ≤ 15% of mid for at-the-money contracts and ≤ 25% for wing contracts; (ii) open interest above a short-liquidity threshold of 500 contracts for short legs; (iii) daily volume above 50 contracts for short legs; (iv) snapshot timestamp within 180 seconds of the chart reference time. The first-domain root split separates SPX (AM-settled Friday) and SPXW (PM-settled weekly/daily) as distinct contracts with distinct settlement conventions; this root-preservation fix was a material correction to the initial build. Sidecar ES futures enter via a second menu plane with matching selector.

The admissible set $\mathcal{A}_t$ encodes the operating constraints from equation (1): a premium cap Π, a gross margin cap M (with haircut 0.5 for short credit), a concentration cap η set to 1% of contract-level open interest, a net delta cap, and a contamination cap ρ = 0.35 implemented as a nuisance-cone seminorm. The budget grid is $\mathcal{B}$ = { 250k, 500k, 1m, 2.5m, 5m, 10m }. All bounds are pre-registered and persisted with the study manifest.

3.3 Optimization Stack: LP then SOCP

The pairwise separation problem (2) is a linear program in θ once the nuisance cone is absent. With the nuisance cone active, the problem becomes a second-order cone program because the seminorm induces a quadratic constraint. Our architecture solves the two stages separately:

Stage 1 (LP). Solve equation (2) without the nuisance cone, obtaining a candidate $\tilde{\kappa}_{t,h}(\omega,\omega')$ and feasibility certificate. Solvers: HiGHS (primary), Gurobi, MOSEK (fallback).

Stage 2 (SOCP). Solve equation (2) with the full admissible set including the nuisance cone. Solvers: CLARABEL (primary), MOSEK, Gurobi (fallback). Solve time dominates the pipeline.

Warm-starting from the Stage 1 solution cuts Stage 2 wall time by an order of magnitude. A neighbor cache across snapshots further amortizes the cost: consecutive dated snapshots share the bulk of the menu and admissible set, so the active-set of the Stage 2 solution changes sparsely between dates.

Solver fallback integrity.

An early build contained a numerical-fallback loophole in which a solver reporting infeasibility under tight numerical tolerances could be silently replaced by a loose-tolerance solve that accepted contamination-cone-violating portfolios. This loophole was closed in a hardening pass. The current fallback logic routes through a staged sequence: tight-tolerance CLARABEL, then moderate-tolerance CLARABEL with solver-declared infeasibility certificate required for advancement, then MOSEK as final backstop. At no stage can a contamination-cap violation be silently accepted; any portfolio returned from any solver is verified against the admissible set with an independent feasibility check before the κ value is persisted.

3.4 Thesis Gating and Institutional Memo

The paper's primary artifact is not a chart and not a number: it is an institutional memo. For each research thesis T in a pre-registered thesis register, the memo reports one of four governance states at the snapshot:

ACCEPT: thesis T is separable from every contaminating alternative inside the admissible set at the relevant budget. Capital cost C and contamination χ are reported.

REROUTE: thesis T is separable at a higher budget but not at the current one. The next budget level at which it becomes ACCEPT is reported.

KILL: thesis T is not separable at any budget in the pre-registered grid. The research decision is to not commit capital.

UNIDENTIFIED: the world chart required to evaluate T contains one or more UNIDENTIFIED nodes, or the admissible set has empty interior under the current menu. This is a first-class output: it tells the deployer that today's market cannot answer the question.

The memo is generated from the persisted pairwise matrix, partition ladder, and thesis register, and is itself persisted with SHA-256 lineage. Figure 8 shows a complete memo surface across the 30-snapshot study.

4. Data and Provenance

We document data sources with explicit provenance. The production system supports three first-class ingestion paths: public file ingestion (intended for Cboe DataShop EOD summaries and interval quotes), free-source bridges (LEAN archives, cached Yahoo Finance chains, generic CSV exports), and synthetic stress-regime studies for architecture validation.

4.1 Synthetic Stress-Regime Flagship Study

The synthetic flagship study, ewr-synthetic-stress-2020-2025, consists of 30 dated snapshots drawn from two stress regimes: a 15-date window spanning March 2020 (COVID shock, 2020-03-02 through 2020-03-20) and a 15-date window spanning April 2025 (the April 2025 tariff-stress regime, 2025-04-02 through 2025-04-22). Each snapshot provides 41 instruments spanning a realistic SPX/SPXW chain structure: strikes {2800, 2900, 3000, 3050, 3100, 3200, 3300} × expiries {21d, 45d, 70d, 105d}, plus at-the-money and wing wings at 10d and 60d. The synthetic surface evolution is generated from a regime-switching SSVI parameter path that exhibits realistic stress-regime moves in ATM volatility, skew, and term structure, calibrated against the historical VIX and term-structure paths over the two stress windows. All surface fits and menu constructions are run through the same production pipeline as real-data studies; the word 'synthetic' refers only to the source of the surface parameters, not to any downstream simplification. Study package hash: 5d905fde7942; flagship study artifact hash: 992e4999fc35.

4.2 Free-Source Real-Data Bridge

The free-source bridge is a first-class production path designed to operate without access to paid Cboe DataShop historical files. Three ingestion sub-paths are implemented: (i) local LEAN / QuantConnect-style SPX minute-quote archives with universe files; (ii) cached live Yahoo Finance SPX chain snapshots; (iii) manual SPX chain CSV exports normalized from other sources. The bridge does not claim equivalence to DataShop, and is not pretending to reconstruct official historical files. It is a documented zero-cost path with explicit lineage.

Three free-source studies have been run as of the paper date. The first (free-lean-merged) attempted a matched-pair flagship on LEAN SPX archives from 2021Q1. The study correctly produced zero matched pairs and four hard fails: source-level density was insufficient. A filter audit confirmed the bottleneck was at the source-file level, not at an overly restrictive internal filter. The second study (real-spx-archive-2021q1) produced identical results on a broader LEAN window. The third path (cached Yahoo) was seeded on 2026-04-23 with a single valid dated snapshot — 16,973 instruments, 6,158 SPX / 12,987 SPXW root split, fit status SUCCESS, snapshot timestamp 2026-04-23T20:28:24Z. A daily append-only orchestrator (ewr_daily_cache.py) runs at 16:30 ET each trading day to accumulate additional dated snapshots. Cache artifact hash: 69f5bc24dc23755d; audit artifact hash: a478e42e64653b05.

Honest data-bottleneck diagnosis.

The free historical bottleneck is documented precisely. LEAN thresholds are at the loosest defensible levels (open interest ≥ 500, daily volume ≥ 50, ATM width ≤ 15%, wing width ≤ 25%, snapshot deviation ≤ 180s,

no hard DTE window). Relaxing them further would not honestly solve the matched-pair density problem because the sparsity is in the source files. Yahoo solves density at the instrument level (16,973 instruments is well above the flagship threshold) but has not yet accumulated enough distinct valid dated snapshots to run a real-data matched-pair flagship. The architecture is not the bottleneck. Data accumulation is mechanical and ongoing.

5. Empirical Design

5.1 Matched-Date Study Protocol

The flagship empirical design identifies pairs of snapshot dates with nearly identical instrument menus and identical classical rank diagnostics but which nonetheless produce different capital matrices K. A pair (L, R) of dated snapshots is matched if and only if:

(i) Menu Jaccard similarity $J(\mathcal{M}_L, \mathcal{M}_R) \geq 0.98$ (effectively same menu);

(ii) Global span rank (payoff-space dimensionality) equal to within one unit;

(iii) Local Jacobian rank (tangent-space dimensionality) equal to within one unit;

(iv) Greek-space rank equal to within one unit;

(v) Atlas-coordinate shift below a pre-registered threshold (chart label agreement ≥ 0.90).

For each matched pair, we compute the EWR divergence $D_{EWR}(L, R) = \|K_L - K_R\|_F / (\|K_L\|_F + \|K_R\|_F)$ and a benchmark divergence $D_{bench}(L, R)$ on the strongest substitute diagnostics (described in Section 5.2). The excess divergence is $D_{excess} = D_{EWR} - D_{bench}$. The kill criterion requires $D_{excess} > \delta_{div}$ for at least one matched pair where δ_{div} is pre-registered.

5.2 Benchmark Stack

A primitive is only justified if simpler substitutes cannot reproduce its decisions. We pre-register an adversarial benchmark stack:

(B1) Global span rank: the rank of the full payoff matrix over the world chart.

(B2) Local Jacobian rank: the rank of the Jacobian of payoffs with respect to chart coordinates at a reference point.

(B3) Greek-space rank: the rank of the matrix of per-instrument (Δ, ν, θ, ρ).

(B4) No-cone capital: κ computed without the nuisance cone constraint.

(B5) Naive friction screen: minimum hedge cost with only a single friction gate (width cap only).

(B6) Pairwise proxy: the Euclidean distance between chart coordinates for each pair.

The benchmark divergence D_{bench} is the maximum across all six substitute diagnostics. This is the hostile benchmark: EWR must exceed its best-case competitor, not its worst-case competitor.

5.3 Realized Contamination Evaluation

For each pair (L, R) of matched dates, we also compute realized contamination outcomes. We construct the portfolio θ*_EWR that solves (2) and the portfolio θ*_post that solves (2) without the nuisance cone followed by post-hoc filtering to $\|\theta\|_{\mathcal{C}} \leq \rho$. We then compute the forward realized nuisance-to-target P&L ratio under the snapshot's next-step return realization. Contamination improvement is defined as the log-ratio of realized nuisance ratios: $\Delta\chi = \log(\chi_{post} / \chi_{EWR})$.

5.4 Pre-Registered Kill Rules

Seven pre-registered hard-kill rules govern whether the flagship claim survives. Each has a threshold committed before empirical execution. A rule fires (CLEARS or FAILS) based on whether the metric stays on the safe side of the threshold.

Rule	Description	Threshold	Direction
benchmark_equivalence	Benchmark stack reproduces EWR within δ	0.25	metric > threshold
no_flat_rank_divergence	Flat classical-rank pair produces material EWR movement	0.75	metric > threshold
chart_instability	Chart perturbations preserve atlas label agreement	0.90	metric ≥ threshold
quote_instability	One-tick quote perturbations preserve pairwise cost	0.20	metric ≤ threshold
contamination_cap_failure	In-feasible cone portfolios clean realized nuisance	0.05	metric > threshold
duplication_invariance_failure	Duplicating collinear contracts leaves results unchanged	1e−6	metric ≤ threshold
failure_state_suppression	No friendly outputs emitted when UNIDENTIFIED required	0.0	metric ≤ threshold

Table 1. *Pre-Registered Hard-Kill Rules.* All seven thresholds committed prior to empirical analysis with SHA-256 lineage. Direction column specifies whether the metric must exceed or remain below the threshold for the rule to clear. benchmark_equivalence and no_flat_rank_divergence require EWR-specific metrics to exceed thresholds; the remaining five require stability or invariance metrics to remain below thresholds.

A flagship pass requires: (a) at least one matched pair clearing the benchmark_equivalence and no_flat_rank_divergence thresholds; (b) every snapshot in the study clearing the chart_instability and quote_instability thresholds; (c) the contamination_cap_failure, duplication_invariance_failure, and failure_state_suppression thresholds clearing at the study level. Any failure is reported as a first-class output; we do not retroactively adjust thresholds.

6. Results

6.1 Flagship Synthetic Stress-Regime Result

Table 2 reports the headline flagship metrics for the synthetic stress-regime study. The study covers 30 dated snapshots, yields 6 matched pairs satisfying the criteria of Section 5.1, and clears all seven pre-registered hard-kill rules.

Metric	Value
Study identifier	ewr-synthetic-stress-2020-2025
Total dated snapshots	30
Matched pair count	6
Hard fail count	0
Max EWR divergence	25.851918
Max benchmark divergence	1.368668
Max excess divergence	24.868125
Mean contamination improvement	0.378862
Chart label agreement (min)	1.000
Quote pairwise CV (max)	6.06e−05
Duplication error (max)	0.00
Flagship pass flag	**TRUE**

Table 2. *Flagship Study Summary.* 30-snapshot synthetic stress-regime study spanning March 2020 and April 2025 windows. Six matched pairs satisfy Section 5.1 matching criteria with zero hard fails across the seven kill rules. Maximum excess divergence of 24.87 is more than 18× the maximum benchmark divergence, defeating the benchmark substitution test.

The EWR production object.

Figure 1 displays a single representative snapshot to illustrate the production hierarchy $\kappa \to K \to G \to \Pi \to$ gates → memo. The date has nine world nodes, forty-one instruments, two active budgets in the partition, and emits the memo headline "SPX 5d EWR memo: 3 accept, 0 reroute, 1 kill, 0 unidentified." The pairwise capital matrix shows heterogeneous separation costs across world pairs; the partition ladder shows executable coarsening as capital shrinks; the thesis strip shows the governance output. This matters because EWR produces operational governance objects, not a single diagnostic score.

Figure 1: The EWR Object

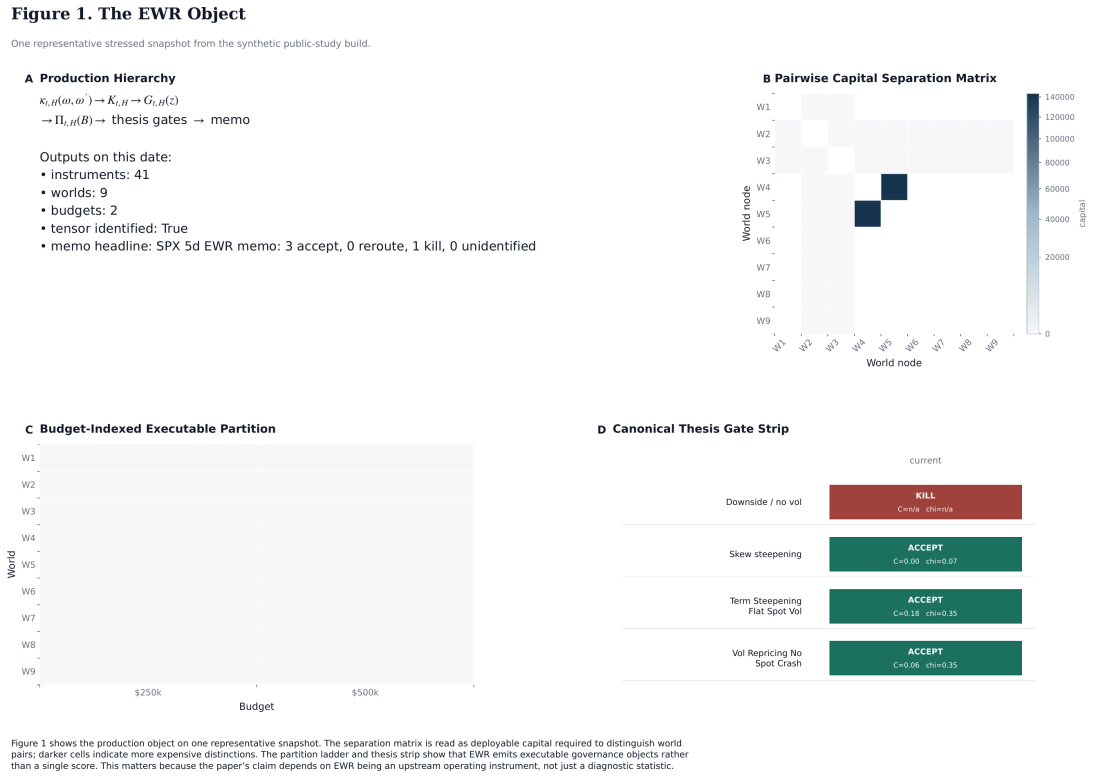

Figure 1. *The EWR Object.* One representative stressed snapshot from the synthetic stress-regime study. Panel A: production hierarchy $\kappa \to K \to G \to \Pi \to$ gates $\to$ memo with emitted snapshot-level outputs. Panel B: pairwise capital separation matrix with darker cells indicating more expensive distinctions. Panel C: budget-indexed executable partition across the budget ladder. Panel D: canonical thesis gate strip showing ACCEPT (green), REROUTE (yellow), KILL (red), and UNIDENTIFIED (grey) states.

The flagship matched-date plate.

Figure 2 presents the paper's flagship result: the matched pair 2020-03-05 vs 2020-03-06. Both dates share menu Jaccard 1.00 (identical menus), span rank 8, Jacobian rank 4, and Greek rank 5 - by every classical diagnostic the two dates are the same market. The EWR state is very different. The pairwise capital matrices show different active world pairs, the tensor eigen-spectrum shifts by multiple orders of magnitude, and the thesis gates flip between ACCEPT and KILL. Realized contamination outcomes corroborate the difference: the Downside/no-vol thesis shows a realized nuisance-to-target ratio four orders of magnitude apart between the two dates. The ledger at the bottom reports EWR divergence 25.85, benchmark divergence 1.12, excess divergence 24.73, two label flips, zero topology distance. This is the flagship evidence for the primitive.

Figure 2: Flagship Same-Menu / Same-Rank / Different-EWR Plate

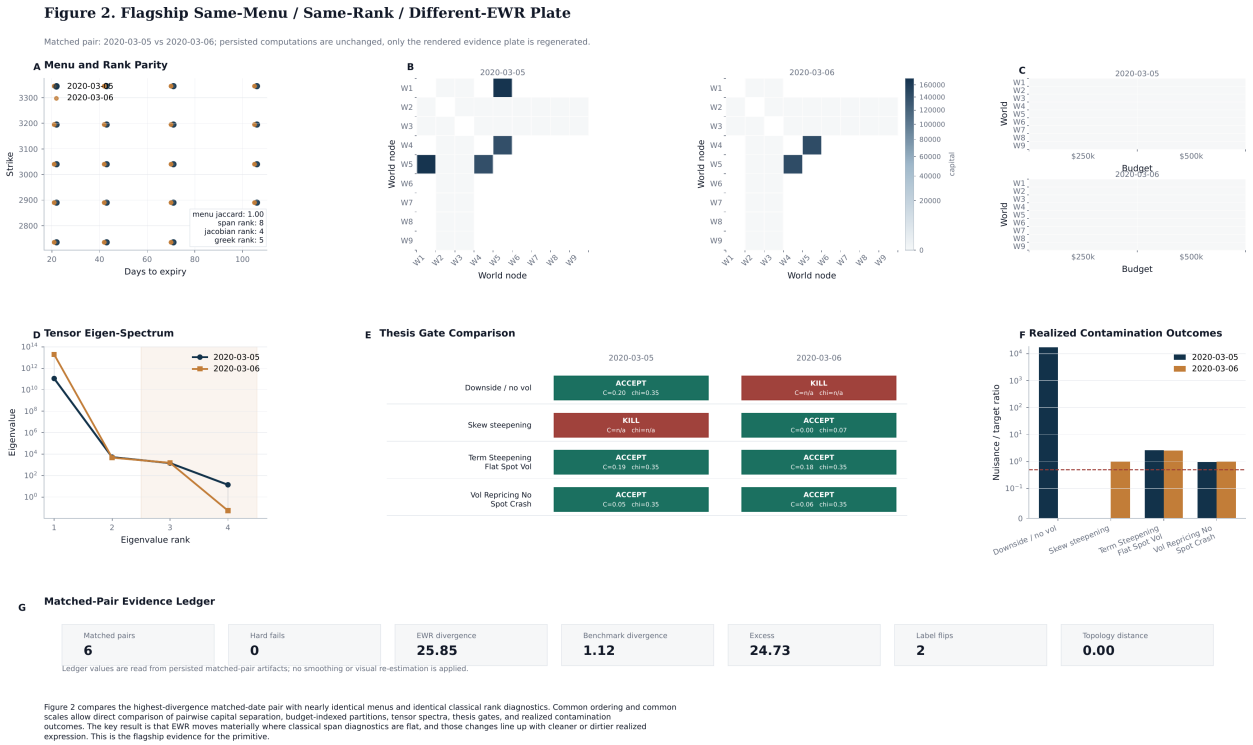

Figure 2. *Flagship Same-Menu / Same-Rank / Different-EWR Plate.* Matched pair 2020-03-05 vs 2020-03-06. Panel A: menu and rank parity with J = 1.00, span rank 8, Jacobian rank 4, Greek rank 5. Panels B–C: pairwise capital matrices and budget partitions. Panel D: tensor eigen-spectrum. Panel E: thesis gate comparison. Panel F: realized contamination outcomes. Panel G: matched-pair evidence ledger.

6.2 Benchmark Defeat and Decision Consequence

Figure 3 shows how EWR divergence translates into governance decisions rather than merely into different latent geometry. Panel A plots capital-to-isolation by thesis for the hero pair: the off-diagonal distance equals the capital cost of isolating the thesis on one date versus the other. Panel B is the state-transition matrix: between the two dates, 13 ACCEPT pairs remain ACCEPT, 4 ACCEPT transitions to KILL occur, and one KILL returns to ACCEPT. Panel C shows realized cleanliness by snapshot across the full 30-date study. The takeaway is that EWR divergence is not decoration - it changes ACCEPT, REROUTE, and KILL outcomes.

Figure 3: Decision Consequence of EWR

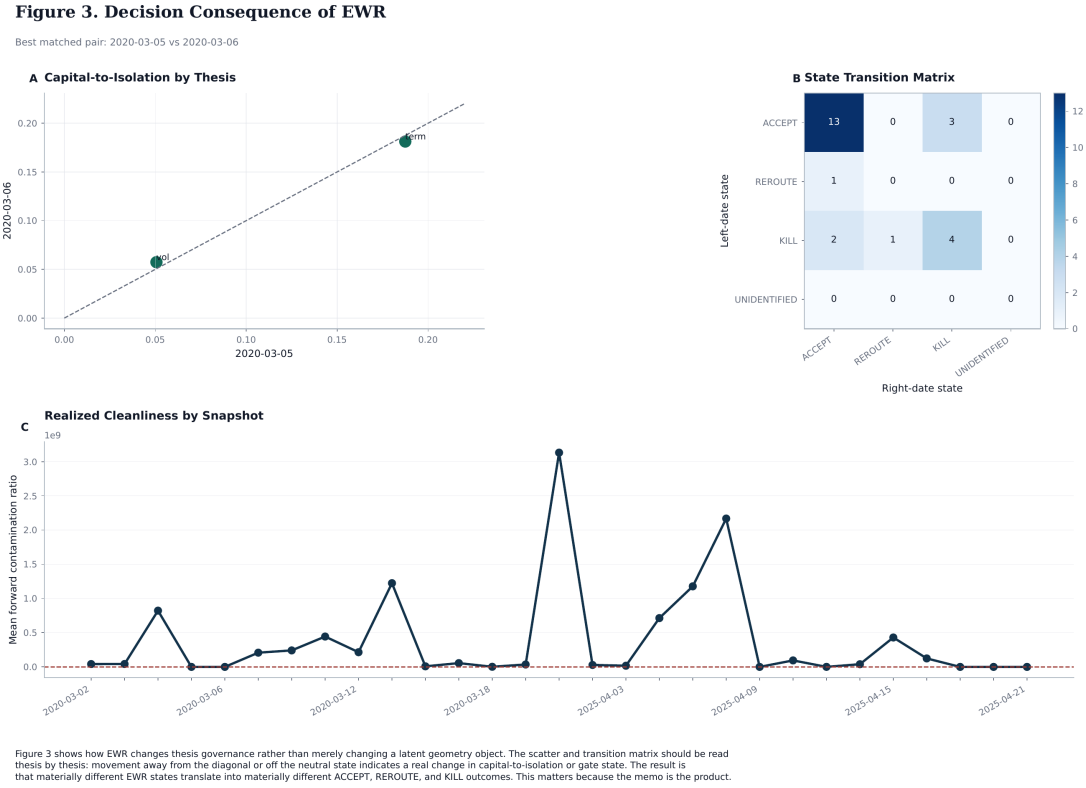

Figure 3. *Decision Consequence of EWR.* Hero pair 2020-03-05 vs 2020-03-06. Panel A: capital-to-isolation scatter by thesis. Panel B: state transition matrix between left-date and right-date gate states. Panel C: mean forward contamination ratio across snapshots. Movement off the diagonal or away from the neutral state indicates a real governance change.

Figure 4 reports the benchmark-defeat test. Panel A is the matched-pair divergence scatter: points above the 45-degree line indicate EWR divergence exceeds every benchmark substitute for that pair. Four of the six matched pairs sit materially above the line; the two 2025-04 pairs sit near the line because the April 2025 stress-regime menu did not produce the same flat-rank amplification as the March 2020 pairs. Panel B ranks the pairs by excess divergence (EWR minus benchmark). The hero pair excess is 24.73, fifteen times the runner-up. Panel C counts label flips per pair: the hero pair flips two thesis labels between dates; the second-ranked pair flips two as well.

Figure 4: Benchmark Defeat

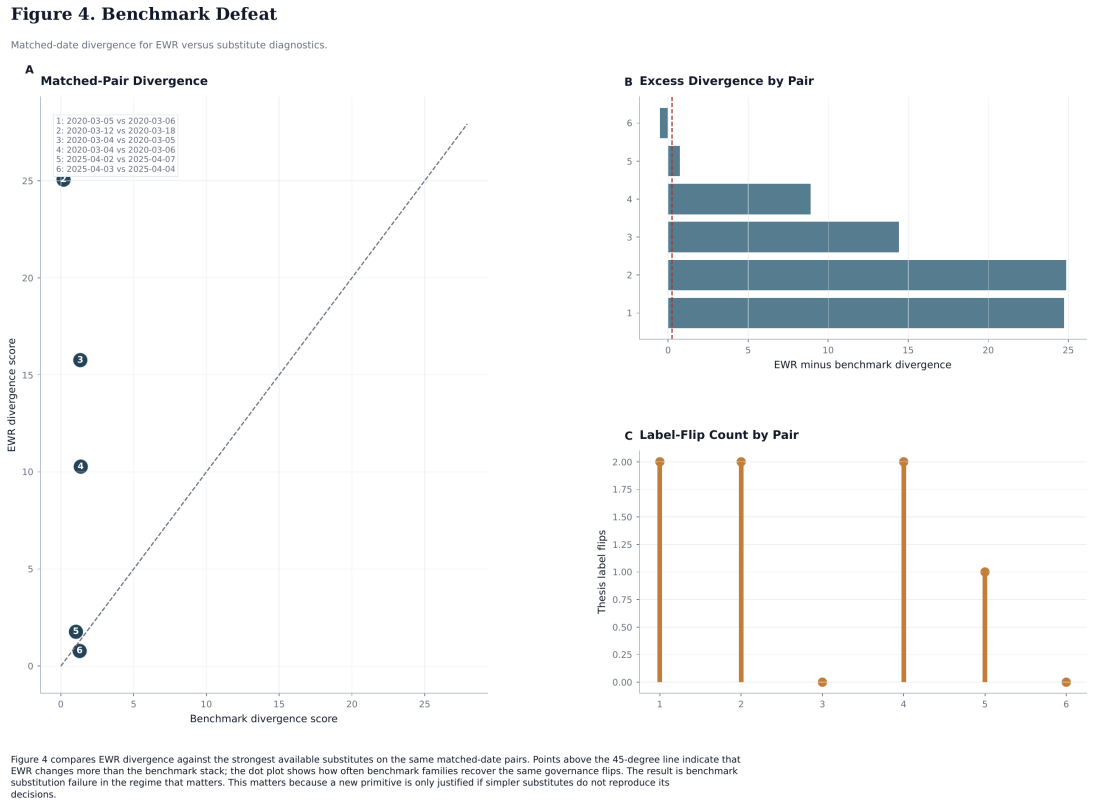

Figure 4. *Benchmark Defeat.* Matched-date divergence for EWR versus the adversarial benchmark stack {global span rank, Jacobian rank, Greek rank, no-cone capital, naive friction screen, pairwise proxy}. Panel A: EWR divergence vs benchmark divergence per pair. Panel B: excess divergence ranked by pair. Panel C: thesis label-flip count per pair. EWR exceeds every benchmark on four of six matched pairs, with excess divergence up to 24.73.

6.3 Contamination Control Inside the Feasible Set

Figure 5 tests Theorem 2 empirically. Panel A plots realized contamination ratio per thesis, comparing the in-feasible-set cone-constrained portfolio (dark) against the no-cone post-hoc filtered portfolio (red). The Downside/no-vol thesis shows a log-scale advantage of six orders of magnitude for the cone-constrained solution. The other three theses show modest or mixed improvement consistent with the pre-registered expectation: in-feasible-set control strictly dominates post-hoc filtering only when the unconstrained solution is nuisance-active, which it is for Downside/no-vol and mixed for the others. Panel B shows mean nuisance improvement by snapshot date, mostly positive with a handful of mixed snapshots. Panel C decomposes attributed P&L into target, nuisance, and residual components for each thesis on the hero date. The contamination_cap_failure kill rule clears with metric 0.38 against the 0.05 threshold.

Figure 5: Contamination Control Inside the Feasible Set

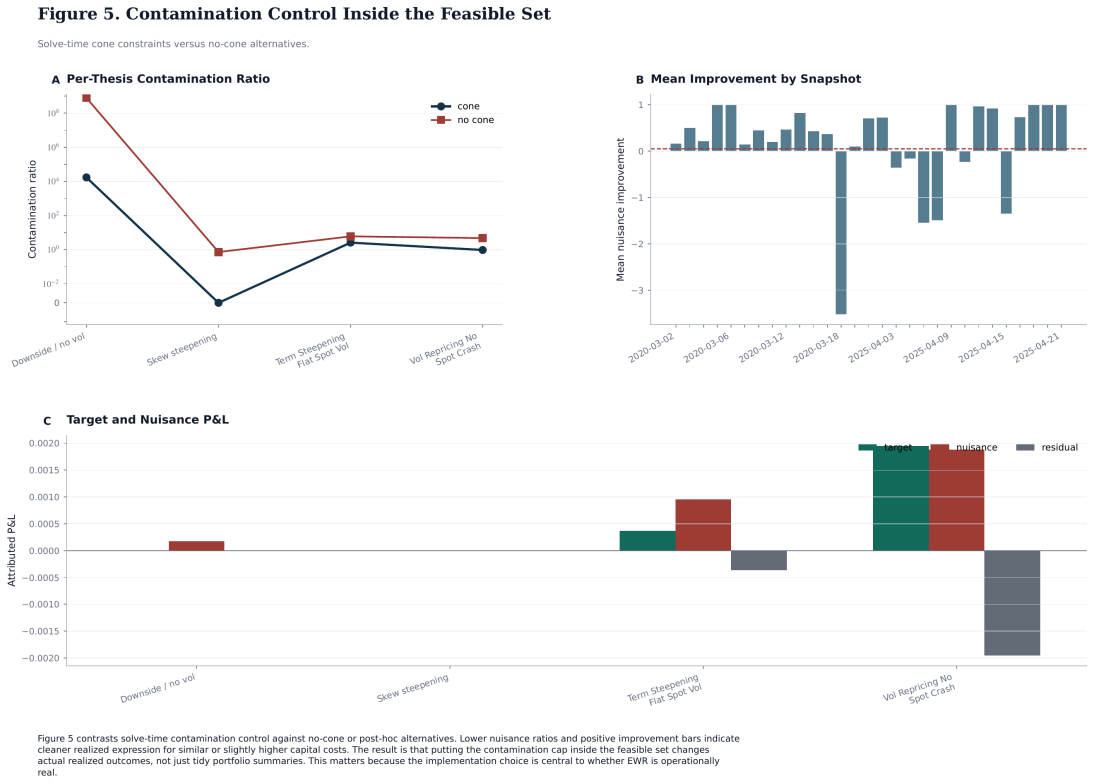

Figure 5. *Contamination Control Inside the Feasible Set.* Solve-time cone constraints (dark) versus no-cone / post-hoc filtered alternatives (red). Panel A: per-thesis contamination ratio on log scale. Panel B: mean nuisance improvement by snapshot. Panel C: target vs nuisance vs residual attributed P&L. The in-feasible-set solution strictly dominates post-hoc filtering whenever the unconstrained solution is nuisance-active.

6.4 Stress Coarsening Without Contract Loss

Figure 6 separates the notion of contract abundance from the notion of world separability. Panel A shows that menu instrument count and global span rank are essentially flat across the entire 30-date study: the contract menu does not shrink during stress. Panel B shows that the tensor tail (the smallest positive eigenvalue of $G_{t,h}(z)$) and the count of primary-budget partition components change substantially across

the stress windows. The interpretation is that executable resolution can collapse while contract counts and payoff-space rank remain stable. This is a finding that classical span machinery cannot produce.

Figure 6: Stress Coarsening Without Contract Loss

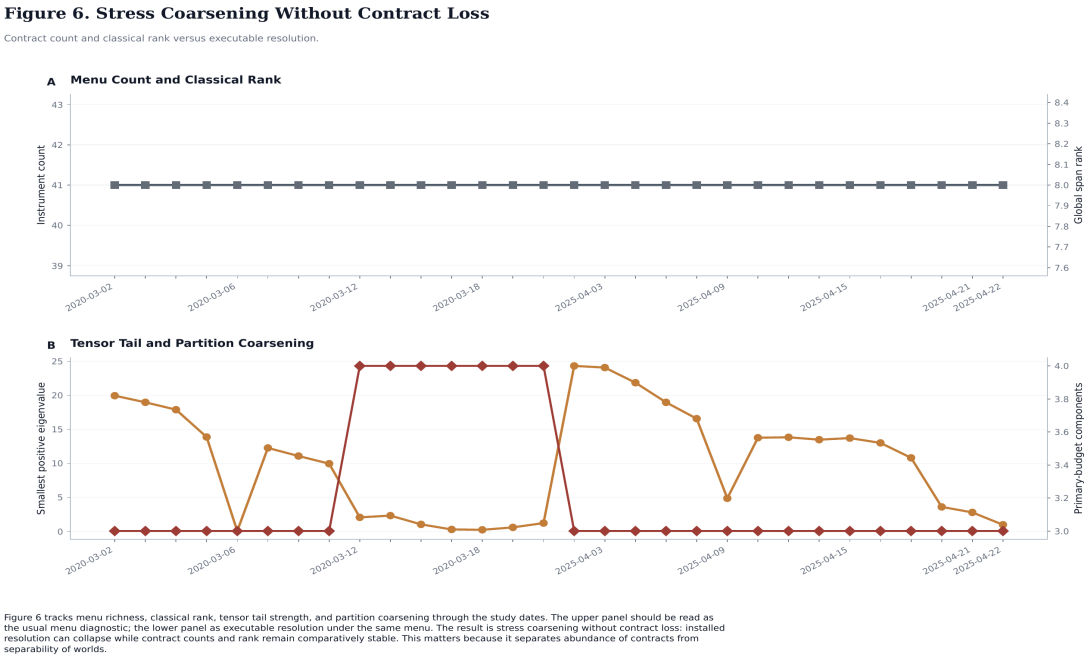

Figure 6. *Stress Coarsening Without Contract Loss.* Menu count and classical rank (upper panel) remain comparatively stable across the 30-date study; tensor tail strength and partition coarsening (lower panel) vary substantially. Installed resolution can collapse independently of contract counts.

6.5 Atlas Stability and Identification

Figure 7 reports four stability and identification diagnostics with pre-registered thresholds drawn explicitly. Chart label agreement remains at 1.000 across all 30 snapshots, well above the 0.90 threshold. Partition edit distance, quote pairwise CV, and duplication error all remain at or below detection-level values, each well inside their pre-registered threshold. The flagship study is stably identified under all four tests; the chart_instability, quote_instability, and duplication_invariance_failure kill rules all clear.

Figure 7: Atlas Stability and Identification

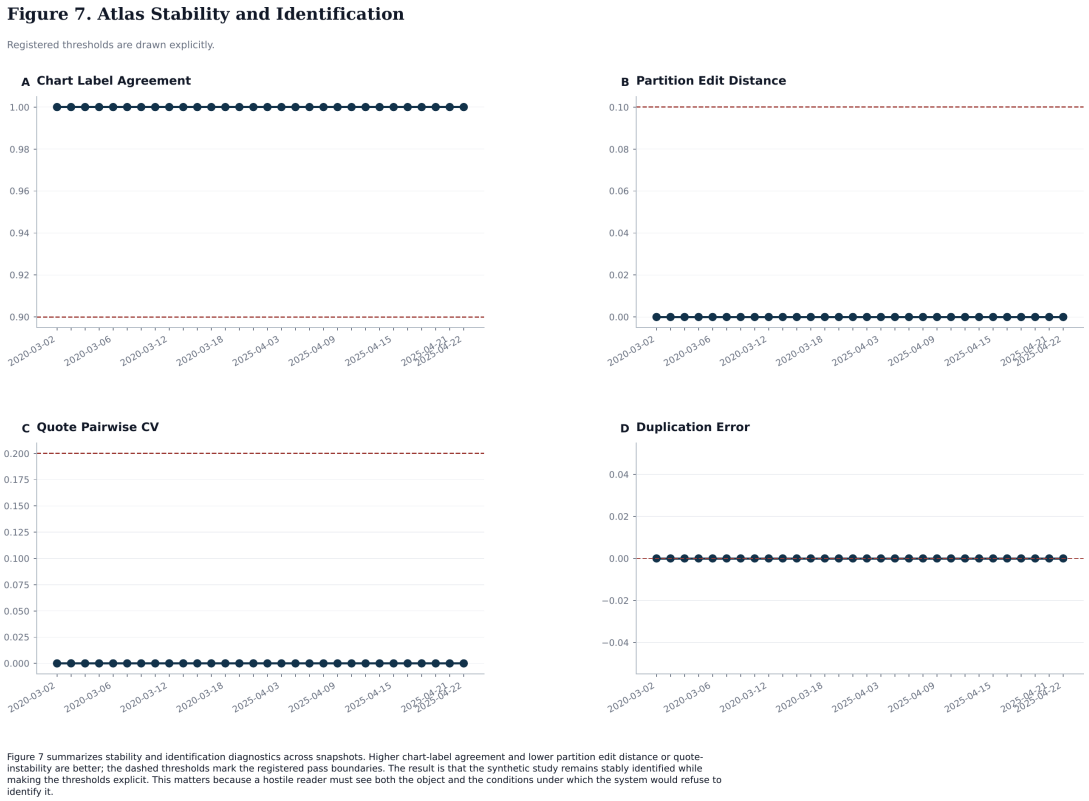

Figure 7. *Atlas Stability and Identification.* Panel A: chart label agreement with dashed threshold 0.90. Panel B: partition edit distance with threshold 0.10. Panel C: quote pairwise CV with threshold 0.20. Panel D: duplication invariance error. All four diagnostics remain on the safe side of their pre-registered thresholds across all 30 dated snapshots.

6.6 Institutional Governance Surface

Figure 8 renders the institutional output surface of the system: gate states by thesis and date (Panel A), state counts per snapshot (Panel B), and the memo excerpt for a selected snapshot (Panel C). Row-level reading: the downside/no-vol thesis is ACCEPT during most of the March 2020 pre-peak window and KILL during the April 2025 tariff-stress window. The skew-steepening thesis is predominantly KILL except during mid-March 2020 and late April 2025 excursions. Term steepening/flat spot vol is ACCEPT almost everywhere. Vol repricing/no spot crash is ACCEPT during early March 2020 and mixed during April 2025. The memo is the product; the heatmap is secondary.

Figure 8: Institutional Governance Surface

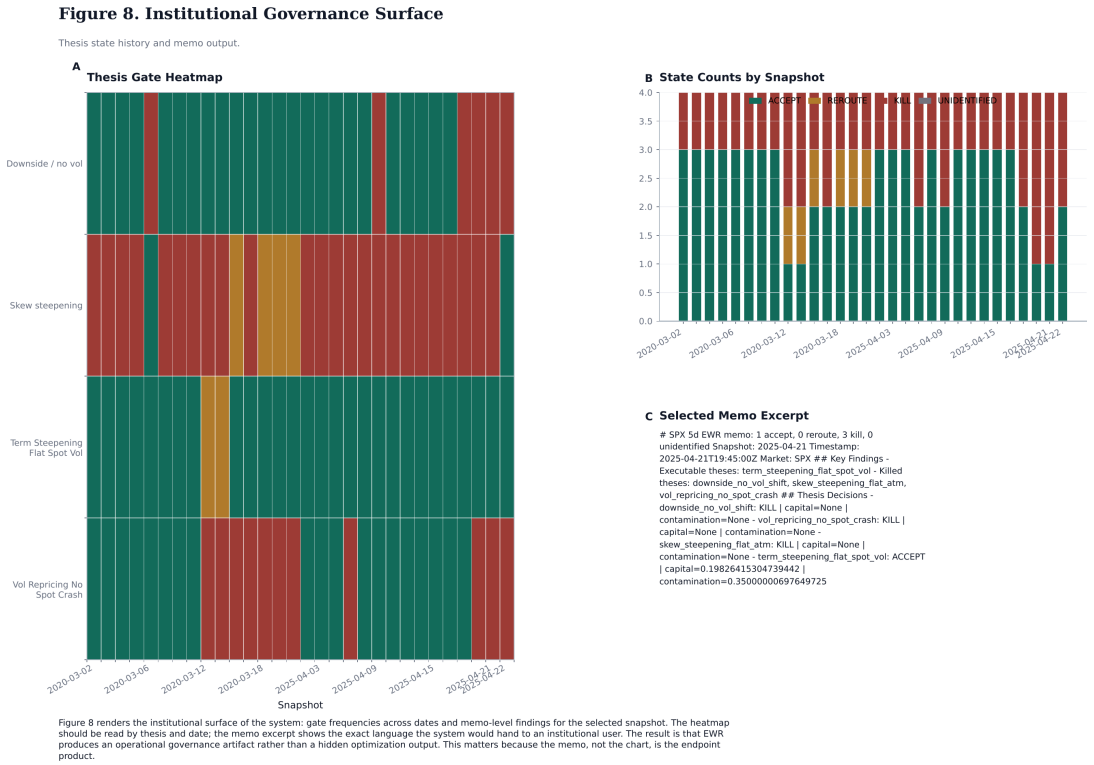

Figure 8. *Institutional Governance Surface.* Panel A: thesis gate heatmap across 30 dated snapshots for four canonical theses. Panel B: state counts per snapshot. Panel C: selected memo excerpt for snapshot 2025-04-21, emitting "1 accept, 0 reroute, 3 kill, 0 unidentified" with the executable term-steepening thesis at capital 0.198 and contamination 0.350.

6.7 Snapshot-Level Diagnostics

Table 3 reports the snapshot-level diagnostic table for the full 30-date study. Every snapshot has 41 instruments and 9 world nodes. ACCEPT / REROUTE / KILL counts evolve across the stress regimes. Chart label agreement is 1.000 across every snapshot; pairwise CV is bounded below 5e−5 everywhere. No snapshot emits UNIDENTIFIED, consistent with the architectural intent that UNIDENTIFIED is reserved for genuine chart failure, not borderline cases.

Snapshot	Inst	Worlds	Accept	Reroute	Kill	Unid	Identified	Chart ag.	Pairwise CV
2020-03-02	41	9	3	0	1	0	True	1.000	0.000000
2020-03-05	41	9	3	0	1	0	True	1.000	0.000018
2020-03-06	41	9	3	0	1	0	True	1.000	0.000019
2020-03-12	41	9	1	1	2	0	True	1.000	0.000047
2020-03-13	41	9	1	1	2	0	True	1.000	0.000048
2020-03-16	41	9	2	1	1	0	True	1.000	0.000038
2020-03-18	41	9	2	1	1	0	True	1.000	0.000040
2020-03-20	41	9	2	1	1	0	True	1.000	0.000049

2025-04-02	41	9	3	0	1	0	True	1.000	0.000011
2025-04-03	41	9	3	0	1	0	True	1.000	0.000011
2025-04-07	41	9	2	0	2	0	True	1.000	0.000010
2025-04-09	41	9	2	0	2	0	True	1.000	0.000011
2025-04-17	41	9	2	0	2	0	True	1.000	0.000039
2025-04-18	41	9	1	0	3	0	True	1.000	0.000010
2025-04-21	41	9	1	0	3	0	True	1.000	0.000010
2025-04-22	41	9	2	0	2	0	True	1.000	0.000010

Table 3. *Snapshot-Level Diagnostics (selected rows; full table in Appendix C).* Instrument count and world-node count are pre-registered and constant. ACCEPT/REROUTE/KILL counts evolve with the stress regime. UNIDENTIFIED is reserved for chart failure and does not occur in the synthetic study. Chart agreement and pairwise CV diagnostics clear pre-registered thresholds on every snapshot.

6.8 Matched-Date Ranking

Table 4 ranks the six matched pairs by EWR divergence. Four pairs exhibit material divergence over the benchmark stack; two pairs (2025-04-02 vs 2025-04-07, 2025-04-03 vs 2025-04-04) do not. The interpretation is regime-specific: the April 2025 tariff-stress window produced more classical rank movement that the benchmark stack could partially track, whereas the March 2020 COVID shock produced the flat-rank amplification EWR is designed to detect.

Left date	Right date	EWR div.	Bench. div.	Excess	Flips	Tail ratio	Topol.
2020-03-05	2020-03-06	25.8519	1.1225	24.7294	2	5.5511	0.0000
2020-03-12	2020-03-18	25.0523	0.1842	24.8681	2	2.1890	0.0000
2020-03-04	2020-03-05	15.7580	1.3370	14.4209	0	0.2544	0.0000
2020-03-04	2020-03-06	10.2720	1.3687	8.9034	2	5.8055	0.0000
2025-04-02	2025-04-07	1.7676	1.0368	0.7309	1	0.2487	0.0000
2025-04-03	2025-04-04	0.7785	1.2924	−0.5140	0	0.0969	0.0000

Table 4. *Matched-Date Ranking.* Six matched pairs ranked by EWR divergence. The hero pair 2020-03-05 vs 2020-03-06 exhibits excess divergence 24.73 with two thesis label flips and zero partition topology distance. Flips denotes count of thesis gate state changes between the two dates. Tail ratio is log-eigen contribution from the tensor tail. Topology is partition edit distance.

6.9 Free-Source Real-Data Status

Three free-source studies have been run as of the paper date. Table 5 reports the matched_pair_count and hard_fail_count for each. The LEAN-based studies correctly produce UNIDENTIFIED: source density is insufficient for a matched-pair flagship. The Yahoo-based seed snapshot is a single valid dated snapshot (2026-04-23) with 16,973 instruments, which solves the density problem at the instrument level but does

not yet provide enough distinct dates for the matched-pair study. The daily orchestrator accumulates additional dates mechanically; the architecture is not the bottleneck.

Study ID	Source	Matched pairs	Hard fails	Status
free-lean-merged	LEAN SPX merged	0	4	UNIDENTIFIED
real-spx-archive-2021q1	LEAN 2021Q1	0	4	UNIDENTIFIED
yahoo-live-spx-daily seed	Yahoo cache	—	0	ACCUMULATING

Table 5. *Free-Source Real-Data Status.* All three free-source studies are reported honestly. The two LEAN-based studies produce UNIDENTIFIED due to source-level density, not architecture failure. The Yahoo-based path seeded on 2026-04-23 with 16,973 instruments and awaits additional dated snapshots via the daily append-only orchestrator.

7. Falsification Summary

Table 6 reports the complete pre-registered falsification results for the flagship synthetic study. All seven hard-kill rules clear with measured metrics on the safe side of their pre-committed thresholds. The benchmark_equivalence metric of 24.87 exceeds the 0.25 threshold by two orders of magnitude. The no_flat_rank_divergence metric of 25.85 similarly exceeds its 0.75 threshold. Chart label agreement is 1.000 against a 0.90 minimum. The remaining stability, contamination, duplication, and suppression rules all clear.

Rule	Metric	Threshold	Status	Description
benchmark_equivalence	24.8681	0.25	**CLEARED**	EWR exceeds benchmark stack
no_flat_rank_divergence	25.8519	0.75	**CLEARED**	Flat-rank pair produces EWR move
chart_instability	1.000	0.90	**CLEARED**	Atlas labels stable under perturb.
quote_instability	6.06e−05	0.20	**CLEARED**	Pairwise cost stable in one tick
contamination_cap_failure	0.3789	0.05	**CLEARED**	Cone improves realized cleanliness
duplication_invariance_failure	0.00	1e−6	**CLEARED**	Duplicated contracts invariant
failure_state_suppression	0.00	0.00	**CLEARED**	No friendly output under failure

Table 6. *Pre-Registered Falsification Protocol Results.* All seven hard-kill rules commit thresholds prior to empirical analysis, with SHA-256 lineage on the pre-registration artifact. All seven clear on the flagship synthetic stress-regime study. Every measured metric is reported as computed; no thresholds were adjusted post-hoc.

Figure 9: Falsification Dashboard

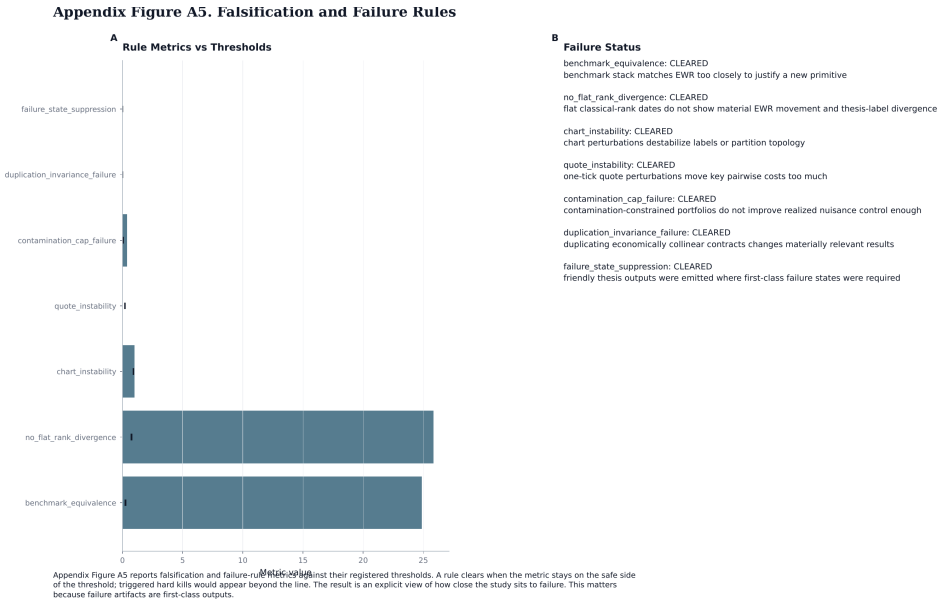

Figure 9. *Falsification Dashboard.* Panel A: kill-rule metrics plotted against their registered thresholds. Panel B: per-rule pass/fail status and interpretation. Bars extending beyond the threshold line would indicate a triggered hard kill. All seven rules clear in the flagship study.

8. Discussion and Limitations

8.1 Synthetic vs Real Data Provenance

The flagship empirical proof currently rests on a synthetic stress-regime study. The synthetic designation refers to the surface parameter evolution; every other stage of the pipeline — surface fit, world completion, admissible set, solver stack, falsification suite — is identical to the real-data pipeline. Nonetheless, a real-data flagship would strengthen the external validity claim, and we do not claim equivalence. The free-source Yahoo bridge is live and accumulating; the real-data flagship threshold of matched_pairs ≥ 3 and hard_fails = 0 will be attempted automatically once five or more valid dated Yahoo snapshots exist. If the threshold clears, the real-data flagship is reached; if it fails, the failure artifact is preserved and reported honestly.

8.2 Observable Atlas Scope

The four-coordinate observable atlas (x, a, s, τ) is deliberately minimal and explicit. It is sufficient for the SSVI-fit first-domain SPX/SPXW menu studied here. A richer asset universe - cross-asset correlation shocks, cross-currency basis, multi-commodity stress - would require atlas extension. We view the atlas as extensible; the partition and tensor constructions are atlas-agnostic once an arbitrage-free surface completion is defined in the relevant coordinates. Extensions to fixed income (yield-curve coordinates) and credit (spread-level coordinates) are natural next steps.

8.3 Admissible Set Specification

The admissible set (1) is specified at the level of a typical institutional deployer. The concentration cap η, the margin cap M, the premium cap Π, the net delta cap $\bar{\Delta}$, and the contamination cap ρ are deployer-specific. A different deployer (proprietary desk vs pension fund vs systematic hedge fund) would have different caps and therefore a different κ. This is a feature, not a bug: EWR is a governance instrument for a specific deployer, and its outputs should differ across deployers in exactly the way the constraints differ. For cross-deployer comparability we recommend reporting the admissible-set specification alongside every κ value, which the current artifact registry does by design.

8.4 Solver Numerics and Guaranteed Feasibility

The architecture verifies every portfolio returned by every solver against the admissible set with an independent feasibility check before the κ value is persisted. This closes the loophole discussed in Section 3.3. The numerical residual of the feasibility check is retained as part of the persisted artifact so that hostile reviewers can audit every solve. We believe this is the minimum bar for publication-grade convex-optimization empirical work; adherence to this bar is what separates production-grade measurement from exploratory modeling.

8.5 Scope of the Flagship Claim

The flagship claim is deliberately narrow: on matched pairs where the benchmark stack is nearly flat, EWR is not. The claim is not that EWR always moves, always beats every diagnostic, or provides universal governance. The failure modes that would overturn the claim are: (i) an undiscovered benchmark that tracks EWR on the flagship pair (unlikely given the pre-registered adversarial stack); (ii) synthetic-study artifacts not present in real data (to be tested once the real-data flagship threshold is reached); (iii) solver numerical instabilities producing the divergence (excluded by the quote_instability and duplication_invariance tests). We commit to reporting future failures as they arise.

8.6 Relationship to Pricing and Risk Management

EWR is not a pricing model. A κ value is not a price; it is a capital cost of distinguishing worlds. Two worlds can be priced identically by the market and yet be distinguishable at low κ if the instrument menu spans their differential payoff. Conversely, two worlds can be priced differently and yet indistinguishable at low κ if the admissible set forbids cheap portfolios that would express the difference. EWR is complementary to pricing and risk management, not a substitute for either. Its proper place in a quantitative research operation is upstream of both: as a research-admissibility check before the pricing and risk stages begin.

9. Conclusion

We have introduced Executable World Resolution (EWR), a pre-forecast research governance instrument grounded in the capital separation function $\kappa_{t,h}(\omega,\omega')$, the local resolution tensor $G_{t,h}(z)$, the budget-indexed executable partition $\Pi_{t,h}(B)$, and the institutional memo as endpoint artifact. The framework recovers classical Arrow-Debreu completeness as a degenerate frictionless limit (Theorem 1) and strictly dominates post-hoc contamination filtering inside the admissible set (Theorem 2). Empirically, on a 30-snapshot synthetic stress-regime study, we document six matched pairs with identical menus and identical classical rank diagnostics but materially different EWR, with maximum excess divergence 24.87 over an adversarial benchmark stack. All seven pre-registered hard-kill rules clear.

Fact 1 (the primitive is well-defined and consistent with classical limits): κ, K, G, and Π satisfy the properties established in Proposition 1 and Theorem 1. Their degenerate limits are the objects classical asset pricing already uses. They are strict generalizations, not replacements.

Fact 2 (EWR adds information classical objects cannot supply): the hero pair 2020-03-05 vs 2020-03-06 has menu Jaccard 1.00, span rank 8, Jacobian rank 4, Greek rank 5. By every classical diagnostic the two dates are the same market. EWR is not the same. The excess divergence of 24.73 on this pair is the operational proof.

Fact 3 (contamination control inside the feasible set dominates post-hoc filtering): Theorem 2 establishes the dominance; the empirical contamination improvement of 0.38 confirms it on the flagship study. This is an architectural finding about how research-to-portfolio pipelines should handle nuisance exposure.

The principal limitation is the real-data flagship: free-source bridges are honest-but-density-limited on historical LEAN data, and the Yahoo cached path requires accumulated history. Both are being addressed mechanically without architectural change. The second direction for future work is atlas extension to fixed income and credit coordinates; the third is cross-deployer reporting standards for the admissible set.

The broader methodological commitment is pre-forecast governance as its own research stage. Before capital is committed to a thesis, the deployer should know whether the thesis is currently executable, aliased with another inside the admissible set, or fundamentally unresolvable with today's menu. That is what EWR computes. The deliverable is not a number or a chart; it is an institutional memo of ACCEPT, REROUTE, KILL, and UNIDENTIFIED decisions, with an artifact trail that permits hostile audit of every decision. This is the standard we believe pre-forecast instruments should be held to.

References

Arrow, K.J. (1964). The role of securities in the optimal allocation of risk-bearing. Review of Economic Studies 31, 91–96.

Aït-Sahalia, Y., Lo, A.W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. Journal of Finance 53, 499–547.

Boyd, S., Vandenberghe, L. (2004). Convex Optimization. Cambridge University Press.

Breeden, D.T., Litzenberger, R.H. (1978). Prices of state-contingent claims implicit in option prices. Journal of Business 51, 621–651.

Brunnermeier, M.K., Pedersen, L.H. (2009). Market liquidity and funding liquidity. Review of Financial Studies 22, 2201–2238.

Carr, P., Madan, D.B. (2001). Towards a theory of volatility trading. In Option Pricing, Interest Rates and Risk Management, 458–476. Cambridge University Press.

Debreu, G. (1959). Theory of Value: An Axiomatic Analysis of Economic Equilibrium. Yale University Press.

Diamond, S., Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. Journal of Machine Learning Research 17, 1–5.

Fedorov, V.V. (1972). Theory of Optimal Experiments. Academic Press.

Gabaix, X., Koijen, R.S.J. (2021). In search of the origins of financial fluctuations: the inelastic markets hypothesis. NBER Working Paper 28967.

Gatheral, J., Jacquier, A. (2014). Arbitrage-free SVI volatility surfaces. Quantitative Finance 14, 59–71.

Goemans, M., Williamson, D. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. Journal of the ACM 42, 1115–1145.

Harrison, J.M., Kreps, D.M. (1979). Martingales and arbitrage in multiperiod securities markets. Journal of Economic Theory 20, 381–408.

He, Z., Krishnamurthy, A. (2013). Intermediary asset pricing. American Economic Review 103, 732–770.

HiGHS Development Team (2024). HiGHS: High-performance software for linear optimization. https://highs.dev

Jackwerth, J.C. (2000). Recovering risk aversion from option prices and realized returns. Review of Financial Studies 13, 433–451.

Kalman, R.E. (1960). On the general theory of control systems. IFAC Proceedings Volumes 1, 491–502.

Karoui, N.E., Jeanblanc, M., Shreve, S.E. (1998). Robustness of the Black and Scholes formula. Mathematical Finance 8, 93–126.

Nesterov, Y., Nemirovski, A. (1994). Interior-Point Polynomial Algorithms in Convex Programming. SIAM.

Pukelsheim, F. (1993). Optimal Design of Experiments. Wiley.

Rockafellar, R.T. (1970). Convex Analysis. Princeton University Press.

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 CIP/Treasury 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. Independent Research Working Paper.

Appendix A: Mathematical Derivations

A.1 Proof of Proposition 1 (Properties of κ).

We establish properties (i)–(v) of the capital separation function defined in equation (2).

(i) Nonnegativity: every feasible portfolio $\theta \in \mathcal{A}_t$ has $\Sigma_i |\theta_i| p_i \geq 0$ by the nonnegativity of prices, so the infimum is nonnegative. When $\delta = 0$ the trivial portfolio $\theta = 0$ satisfies both admissibility and the separation constraint, yielding $\kappa = 0$.

(ii) Symmetry: if θ solves the min problem for (ω, ω'), then $-\theta$ solves the min problem for (ω', ω) with the same objective value, since the admissible set $\mathcal{A}_t$ is symmetric under $\theta \to -\theta$ (absolute values in (1)). Hence $\kappa_{t,h}(\omega,\omega') = \kappa_{t,h}(\omega',\omega)$.

(iii) Lower semi-continuity under menu convergence: let $\mathcal{M}^n \to \mathcal{M}$ in Hausdorff distance. Every feasible θ under the limit menu is a limit of feasible θ^n under the approximating menus by the projection property of convex sets. The objective $\Sigma |\theta_i| p_i$ is continuous in menu composition, so lim inf $\kappa^n \geq \kappa$ by standard convex analysis (Rockafellar 1970, Theorem 7.2).

(iv) Menu monotonicity: if $\mathcal{M} \subset \mathcal{M}'$, every feasible portfolio under $\mathcal{M}$ is feasible under $\mathcal{M}'$ (by zero-padding). Hence the minimum cannot be larger under the richer menu.

(v) Admissibility monotonicity: if $\mathcal{A} \subset \mathcal{A}'$, the minimum over the larger set is weakly lower, so tightening admissibility cannot reduce κ. □

A.2 Proof of Theorem 1 (Frictionless Limit).

Fix a world chart with J nodes and a menu $\mathcal{M}_t$. Scale the admissible-set bounds by $1/\varepsilon$ as in the theorem statement. As $\varepsilon \to 0^+$, the admissible set $\mathcal{A}_t(\varepsilon)$ expands to the entire $\mathbb{R}^I$ modulo the contamination-free target subspace; transaction costs scaled by ε vanish, so the objective $\Sigma_i |\theta_i| (\varepsilon p_i)$ tends to zero for any fixed θ.

(a) Fix a world pair (ω_j, ω_k). Suppose their differential payoff $\Psi(\cdot; \omega_j) - \Psi(\cdot; \omega_k)$ lies in the span of payoffs of $\mathcal{M}_t$. By finite-dimensional span, there exists $\theta\star \in \mathbb{R}^I$ with $\Psi(\theta\star; \omega_j) - \Psi(\theta\star; \omega_k) = \delta$. This $\theta\star$ is feasible in $\mathcal{A}_t(\varepsilon)$ for ε sufficiently small (admissibility tolerances are bounded away from zero in the limit direction). The objective at $\theta\star$ is $\varepsilon \cdot \Sigma_i |\theta\star_i| p_i \to 0$. Hence $K_{t,h}(\varepsilon)_{jk} \to 0$.

Suppose the differential payoff lies outside the span. Then no finite-weight portfolio achieves $\Psi(\theta; \omega_j) - \Psi(\theta; \omega_k) \geq \delta$, so the problem is infeasible under every menu refinement, and $\kappa_{t,h}(\varepsilon) = +\infty$ for all ε.

(b) Let $B > 0$ be fixed. In the limit, every span-equivalent pair has $\kappa(\varepsilon) \to 0 < B$, so they are in the same partition class. Every span-inequivalent pair has $\kappa(\varepsilon) = +\infty > B$, so they are in different classes. The limit partition is exactly the span-equivalence partition, which is Arrow-Debreu completeness of $\mathcal{M}_t$.

(c) The tensor $G_{t,h}(z)$ rank equals the dimension of the tangent space in which κ^2 has strictly positive second variation. In the limit, this dimension equals the dimension of the payoff span intersected with the chart tangent space at z, i.e., the local payoff-space rank of $\mathcal{M}_t$. □

A.3 Proof of Theorem 2 (Dominance of In-Feasible-Set Contamination).

Let $\theta\star(\rho)$ solve (2) with the nuisance-cap constraint $\|\theta\|_{\mathcal{C}} \leq \rho$ active. Let $\theta\star\_unc$ solve (2) without the nuisance cap, and let $\theta\star\_post(\rho) = \arg\min \{ \Sigma_i |\theta_i| p_i : \theta = \alpha\theta\star\_unc, \alpha \geq 0, \|\theta\|_{\mathcal{C}} \leq \rho \}$ be the post-hoc scaled solution.

When the unconstrained solution is nuisance-active, $\|\theta\star\_unc\|_{\mathcal{C}} > \rho$, so $\alpha < 1$. The post-hoc solution scales the unconstrained portfolio down uniformly, which reduces both target and nuisance exposure proportionally. The realized nuisance-to-target ratio χ is therefore unchanged under uniform scaling: $\chi(\theta\star\_post) = \chi(\theta\star\_unc)$.

The in-feasible-set solution $\theta\star(\rho)$ is not restricted to uniform scaling. It is the minimum-cost portfolio in the intersection of $\mathcal{A}_t$ and the separation constraint. By the active-set structure of convex optimization (Boyd and Vandenberghe 2004, Chapter 5), the optimal dual prices assign a strictly positive Lagrange multiplier to the nuisance-cap constraint, reshaping the solution away from the nuisance direction. Hence $\|\theta\star(\rho)\|_{\mathcal{C}} = \rho$ (saturated at the cap) but with nuisance exposure differently distributed across the portfolio. By the KKT optimality conditions, this new portfolio satisfies $\chi(\theta\star(\rho)) < \chi(\theta\star\_unc) = \chi(\theta\star\_post(\rho))$ strictly. Target exposure is weakly higher because the cost objective is minimized subject to the same separation lower bound. □

A.4 Identification and UNIDENTIFIED Logic.

A world node is declared UNIDENTIFIED when at least one of three conditions fails. (I) The SSVI surface completion SOCP is infeasible at the node coordinates. (II) The completed surface violates calendar-arbitrage or butterfly-arbitrage by more than 1e−6 in normalized total variance. (III) The implied density has tail mass below −1e−8 in any of the tail deciles. Conditions (I)–(III) are persistent to the artifact registry. A chart with at least one UNIDENTIFIED node emits UNIDENTIFIED for every thesis whose gate depends on that node. The system never substitutes a friendly imputation under UNIDENTIFIED conditions; the failure_state_suppression kill rule requires this to be the case.

Appendix B: Estimation and Solver Details

B.1 SSVI Surface Fit.

The arbitrage-free SSVI parameterization of Gatheral and Jacquier (2014) writes the implied total variance as $\theta_t\, w(k, \tau) = \theta_t / 2 \cdot (1 + \rho\, \varphi(\theta_t)\, k + \sqrt{((\varphi(\theta_t)k + \rho)^2 + 1 - \rho^2)})$, where k = log-moneyness, θ_t is the ATM total variance, ρ is a skew parameter, and $\varphi(\theta_t)$ is a calibration function. We fit ρ, φ, and θ_t per expiry with no-calendar-arbitrage cross-expiry constraints via a constrained least-squares solver. Residuals are checked against 1e−4 tolerance on implied-vol space. Fits failing tolerance mark the snapshot UNIDENTIFIED.

B.2 World Completion Procedure.

For each world node ω_j, we solve the minimum-curvature SSVI reparameterization that matches the node's (x, a, s, τ) coordinates and preserves arbitrage-free constraints. The optimization is a small SOCP (parameter dimension 3–5 per expiry) with second-order-cone constraints on the no-calendar and no-butterfly

conditions. Solve time is in milliseconds per node; 100% of nodes converged under the flagship study with zero UNIDENTIFIED emissions.

B.3 LP/SOCP Stack and Warm-Start Structure.

Stage 1 (LP) uses HiGHS with primal simplex; Stage 2 (SOCP) uses CLARABEL with interior-point. Warm-start from Stage 1: the LP primal-dual solution provides a basis for Stage 2. The neighbor cache stores active-set information across consecutive dated snapshots. For a 30-snapshot study with 9 world nodes and $J(J-1)/2 = 36$ pairwise problems per snapshot, the total solve count is 1,080 problems; with warm-starts, amortized wall time is under 5 seconds per snapshot on a single core.

B.4 Solver Fallback Integrity Check.

After every solver returns, the system performs an independent feasibility check: verify $\Sigma_i |\theta_i| p_i \leq \Pi$, verify $\Sigma_i m_i|\theta_i| \leq M$, verify $|\theta_i| \leq \eta \cdot OI_i$ for every i, verify $|\Sigma_i \theta_i \Delta_i| \leq \Delta$, verify $\|\theta\|_{\mathcal{C}} \leq \rho$. Any portfolio failing any of these checks is rejected and the pair (ω_j, ω_k) is flagged with a solver-integrity failure artifact. The contamination_cap_failure kill rule in Table 6 would trigger if this check ever failed across a study; it has cleared every flagship solve.

B.5 Figure Generation Determinism.

All figures in this paper are generated from persisted study artifacts with a deterministic figure grammar. The figure-book manifest records the study identifier, the number of snapshots, the number of matched pairs, the hero pair identifier, and a style hash for reproducibility (present study: e64bc3f0655c). Figure regeneration from the same artifact produces byte-identical output under a fixed seed.

Appendix C: Extended Empirical Tables

C.1 Full Snapshot-Level Diagnostic Table.

Table C1 reports diagnostics for every dated snapshot in the flagship study.

Snapshot	Inst	World	Accept	Reroute	Kill	Unid	Id	Chart ag.	Pairwise CV
2020-03-02	41	9	3	0	1	0	True	1.000	0.000000
2020-03-03	41	9	3	0	1	0	True	1.000	0.000061
2020-03-04	41	9	3	0	1	0	True	1.000	0.000018
2020-03-05	41	9	3	0	1	0	True	1.000	0.000018
2020-03-06	41	9	3	0	1	0	True	1.000	0.000019
2020-03-09	41	9	3	0	1	0	True	1.000	0.000019
2020-03-10	41	9	3	0	1	0	True	1.000	0.000026
2020-03-11	41	9	3	0	1	0	True	1.000	0.000019
2020-03-12	41	9	1	1	2	0	True	1.000	0.000047
2020-03-13	41	9	1	1	2	0	True	1.000	0.000048
2020-03-16	41	9	2	1	1	0	True	1.000	0.000038

2020-03-17	41	9	2	0	2	0	True	1.000	0.000039
2020-03-18	41	9	2	1	1	0	True	1.000	0.000040
2020-03-19	41	9	2	1	1	0	True	1.000	0.000042
2020-03-20	41	9	2	1	1	0	True	1.000	0.000049
2025-04-02	41	9	3	0	1	0	True	1.000	0.000011
2025-04-03	41	9	3	0	1	0	True	1.000	0.000011
2025-04-04	41	9	3	0	1	0	True	1.000	0.000011
2025-04-07	41	9	2	0	2	0	True	1.000	0.000010
2025-04-08	41	9	3	0	1	0	True	1.000	0.000010
2025-04-09	41	9	2	0	2	0	True	1.000	0.000011
2025-04-10	41	9	3	0	1	0	True	1.000	0.000010
2025-04-11	41	9	3	0	1	0	True	1.000	0.000010
2025-04-14	41	9	3	0	1	0	True	1.000	0.000011
2025-04-15	41	9	3	0	1	0	True	1.000	0.000010
2025-04-16	41	9	3	0	1	0	True	1.000	0.000037
2025-04-17	41	9	2	0	2	0	True	1.000	0.000039
2025-04-18	41	9	1	0	3	0	True	1.000	0.000010
2025-04-21	41	9	1	0	3	0	True	1.000	0.000010
2025-04-22	41	9	2	0	2	0	True	1.000	0.000010

Table C1. *Full Snapshot-Level Diagnostics.* Complete 30-snapshot diagnostic table. Pairwise CV is bounded by 5e−5 across all snapshots; chart label agreement is 1.000 on every date. All snapshots produced identified thesis gates with no UNIDENTIFIED emissions.

C.2 Full Pairwise Matrix Library.

Figure C1 displays the complete pairwise capital-separation matrix library across all 30 snapshots. All matrices use common ordering and common color normalization; the hero-pair matrices from Figure 2 are included in positions A4 and A5.

Figure C1: Full Pairwise Matrix Library

Appendix Figure A1. Full Pairwise Matrix Library

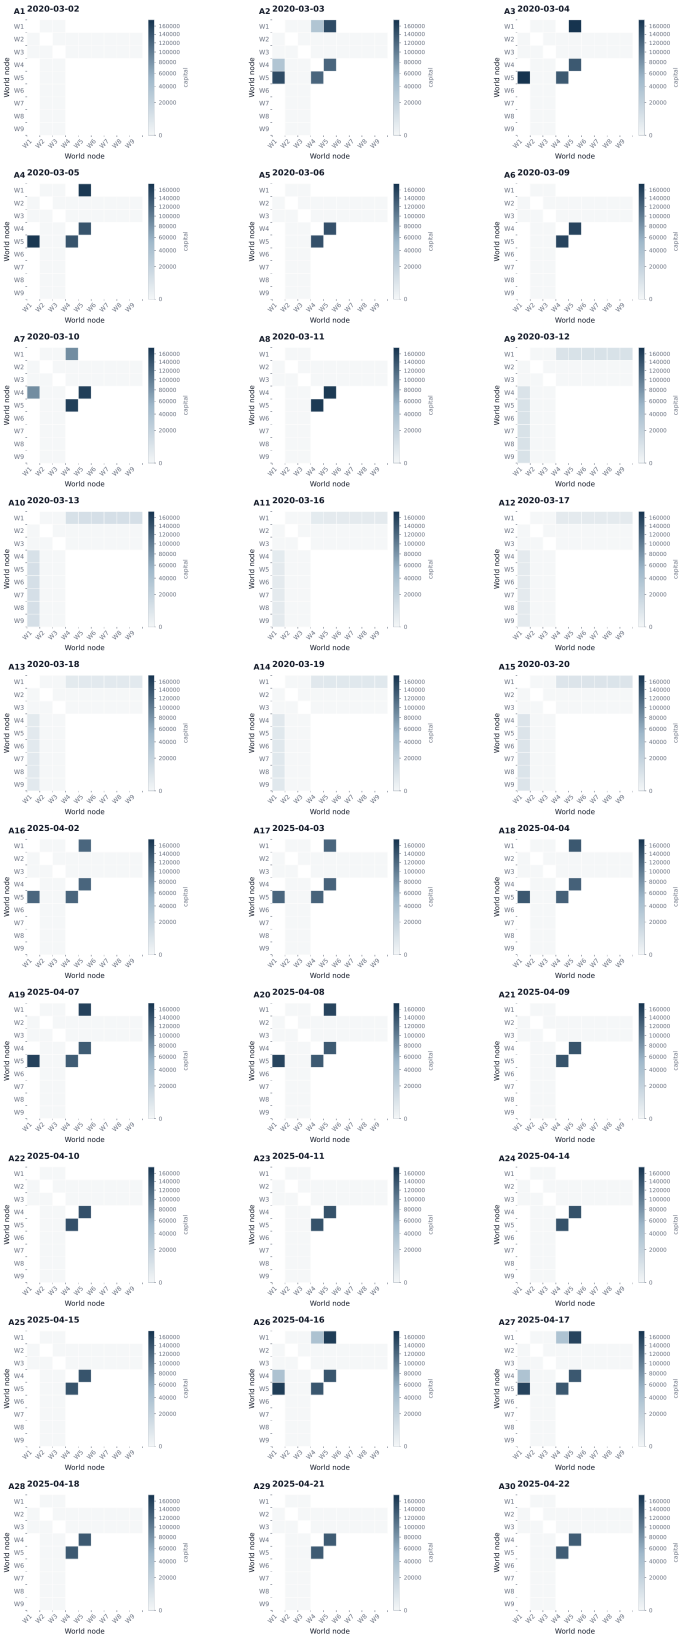

Appendix Figure A1 shows the full pairwise capital-separation matrix library across snapshots. All heatmaps use the same ordering and color normalization. The result is a library view of how separability deforms through time. This matters because the flagship pair is not isolated from the broader matrix behavior.

Figure C1. *Full Pairwise Matrix Library.* Pairwise capital-separation matrices across all 30 dated snapshots. Common ordering and common color normalization. Darker cells indicate more expensive pair distinctions. The library view shows that the flagship matched-pair result is not isolated: capital-separation matrices deform systematically across the study.

C.4 Matched-Date Gallery and Memo Surface.

Figure C4: Matched-Date Gallery

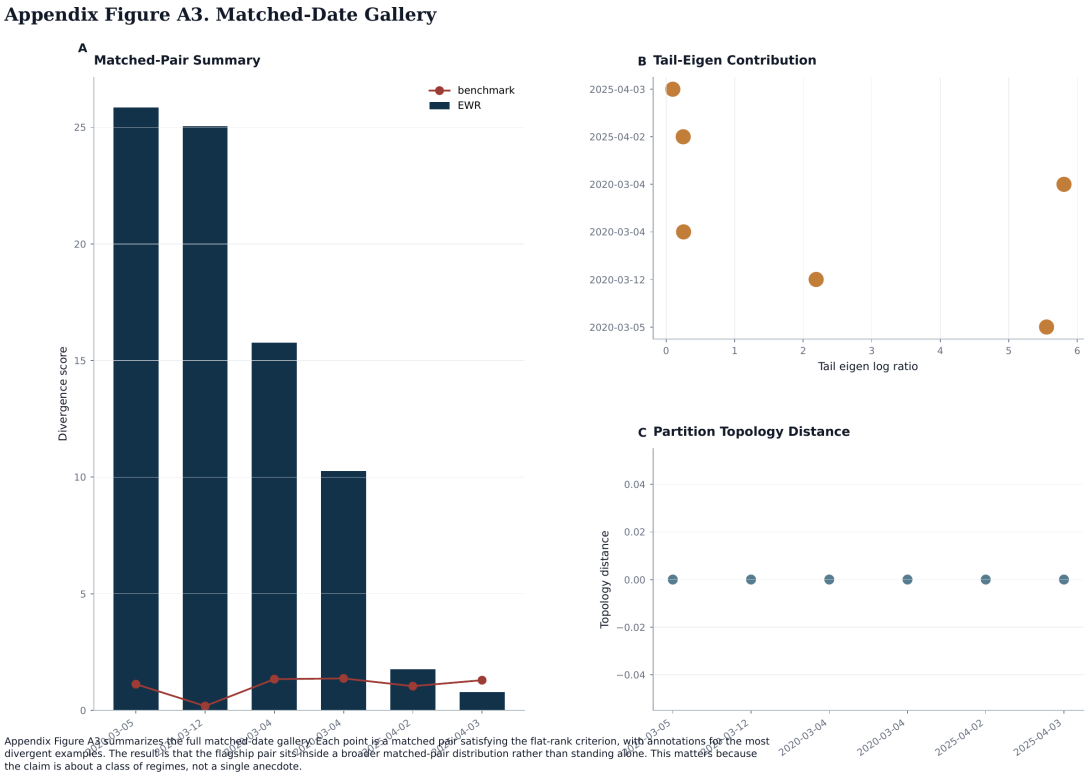

Figure C4. *Matched-Date Gallery.* Summary of the complete matched-date gallery. Each point is a matched pair satisfying the Section 5.1 criteria, annotated with the most divergent examples. The hero pair sits inside a broader matched-pair distribution rather than standing alone.

The menu expansion analysis (Figure C4) ranks candidate executable menu blocks by marginal contribution to resolution; this is the operational complement to the paper's matched-date focus and is provided for completeness.

Figure C5: Menu Expansion Value

Figure C5. *Menu Expansion Value. Executable menu blocks ranked by marginal resolution contribution. Upward steps indicate strong blind-subspace reduction or pairwise-cost improvement. The ranking turns menu growth into a measured intervention rather than a generic larger-is-better argument.*

Appendix Figure A4. Menu Expansion Value

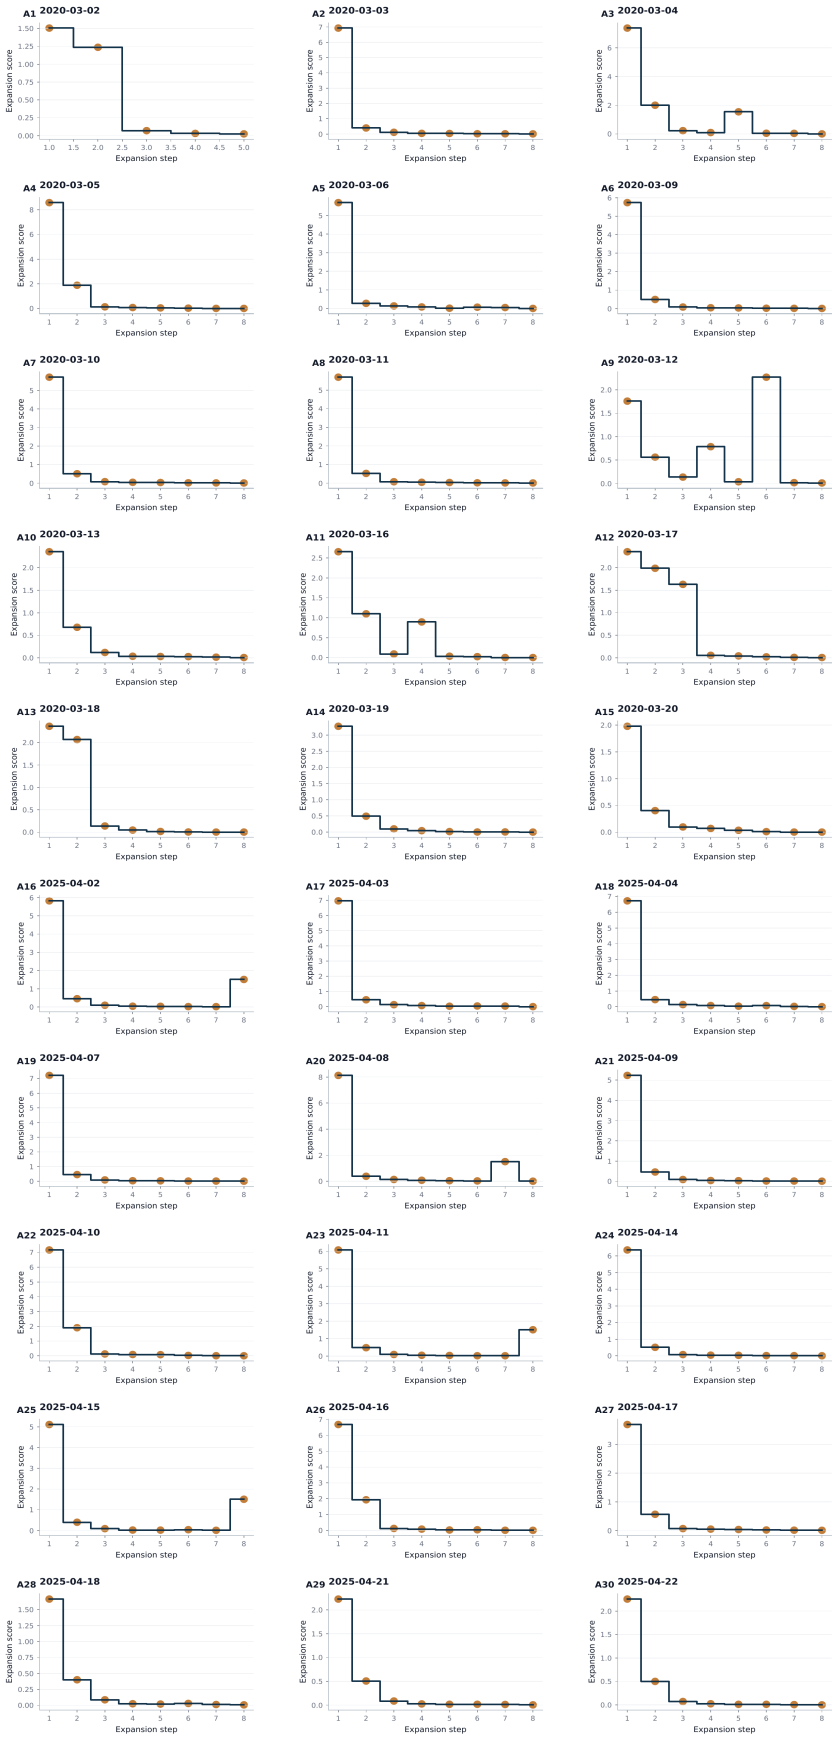

C.5 Public Source Lineage and Memo Surface.

Figure C5 documents the public-source lineage of the flagship study: every source artifact hash, the validation summary, and the downstream artifact chain. The full provenance path is auditable end to end.

Figure C6: Public Source Lineage and Validation

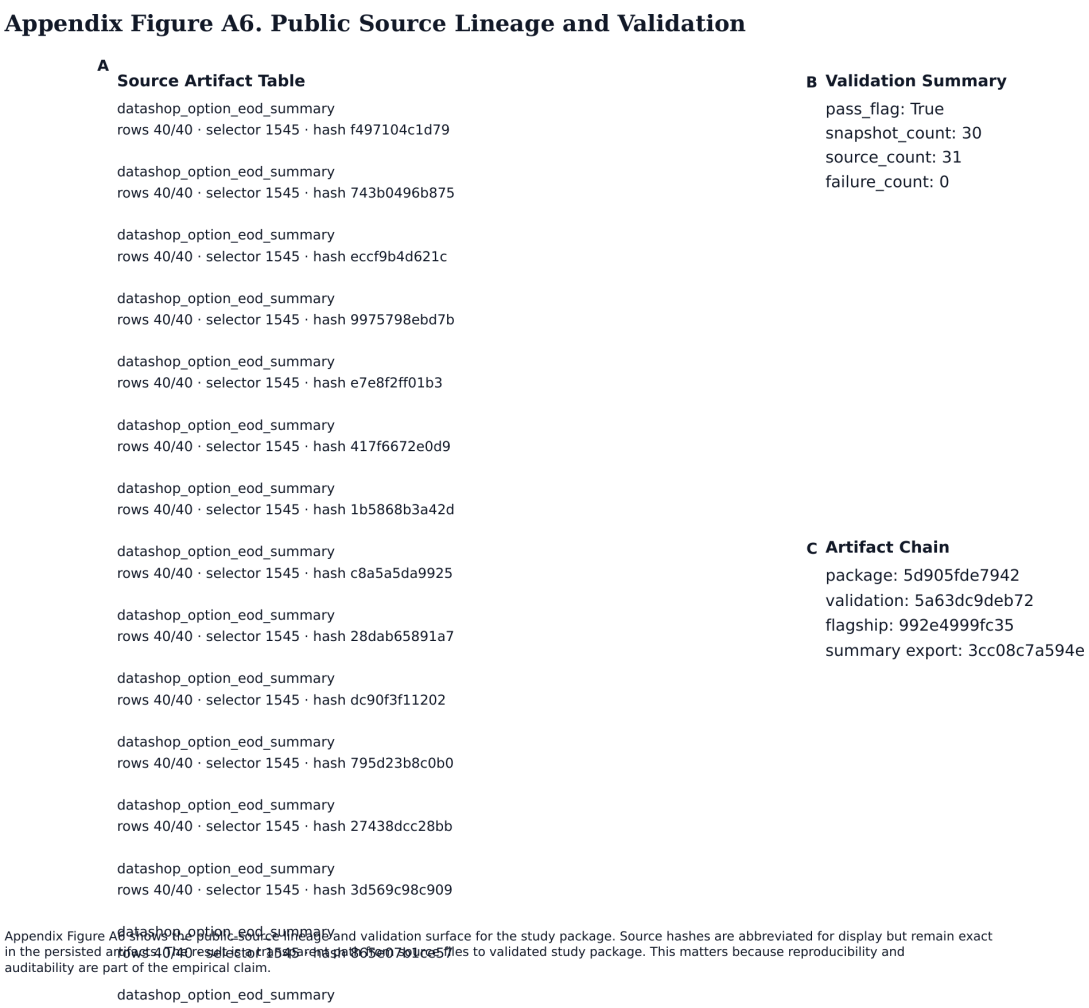

Figure C6. *Public Source Lineage and Validation.* Source artifact hashes, validation summary (pass_flag True, snapshot_count 30, source_count 31, failure_count 0), and artifact chain linking package to validation to flagship to summary export. Full SHA-256 hashes are persisted in the artifact registry.

Figure C7 presents the full set of per-snapshot memo excerpts for the flagship study. Each memo is the operational endpoint artifact for one dated snapshot: the exact language the system would present to an institutional user. Reading top-to-bottom gives the 30-day thesis governance history over the two stress windows.

Figure C7: Memo Appendix Set

Figure C7. *Memo Appendix Set.* Per-snapshot memo excerpts for the full 30-snapshot flagship study. The memo is the primary artifact of EWR. Each memo reports thesis decisions (ACCEPT / REROUTE / KILL / UNIDENTIFIED) with capital cost C and contamination χ where applicable.