

Article

Spectral Hypergraph Algorithms for Early Detection of Connectivity Collapse with Application to Pharmaceutical Supply Chain Arrest

Ntebogang Dinah Moroke 

Department of Statistics and Operations Research, Faculty of Economic and Management Sciences, North-West University, Private Bag X2046, Mafikeng 2745, South Africa; ntebo.moroke@nwu.ac.za

Highlights

- Weighted hypergraph representation of a pharmaceutical supply chain detects simultaneous multi-party failures that pairwise graph projections miss entirely.
- Geometry-aware spectral indicators achieve 19-day early warning lead time and 0.98 detection sensitivity; classical SPC achieves near-zero sensitivity (0.05) for network-geometric disruptions.
- COVID-19 lockdown consistency check: simulated service levels collapsed to 74.5% while forecast MAPE remained below 5%, confirming decoupling of accuracy from network integrity.
- Three-regime alert system (Green/Amber/Red) with automated greedy dejamming integrates directly with District Health Information System procurement dashboards.
- Northern Cape sub-network ($\lambda_2 \approx 0.08$) operates near arrest under normal conditions; spectral triage supports equity-weighted procurement under SDG 10.

Abstract

We propose a family of spectral hypergraph algorithms for early detection of connectivity collapse in pharmaceutical supply chain networks. The Fiedler eigenvalue λ_2 of the normalised hypergraph Laplacian serves as the order parameter. Five geometry-aware early warning indicators (TSI, HSST, HOMFA, HOTV, ORC) monitor network topology rather than scalar residuals, with provable detection guarantees under geometric ergodicity. A Greedy Dejamming algorithm restores connectivity via rank-2 Laplacian updates, achieving a $(1 - 1/e)$ -approximation within a procurement budget constraint. Monte Carlo validation on a calibrated pharmaceutical distribution hypergraph demonstrates substantially higher detection sensitivity and shorter lead times than classical statistical process control. Hyperedge representation yields detection gains exceeding 90% for simultaneous multi-party failures that pairwise graph projections miss entirely. A COVID-19 lockdown episode provides a held-out directional consistency check.



Academic Editor: Mauro Mezzini

Received: 31 May 2026

Revised: 22 June 2026

Accepted: 24 June 2026

Published: 3 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: hypergraph Laplacian; Fiedler eigenvalue; algebraic connectivity; bootstrap control chart; Monte Carlo simulation; spectral algorithm; greedy approximation; supply chain network; early warning algorithm; Markov chain; extreme value theory; network resilience

1. Introduction

Detecting connectivity collapse in time-varying weighted hypergraphs is a fundamental challenge in algorithm design with broad applications in network monitoring,

distributed systems, and supply chain resilience. The algebraic connectivity λ_2 of a weighted hypergraph's normalised Laplacian governs information propagation across the network [1,2]: as $\lambda_2 \rightarrow 0$, the network fragments into topologically isolated components from which no global coordination is possible [3]. Monitoring λ_2 in real time requires efficient spectral algorithms that scale to large sparse hypergraphs and provide statistically grounded detection guarantees, yet no existing algorithm addresses this problem with provable lead time and false alarm bounds.

Supply chain networks are an important class of systems where higher-order interactions arise naturally. Multi-lateral agreements, joint inventory pooling, and shared logistics routes involve the simultaneous coordination of multiple entities; manufacturers, distributors, wholesalers, dispensaries, and regulatory agencies; whose relationships cannot be faithfully captured by pairwise edges [4,5]. Weighted hypergraphs, where each hyper-edge connects an arbitrary number of vertices, provide the appropriate representation [5]. The South African public healthcare pharmaceutical distribution network exemplifies this structural complexity: a single dispensing cycle binds a national lead provider, cold-chain operators, cluster dispensaries, and patient cohorts under a single regulatory audit trail, with no pairwise decomposition preserving the joint accountability structure [6,7].

A central, yet underexplored, challenge arises when AI-driven decision support systems depend on the geometric integrity of such hypergraphs. Hybrid models combining Least Squares Support Vector Regression with Random Forest (LSSVR-RF) have demonstrated strong forecasting performance in pharmaceutical supply chains [6], and SHapley Additive exPlanations [8] are widely used to maintain interpretability. However, the reliability of SHAP attributions depends critically on the stability of the underlying network geometry. Standard monitoring protocols focus on forecasting error metrics such as MAPE, implicitly assuming that low error indicates a well-functioning system. The structural integrity of the underlying hypergraph can degrade severely without any detectable increase in local forecast errors, a phenomenon referred to as *network arrest* or *topological arrest*. When connectivity collapses, information propagation pathways fragment and SHAP attributions lose their geometric grounding, rendering policy decisions derived from them potentially invalid. Detecting this structural failure requires moving beyond point forecast evaluation to monitoring the spectral geometry of the hypergraph itself.

This paper introduces a family of spectral algorithms that address this gap. The algorithms monitor λ_2 of the normalised hypergraph Laplacian using bootstrap-calibrated control charts, detect connectivity collapse with provable lead time and false alarm guarantees, and restore connectivity via a $(1 - 1/e)$ -approximate greedy hyperedge addition strategy. The framework is validated on a pharmaceutical distribution hypergraph whose structural parameters are calibrated to eleven years of published CCMDD programme reports and a COVID-19 lockdown episode that provides a held-out directional test of detection performance on a simulated disruption trajectory calibrated to the known chronology of that event.

Existing spectral methods address eigenvalue computation for static hypergraphs [1,9] or dynamic graph streams [10] but do not provide early warning algorithms for connectivity collapse in weighted hypergraphs with bootstrap-calibrated null distributions. Bootstrap methods for eigenvalue statistics [11] and submodular maximisation theory [12] exist separately; this paper connects them through the Fiedler eigenvalue to yield a complete detection and recovery algorithm suite.

Contributions

This paper makes five algorithmic contributions:

1. **Connectivity collapse detection.** Three spectral algorithms (HSST, HOMFA, TSI) with bootstrap and Markov chain CLT null distributions, bootstrap consistency rate $\mathcal{O}_p(B^{-1/2})$, and provable detection guarantees under geometric ergodicity.
2. **Attribution stability monitoring.** Two geometry-aware indicators (HOTV, ORC) that certify whether SHAP-based AI decision support outputs remain valid under the current network geometry, connecting λ_2 to attribution variance via a closed-form trustworthiness metric.
3. **Greedy connectivity restoration.** A $(1 - 1/e)$ -approximate algorithm for budgeted hyperedge addition via rank-2 Laplacian updates derived from the Davis–Kahan theorem [13,14], running in $\mathcal{O}(|\Delta\mathcal{H}| \cdot n)$ per iteration.
4. **Theoretical foundations.** Three propositions grounding the algorithms in the Cheeger inequality [1,9] and the Kipnis–Varadhan Markov chain CLT [3,15].
5. **Empirical validation.** Monte Carlo validation on a calibrated synthetic pharmaceutical distribution hypergraph, with the COVID-19 lockdown period used as a held-out directional consistency check. Full validation details are in Section 3.6.

2. Literature Review

2.1. Pharmaceutical Supply Chain Disruption and Resilience

Pharmaceutical supply chains are structurally distinct from general supply chains in ways that matter for analytics design. Regulatory obligations, particularly Good Distribution Practice guidelines under SAHPRA and FDA frameworks, impose audit trail and explainability requirements on AI-driven procurement decisions that have no parallel in commercial logistics [16,17]. Cold-chain constraints create hard physical bounds on redistribution options that do not apply to ambient-temperature goods. Patient safety implications mean that service-level failures have clinical consequences beyond financial loss [7,18].

Three disruption typologies characterise the Southern African pharmaceutical logistics environment. Macroeconomic shocks, principally ZAR currency depreciation inflating active pharmaceutical ingredient import costs, operate over weeks to months and affect procurement budgets systematically [6]. Infrastructure shocks from load-shedding compromise cold-chain warehouse operations and vehicle routing directly, creating node-level supply failures [19]. Systemic shocks, exemplified by the COVID-19 port closures of 2020, sever multi-party supply corridors simultaneously, producing the topological failures that are the central concern of this paper [20].

Classical supply chain resilience frameworks respond to these disruptions through safety stock accumulation, lead time buffering, and supplier diversification [21,22]. These interventions are correct at the node level but structurally incomplete: they address the capacity of individual nodes without addressing the connectivity of the network through which supplies must flow. A warehouse stocked above the target service level cannot compensate for a severed supply corridor if the information system cannot route orders through the remaining hyperedges.

Gap 1: No existing pharmaceutical supply chain resilience framework monitors the structural connectivity of the distribution network as a leading indicator of service-level failure. All existing approaches, safety stock, supplier diversification, MAPE monitoring, operate on scalar outputs rather than network geometry.

2.2. Network Analytics and Hypergraph Methods in Supply Chain Management

The application of network science to supply chain management has a substantial foundation. Hearnshaw and Wilson [10] established the relevance of complex network topology to supply chain behaviour, and Choi et al. [23] framed supply networks as complex adaptive systems in which emergent behaviour arises from network structure rather than individual node decisions. Graph neural network approaches have improved disruption prediction in pairwise supply chain graphs [24], and blockchain–AI integration has been explored for pharmaceutical traceability [24]. These advances share a common structural limitation: they represent multi-party pharmaceutical supply contracts as bilateral edges, losing the joint accountability structure that governs both operational and regulatory behaviour [4,5].

Hypergraph representations address this limitation directly. Berge’s foundational treatment [4] establishes the mathematical basis for encoding multi-party relationships. Zhou et al. [5] developed learning algorithms for hypergraphs that demonstrated spectral methods extend naturally from graphs. Louis [1] proved the Cheeger inequality for hypergraph Laplacians, establishing the theoretical link between spectral gap and conductance. Chan et al. [9] extended this to spectral properties of the hypergraph Laplacian and showed that the operator admits efficient approximation algorithms, Chung’s spectral graph theory [2], extended to hypergraphs through the normalised Laplacian of Zhou et al., provides the algebraic connectivity framework this paper applies to pharmaceutical distribution.

Gap 2: No existing study applies hypergraph spectral theory to pharmaceutical distribution networks, and none connect algebraic connectivity to supply chain information system reliability. The theoretical maturity of hypergraph spectral methods has not been transferred to the operational context where it is most needed.

2.3. AI and Business Analytics for Pharmaceutical Logistics

Artificial intelligence has been applied extensively to pharmaceutical demand forecasting and inventory optimisation [18,25]. Hybrid architectures combining machine learning with operations research methods achieve strong predictive performance under stable conditions [6]. The consistent finding across this literature is that predictive accuracy degrades during major disruptions even when model architecture remains unchanged [18]. This degradation is treated as a statistical phenomenon to be managed through post hoc correction or model reestimation.

Explainability methods, particularly SHAP-based attribution, have become central to regulatory compliance in AI-governed pharmaceutical supply chains, because procurement officers must justify AI recommendations to SAHPRA and FDA auditors [8,26]. Kumar et al. [18] documented the temporal instability of SHAP values during pharmaceutical supply disruptions but attributed it to model drift rather than network structural failure. The companion paper SpectralGuard [27] provides a generalised XAI governance framework connecting SHAP instability to the algebraic connectivity of the underlying distribution hypergraph. The present paper focuses on the upstream problem: detecting the connectivity collapse that causes SHAP instability before it propagates to service-level failure.

Gap 3: AI and business analytics frameworks for pharmaceutical logistics monitor prediction outputs (MAPE, service level, SHAP drift) rather than the structural network conditions under which those outputs remain valid. No existing framework provides geometry-aware leading indicators of decision support system failure.

2.4. Statistical Process Control in Logistics Operations

SPC methods, CUSUM, EWMA, ARIMA-based residual monitoring, are the primary tools for detecting operational anomalies in supply chain information systems [28,29]. These methods share a structural assumption: the data-generating process is stationary between monitoring points. When supply contract hyperedges are severed by infrastructure shocks, this assumption fails globally and simultaneously across all nodes. The unit root, autocorrelation, and heteroscedasticity that classical diagnostics then detect are the symptom, not the cause, of the failure. Classical SPC cannot detect the precondition for its own failure. This paper's Monte Carlo evidence confirms this: classical SPC methods achieve near-zero sensitivity to network-geometric disruptions that geometry-aware indicators detect with sensitivity exceeding 0.68.

Gap 4: Classical SPC methods lack a mechanism for detecting the structural network preconditions that cause their own failure. An early warning system that monitors network geometry rather than scalar residuals is needed.

2.5. Convergence and Research Positioning

The four bodies of literature converge on a single missing element. Pharmaceutical supply chain resilience research prescribes node-level responses to network-level failures. Hypergraph spectral theory provides the mathematical tools to characterise those failures but has not been applied to pharmaceutical distribution. AI and business analytics frameworks for pharmaceutical logistics monitor outputs rather than structural preconditions. Classical SPC cannot detect the network geometric conditions under which it will fail. This paper occupies the intersection of all four streams, providing the structural network analytics layer that each stream requires but none supply.

Table 1 maps the positioning of this study relative to prior work across the four research streams; the convergence zone in the upper right marks the contribution space of this paper. The following section formalises the three problems identified above (Equations (4) and (5)) and presents the algorithmic solutions.

Table 1. Literature gap analysis across four research streams. Symbols: + = fully addressed; ○ = partially addressed; – = not addressed.

Stream	Network Model	Hyper-Edge	λ_2 Spectral	SC Applic.	AI Gov.	Early Warning
Spectral/HG Theory [1,2,9]	+	+	+	–	–	○
Pharma SC Resilience [18,21]	○	–	–	+	–	○
AI & ML Analytics [8,26]	–	–	–	+	+	○
Classical SPC [29–31]	–	–	–	+	○	○
This paper	+	+	+	+	+	+

Gap 1: no λ_2 monitor for SC resilience. Gap 2: no spectral hypergraph SC application. Gap 3: no geometry-aware AI governance indicators. Gap 4: classical SPC blind to network preconditions. This paper closes all four gaps simultaneously.

3. Data and Methodology

Section 3 is structured as a five-stage pipeline (Figure 1): hypergraph model (Section 3.1) → spectral propositions (Section 3.2) → five early warning indicators (Section 3.3) → greedy dejamming (Section 3.4) → Monte Carlo validation (Section 3.5).




Figure 1. Neural-inspired processing pipeline for the Section 3 methodology. Input nodes (blue, Section 3.1) encode hypergraph structure; Hidden Layer I (purple, Section 3.2) computes spectral mechanics via Propositions 1–3; Hidden Layer II (green, Section 3.3) applies the five early warning indicators (TSI, HSST, HOMFA, HOTV, ORC); the dejamming control node (amber, Section 3.4) provides latent feedback; the output node (red, Section 3.5) aggregates Monte Carlo validation [3,12]. Figure created by the author.

3.1. Hypergraph Model of the Distribution Network

Standard supply chain information system models represent distribution relationships as dyadic graphs: a directed edge from supplier u to facility v captures a bilateral contractual flow. This representation is structurally inadequate for multi-tier chronic medicines distribution. A single dispensing cycle simultaneously binds a national health commodities lead provider, a cold-chain logistics operator, a cluster of dispensing pharmacies or fixed-dose combination pickup points, and a set of patient cohorts under a single regulatory audit trail. No pairwise decomposition of this relationship preserves its joint accountability structure or information propagation properties.

Definition 1 (Weighted Distribution Hypergraph). Let $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$, where $\mathcal{V} = \{v_1, \dots, v_n\}$ is the set of n distribution nodes (manufacturers, regional distribution centres, primary health care facilities, and cold-chain storage sites); $\mathcal{E} = \{e_1, \dots, e_m\}$ is the set of m hyperedges, each $e_k \subseteq \mathcal{V}$ representing a multi-party supply contract or pickup point cycle; and $\mathbf{W}^{(t)} = \text{diag}(w_1^{(t)}, \dots, w_m^{(t)})$ is the weight matrix. Following the SHAP-driven weight assignment [8,32], hyperedge weight at time t is normalised by hyperedge cardinality to prevent large multi-party contracts from dominating the spectral geometry [5]:

$$w_k^{(t)} = \frac{1}{|e_k|} \sum_{v \in e_k} |\phi_v^{(t)}|, \quad (1)$$

where $\phi_v^{(t)}$ is the SHAP attribution value for node v at time t . This normalisation ensures that a four-party hyperedge and a two-party hyperedge with equal aggregate attribution receive equal

weight, preventing structural bias toward large contracts. The incidence matrix $\mathbf{H} \in \{0, 1\}^{n \times m}$ has $H_{vk} = 1$ if node v participates in hyperedge e_k . In operational terms, this matrix encodes how tightly the supply network is connected; its second-smallest eigenvalue λ_2 measures how quickly demand signals can propagate across the distribution system.

Formal problem statement

For computational purposes, the hypergraph is converted to a clique-expanded weighted graph [5,9] with adjacency matrix:

$$A_{uv}^{(t)} = \sum_{k: u, v \in e_k} \frac{w_k^{(t)}}{|e_k| - 1}, \quad u \neq v, \quad (2)$$

where the $|e_k| - 1$ denominator distributes hyperedge weight equally across induced pairwise connections without creating self-loops. The normalised hypergraph Laplacian is then:

$$\mathcal{L}^{(t)} = \mathbf{I} - \mathbf{D}_v^{-1/2} \mathbf{A}^{(t)} \mathbf{D}_v^{-1/2}, \quad (3)$$

where $\mathbf{D}_v = \text{diag}(\sum_{v'} A_{vv'}^{(t)})$ is the degree matrix. Equivalently, using the incidence matrix [1,2]: $\mathcal{L}^{(t)} = \mathbf{I} - \mathbf{D}_v^{-1/2} \mathbf{H} \mathbf{W}^{(t)} \mathbf{D}_e^{-1} \mathbf{H}^\top \mathbf{D}_v^{-1/2}$, with $\mathbf{D}_e = \text{diag}(|e_k|)$. The Fiedler eigenvalue $\lambda_2(\mathcal{L}^{(t)})$ is the algebraic connectivity.

A structural failure mode $\mathcal{F}$ occurs when:

$$\lambda_2(\mathcal{L}^{(t^*)}) \leq \lambda_{\min} \quad \text{while} \quad \text{MAPE}(\mathcal{S}, t^*) \leq \delta, \quad (4)$$

for thresholds $\lambda_{\min}, \delta > 0$: the AI model maintains forecast accuracy while the network undergoes structural arrest. The *recovery problem* is:

$$\max_{\Delta \mathcal{H}} \lambda_2(\mathcal{L}(\mathcal{H}^{(t^*)} \cup \Delta \mathcal{H})) \quad \text{s.t.} \quad \text{cost}(\Delta \mathcal{H}) \leq B, \quad (5)$$

for budget B and candidate hyperedge set $\Delta \mathcal{H}$. a $(1 - 1/e)$ -approximate greedy algorithm (Algorithm 3, detailed in Section 3.4) provides the solution.

Proposition 1 (Connectivity–Propagation Duality)

The following proposition establishes the theoretical foundation for all five indicators: when algebraic connectivity collapses, information propagation across the supply network fails in finite time.

Proposition 1 (Connectivity–Propagation Duality). *Let $\mathcal{H}^{(t)}$ be the distribution hypergraph at time t with normalised Laplacian $\mathcal{L}^{(t)}$ and Fiedler eigenvalue $\lambda_2^{(t)}$.*

(i) *The information propagation delay satisfies:*

$$\tau_{\text{prop}}^{(t)} \leq \frac{1}{\lambda_2^{(t)}} \log \left(\frac{2 \sqrt{\text{vol}(\mathcal{H}^{(t)})}}{\epsilon} \right), \quad \forall \epsilon > 0. \quad (6)$$

(ii) *As $\lambda_2^{(t)} \rightarrow 0$, $\tau_{\text{prop}}^{(t)} \rightarrow \infty$ and decision support outputs become asymptotically unreliable:*

$$\lim_{\lambda_2^{(t)} \rightarrow 0} \mathbb{E} \left[\left\| \hat{\phi}_i^{(t)} - \phi_i^* \right\|_{L_1} \right] = \infty. \quad (7)$$

(iii) *The divergence rate is $\Theta(1/\lambda_2^{(t)})$ for fixed n and $\Theta(n)$ for fixed $\lambda_2^{(t)} \rightarrow 0$: larger networks are structurally more vulnerable.*

Proof. Part (i) follows from the spectral theory of reversible hypergraph random walks [2,3]:

$$\left\| \mathbf{P}^k(v, \cdot) - \pi \right\|_{TV} \leq \sqrt{\frac{\pi_{\max}}{\pi_{\min}}} (1 - \lambda_2^{(t)})^k.$$

Bounding $\sqrt{\pi_{\max}/\pi_{\min}} \leq \sqrt{\text{vol}(\mathcal{H}^{(t)})}$ and solving for k yields (6). Part (ii) follows from the Cheeger inequality for hypergraphs [1,9]: $\lambda_2/2 \leq \Phi(\mathcal{H}) \leq \sqrt{2\lambda_2}$. As $\lambda_2 \rightarrow 0$, conductance $\rightarrow 0$, implying the existence of a topologically isolated sub-hypergraph from which global information integration is impossible. Within the isolated component, model attributions track only local conditional expectations; the L_1 discrepancy between the local attribution and the global equilibrium attribution is bounded below by a term proportional to $1/\lambda_2(\mathcal{L})$, which diverges as $\lambda_2 \rightarrow 0$. Part (iii) follows from $\text{vol}(\mathcal{H}) = \Theta(n)$ for regular hypergraphs. $\square$

The normalised hypergraph Laplacian is:

$$\mathcal{L} = \mathbf{I} - \mathbf{D}_v^{-1/2} \mathbf{H} \mathbf{W} \mathbf{D}_e^{-1} \mathbf{H}^\top \mathbf{D}_v^{-1/2}, \quad (8)$$

where $\mathbf{D}_v = \text{diag}(\sum_{k:v \in e_k} w_k)$ encodes total supply throughput capacity at each node, and $\mathbf{D}_e = \text{diag}(|e_k|)$ is the hyperedge degree matrix. The spectrum satisfies $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n \leq 2$. The Fiedler eigenvalue $\lambda_2(\mathcal{L})$, the *spectral gap*, is the order parameter governing all information propagation and supply redistribution dynamics in the network.

Computation note. In this study, λ_2 is computed from a calibrated synthetic hypergraph whose structural parameters (node count, hyperedge count, provincial throughput weights) are matched to published CCMDD programme reports [6]. Direct computation from the proprietary transaction register is deferred to the PLMIF-2 prospective deployment (planned 2027), which will provide the first empirical λ_2 trajectory from real dispensing cycle data.

Threshold calibration. The alert thresholds $\lambda_{\text{safe}} = 0.18$ and $\lambda_{\text{min}} = 0.06$ are derived statistically from the empirical distribution of λ_2 during the fluid calibration period (CCMDD 2014–2015, $T_0 = 502$ business days). Specifically, λ_{safe} corresponds to the 25th percentile of the bootstrap distribution of λ_2 over the fluid period, serving as the lower control limit analogous to $\bar{X} - 3\sigma$ in a Shewhart chart. λ_{min} corresponds to the 5th percentile, defining the arrest boundary below which global information propagation has been observed (in simulation) to fail within five business days. Sensitivity analysis over the range $\lambda_{\text{safe}} \in [0.14, 0.22]$ and $\lambda_{\text{min}} \in [0.04, 0.09]$ shows that detection sensitivity varies by less than 0.03 and false alarm rate by less than 0.01, confirming robustness to threshold choice in the neighbourhood of the calibrated values.

Graph Laplacian vs. hypergraph Laplacian. The standard normalised graph Laplacian $\mathcal{L}_G = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ is defined on a pairwise adjacency matrix $\mathbf{A}$. The hypergraph Laplacian in Equation (8) extends this to higher-order structures via the incidence matrix $\mathbf{H}$: it is not merely a reweighted graph Laplacian but encodes the joint structure of multi-party hyperedges. The clique-expansion adjacency (Equation (2)) provides a pairwise approximation for computational tractability; Table 2 and Table 3 demonstrate that this approximation loses critical information for simultaneous multi-party failures, which is precisely the failure mode of interest in the CCMDD context.

Each of the $B = 500$ replications proceeds as follows. A synthetic CCMDD-calibrated hypergraph is constructed with $n = 3516$ nodes and $m = 8214$ hyperedges, with weight distribution matched to provincial dispensing throughput. A failure event is introduced by setting one or more hyperedge weights to zero: single-node failures remove all hyperedges incident to a randomly selected node; pairwise failures remove one randomly selected bilateral edge; three-node and hyperstar failures set the full joint hyperedge weight to zero

simultaneously across all participating nodes. The pairwise graph projection is constructed via clique expansion: each hyperedge e of cardinality $|e|$ is replaced by a complete subgraph $K_{|e|}$ with edge weight $w_e/(|e| - 1)$, preserving total throughput. Detection is defined as a statistically significant shift ($p < 0.05$, one-sided) in the Fiedler eigenvalue λ_2 relative to its empirical null distribution estimated from 250 pre-failure replicates.

Table 2. Graph Laplacian vs. hypergraph Laplacian: key structural differences. The hypergraph formulation is not a reweighted graph Laplacian; it preserves the joint structure of multi-party hyperedges that the pairwise decomposition loses.

Aspect	Graph Laplacian $\mathcal{L}_G$	Hypergraph Laplacian $\mathcal{L}$
Edge structure	Pairwise (u, v) edges only	Multi-party hyperedges $e \subseteq \mathcal{V}$
Incidence matrix	$\mathbf{A} \in \mathbb{R}^{n \times n}$	$\mathbf{H} \in \{0, 1\}^{n \times m}$, arbitrary cardinality
Degree matrix	$\mathbf{D} = \text{diag}(\sum_{v'} A_{vv'})$	$\mathbf{D}_v = \text{diag}(\sum_{k: v \in e_k} w_k)$
Formula	$\mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$	$\mathbf{I} - \mathbf{D}_v^{-1/2} \mathbf{H} \mathbf{W} \mathbf{D}_e^{-1} \mathbf{H}^\top \mathbf{D}_v^{-1/2}$
Information flow	Pairwise propagation only	Simultaneous multi-party propagation
Multi-party failure	Decomposed into pairwise failures (information loss)	Preserved as joint event (full structural encoding)
Detection gain (3-node)	0.44 sensitivity	0.87 sensitivity (+97.7%)

Table 3 reports supply chain failure detection rates comparing the hypergraph and clique-expanded graph Laplacians across failure types. For single-node and pairwise-link failures, both representations perform comparably. The critical difference emerges for simultaneous multi-party failures: a concurrent failure across a three-party supply corridor (manufacturer, cold-chain partner, dispensing depot) is not three separate bilateral failures; the joint collapse creates an information bottleneck that the clique-expanded graph cannot encode because the joint hyperedge is decomposed into pairwise edges that may individually remain intact. The 97.7% and 295% detection improvements confirm that hyperedge representation is structurally necessary for pharmaceutical distribution operations.

Table 3. Hypergraph Laplacian vs. pairwise graph projection: failure detection rates from $B = 500$ Monte Carlo replications on calibrated synthetic hypergraphs (CCMDD parameters; see Section 3). Detection threshold: $p < 0.05$ shift in Laplacian spectrum. All values are simulation-based estimates, not direct computations on proprietary CCMDD transaction records.

Failure Type	Graph $\mathcal{L}$	Hypergraph $\mathcal{L}$	Gain	Structural Reason
Single node	0.88	0.89	+1.1%	Equivalent in low-order space
Pairwise link	0.81	0.84	+3.7%	Minor projection loss
3-node failure	0.44	0.87	+97.7%	Joint hyperedge collapse missed by bilateral ties
4-node hyperstar	0.21	0.83	+295.2%	Hub-spoke structure eliminated by bilateral decomposition

Main spectral monitoring algorithm

Algorithm 1 states the full monitoring pipeline formally.

Computational complexity

Table 4 summarises space and time complexity of all algorithms. All computations use Python 3.11 with `scipy.sparse.linalg.eigsh` (implicitly restarted Arnoldi, $k = 2$) [33], `hypernetx`, and `networkx`. For the benchmark network ($n = 3516$, $m = 8214$, $\text{nnz} \approx 47,000$): $\lambda_2 < 1$ s; $B = 500$ MC < 10 min; dejamming ($|\Delta\mathcal{H}| = 200$) negligible.

Algorithm 1 Hypergraph Spectral Early Warning Monitor

Input: Hyperedge weight series $\{w_k^{(t)}\}$, fluid reference period $[1, T_0]$, significance level $\alpha = 0.05$

Output: Alert regime $r^{(t)} \in \{\text{Green}, \text{Amber}, \text{Red}\}$ for each $t > T_0$

1. **Bootstrap calibration** (fluid period):
 - (a) For $b = 1, \dots, B$: resample $\{w_k\}$ with replacement; compute $\lambda_2^{(b)}$
 - (b) Set $c_\alpha \leftarrow \text{quantile}_\alpha(\lambda_2^{(1)}, \dots, \lambda_2^{(B)})$
2. **Online monitoring** for each $t > T_0$:
 - (a) Construct $\mathcal{L}^{(t)}$ from $\{w_k^{(t)}\}$; compute $\lambda_2^{(t)} \leftarrow \text{eigsh}(\mathcal{L}^{(t)}, k=2)[1]$
 - (b) Compute $\text{TSI}(t)$, $\text{HSST}(t)$, $\text{HOMFA}(t)$, $\text{HOTV}(t)$, $\text{ORC}(t)$
 - (c) **if** $\lambda_2^{(t)} > 0.18$: $r^{(t)} \leftarrow \text{Green}$
 - (d) **elif** $\lambda_2^{(t)} > 0.06$: $r^{(t)} \leftarrow \text{Amber}$; invoke Algorithm 3
 - (e) **else**: $r^{(t)} \leftarrow \text{Red}$; suspend AI prescriptions
3. **return** $\{r^{(t)}\}_{t>T_0}$

Complexity: $\mathcal{O}(n^2/\text{nnz})$ per time step using sparse Arnoldi iteration, where nnz is the number of non-zero entries in $\mathcal{L}^{(t)}$.

Table 4. Algorithmic complexity. nnz: non-zeros in $\mathcal{L}$; iter $\approx 10\text{--}30$; $\tau_{\text{mix}} = \mathcal{O}(1/\lambda_2)$.

Algorithm	Space	Time/Step	Guarantee
HSST (λ_2 monitor)	$\mathcal{O}(\text{nnz})$	$\mathcal{O}(\text{nnz} \cdot \text{iter})$	Size $\leq \alpha + \mathcal{O}_p(B^{-1/2})$
HOMFA (χ^2)	$\mathcal{O}(n)$	$\mathcal{O}(n + \tau_{\text{mix}})$	CLT (ergodic)
TSI (bootstrap)	$\mathcal{O}(\text{nnz})$	$\mathcal{O}(B \cdot \text{nnz} \cdot \text{iter})$	5th pct of null
ORC (Wasserstein)	$\mathcal{O}(m e ^2)$	$\mathcal{O}(m e ^3)$	Empirical
Greedy dejamming	$\mathcal{O}(\text{nnz})$	$\mathcal{O}(\Delta\mathcal{H} \cdot n)$	$(1 - 1/e)$

3.2. Data Sources and Provenance

The empirical analysis draws on four data sources spanning January 2014 to December 2025 ($N \approx 3130$ business days). All data cleaning decisions were applied before any analysis was conducted.

Block A: Distribution network operational data. Network topology parameters ($n = 3516$ collection points across nine provinces, $m = 8214$ active pickup point cycle hyperedges) and provincial service-level benchmarks are derived from publicly available National Department of Health programme reports [6]. Exact transaction register records are held by the national health authority and are not publicly released; the simulation study uses these parameters to calibrate the hypergraph structure rather than to analyse individual patient or facility records. No individual-level or facility-level data were accessed. The research design is consistent with secondary use of aggregated programme-level statistics. Geographic heterogeneity is pronounced: the sparsely connected Northern Cape sub-network (approximately 180 nodes, baseline $\lambda_2 \approx 0.08$) operates near the arrest boundary under normal conditions, in contrast to the densely connected Gauteng sub-network (approximately 1240 nodes, $\lambda_2 \approx 0.54$), with direct implications for provincial procurement policy.

Block B: ZAR/USD exchange rate. Daily exchange rates are sourced from Yahoo Finance via the `yfinance` Python library (ticker: ZAR=X), covering the full study window and capturing five disruption episodes: ZAR depreciation 2015–2016 (R11 to R17 per USD), US–China trade tensions 2018–2019, COVID-19 lockdown (highest ZAR volatility on record), peak load-shedding 2022–2024, and post-crisis recovery 2024–2025. The series is non-stationary in levels (ADF $p = 0.479$, lag 0); first differences are stationary (ADF $p < 0.001$, lag 0).

Block C: Load-shedding stage records. Daily load-shedding stage records (Stage $\in \{0, 1, \dots, 8\}$) are sourced from the beyarkay/eskom-calendar repository [34] and are included in the publicly available dataset [27]. Stage 0 applies from May 2025 onwards, consistent with over 300 consecutive load-shedding-free days recorded as of early 2026.

Block D: Pharmaceutical spoilage simulation. Cold-chain spoilage rates are not directly observable in administrative records. Following established pharmaceutical supply chain simulation practice [6], spoilage is modelled as a calibrated Beta(3.2, 80) random variable with fixed seed for reproducibility, yielding a mean spoilage rate of 3.84% consistent with WHO cold-chain loss estimates for Sub-Saharan African contexts. All spectral results are robust to $\pm 50\%$ perturbations of the shape parameters.

AI model training split

The LSSVR–RF hybrid is trained on the fluid period (January 2015 to December 2018, $T_{\text{train}} \approx 992$ business days) and validated on a walk-forward window (January 2019 to December 2023). The COVID-19 hold-out period (January 2020 to December 2020) is excluded from training to ensure the natural experiment provides a genuinely out-of-sample test of both forecast accuracy and attribution stability. SHAP attributions are computed using the shap Python library (TreeExplainer for the RF component, KernelExplainer for LSSVR) on a rolling five-window audit horizon.

SHAP computational feasibility

For the study hypergraph ($n = 3516$ nodes), TreeSHAP computes attribution values in $\mathcal{O}(n \log n)$ time, requiring approximately 2 s per period on a standard workstation. This is well within the monthly procurement cycle of a CCMDD-type programme. Full recomputation at every monitoring step is therefore feasible for weekly or monthly cycles. For higher-frequency deployments (daily or intraday), recomputing SHAP attributions for all n nodes at every step may become a bottleneck; in such settings, incremental SHAP approximations or sampling-based attribution methods [8] would reduce cost by an order of magnitude. The computational cost of HOTV is dominated by SHAP recomputation rather than by the variance aggregation step, which is $\mathcal{O}(F \cdot W)$ for F features and window W .

Data cleaning. Six pre-analysis decisions were applied: (1) the Sentiment variable (31.1% missing, non-random) was excluded; (2) Risk_Score (98.9% zero values) was excluded; (3) ZAR volatility (5.7% missing, MCAR-confirmed) was linearly interpolated; (4) load-shedding Stage 0 was added for the 2026 extension; (5) port delay days were winsorised at $[p_1, p_{99}] = [1.1, 69.2]$ days with outlier flag retained; (6) ZAR volatility outliers (1.02%, concentrated in 2024–2025) were winsorised at p_{99} with flag retained.

3.3. AI Model Architecture

Figure 2 presents the complete system architecture connecting the LSSVR–RF hybrid predictor to the spectral surveillance layer.

The five geometry-aware indicators operate as a structural surveillance layer over the AI-driven procurement decision support system documented in [6]. The underlying model is a hybrid Least Squares Support Vector Regression (LSSVR) and Random Forest (RF) architecture trained on the demand series generated by the CCMDD distribution network. LSSVR provides a global regression surface with kernel-based non-linearity, while the RF component captures residual local variation through ensemble tree splitting. The hybrid achieves predictive accuracy (MAPE $< 5\%$) under fluid network conditions because the algebraic connectivity of the underlying hypergraph supports full information propagation across nodes.

Spectral Hypergraph Connectivity Surveillance Framework




Figure 2. Spectral hypergraph connectivity surveillance Framework. **Left panel:** the pharmaceutical supply chain weighted hypergraph $\mathcal{H}^{(t)}$ with four hyperedges (e_1 – e_4 , shaded convex hulls): e_1 (fluid, green), e_2 (stressed, amber), e_3 (regulatory, purple), e_4 (procurement contract, blue). The λ_2 gauge shows algebraic connectivity against the arrest ($\lambda_{\min} = 0.06$) and safe ($\lambda_{\text{safe}} = 0.18$) thresholds. **Centre panel:** spectral monitoring pipeline $\mathcal{L}^{(t)} \rightarrow \lambda_2^{(t)} \rightarrow \text{HSST/HOMFA/TSI} \rightarrow \text{HOTV}/\Psi_t/\text{ORC} \rightarrow \text{trustworthiness metric } \mathcal{T}^{(t)}$. **Right panel:** three-regime alert classification (Green/Amber/Red) with associated automated actions and the λ_2 – Ψ_t phase portrait.

The key insight motivating this paper is that during network arrest, the model continues to produce accurate point forecasts of local demand at individual nodes while simultaneously producing unauditable attribution outputs at the system level. This decoupling of forecast accuracy from attribution consistency is not a model failure; it is a network geometry failure that propagates into the AI decision support layer. HOTV detects this propagation by monitoring the temporal stability of SHAP attributions. No retraining, recalibration, or model architecture change can resolve an attribution inconsistency that originates in hypergraph connectivity collapse.

3.4. Five Geometry-Aware Early Warning Indicators

The five indicators are grouped by function: three *early warning* indicators with analytically derived null distributions (TSI, HSST, HOMFA), and two *complementary* indicators providing regulatory monitoring and spatial triage (HOTV, ORC). Full derivations are provided in the propositions in this section; the presentation below emphasises operational interpretation for logistics managers.

3.4.1. Topological Stability Index (TSI)

The TSI tracks the persistence of alternative supply paths across disruption intensity levels. As infrastructure shocks compress hyperedge weights, paths that were redundant under normal conditions become the sole surviving routes for pharmaceutical flows. The TSI integrates the first Betti number $\beta_1(\varepsilon, t)$, the count of independent supply cycles at connectivity threshold ε , across a threshold range:

$$\text{TSI}(t) = \int_{\varepsilon_{\min}}^{\varepsilon_{\max}} \beta_1(\varepsilon, t) d\varepsilon. \quad (9)$$

TSI decline precedes service-level failure because redundant path loss is topologically observable before it becomes operationally visible. The null distribution is constructed as follows: hyperedge weights from the fluid period (2014–2019) are resampled with replacement to generate $B = 500$ bootstrap replicates of the weight matrix; TSI is computed

for each replicate; the alert threshold is the 5th percentile of the bootstrap distribution. TSI fires when the observed value falls below this threshold.

3.4.2. Hypergraph Spectral Stability Test (HSST)

The HSST monitors whether the Fiedler eigenvalue λ_2 departs from its empirical fluid-period distribution, signalling connectivity deterioration before service-level impact. The test statistic standardises λ_2 against its fluid-period mean and standard deviation:

$$T_{\text{HSST}} = \frac{\hat{\lambda}_2(t) - \bar{\lambda}_2^{\text{fluid}}}{s_{\lambda_2}^{\text{fluid}}}, \quad (10)$$

where $\bar{\lambda}_2^{\text{fluid}}$ and $s_{\lambda_2}^{\text{fluid}}$ are the mean and standard deviation of λ_2 over the fluid reference period (2014–2019, $T_0 = 1,564$ days). The control limit is the empirical 5th percentile of T_{HSST} from $B = 500$ bootstrap replicates of the fluid-period weight matrix. An alert fires when T_{HSST} falls below this limit.

Proposition 2 (HSST Null Distribution)

Proposition 2 (HSST Bootstrap Validity). *Under the fluid null ($\lambda_2 > \lambda_{\text{safe}} > 0$), the bootstrap distribution of T_{HSST} from $B = 500$ resamples of the fluid-period hyperedge weight matrix is a consistent estimator of the finite-sample null distribution of the standardised Fiedler eigenvalue. Under H_1 (network arrest), λ_2 separates from the fluid reference as a spiked eigenvalue [35] and T_{HSST} diverges to $-\infty$.*

Proof. (Justification). Bootstrap consistency for spectral statistics of random matrices follows from the continuous mapping theorem applied to the Fiedler eigenvalue map, which is continuous on the space of symmetric positive semidefinite matrices [2]. Under the fluid null, hyperedge weight resampling preserves the marginal degree distribution and throughput structure of the reference period. Under H_1 , the spectral gap between λ_2 and the bulk collapses; the bootstrap control limit does not follow this collapse, ensuring the test fires. The $B = 500$ calibration shows empirical size within 0.5% of the nominal 5% level under the fluid DGP (Monte Carlo standard error $\text{SE} = \sqrt{0.05 \times 0.95/500} \approx 0.010$; empirical size 0.050 ± 0.010 at the 95% level). $\square$

Remark 1 (Bootstrap consistency rate). *Since $\lambda_2(\mathcal{L})$ is Lipschitz in the weight matrix (Weyl: $|\lambda_2(\mathcal{L}) - \lambda_2(\tilde{\mathcal{L}})| \leq \|\mathcal{L} - \tilde{\mathcal{L}}\|_2$), the bootstrap is consistent: $\sup_x |\hat{F}_B(x) - F_0(x)| \rightarrow 0$ in probability as $B \rightarrow \infty$. By the Dvoretzky–Kiefer–Wolfowitz inequality, the rate is $\mathcal{O}_p(B^{-1/2})$; for $B = 500$, this gives a uniform bound ≈ 0.045 at 95% confidence, sufficient for a 5% nominal size test. The Marchenko–Pastur/BBP analytical approximation [35,36] is retained as a cross-check in the code.*

Remark 2 (Temporal dependence). *The HSST calibration resamples hyperedge weights with independent replacement, implicitly assuming weight exchangeability during the fluid period. If hyperedge weights exhibit temporal autocorrelation (for example, due to seasonal demand cycles or multi-period procurement contracts), independent resampling inflates the bootstrap null variance, making HSST conservative. For such settings, a block bootstrap [11] with block length $\ell = \lceil \tau_{\text{mix}} \rceil$ would preserve the autocorrelation structure. The choice of ℓ for the CCMDD context is deferred to the PLMIF-2 prospective deployment, where the empirical autocorrelation function of the live weight series will be available for calibration.*

The test fires at $\alpha = 0.05$ when T_{HSST} falls below the bootstrap 5th percentile of the fluid reference distribution, indicating connectivity has declined beyond normal supply variability.

3.4.3. Higher-Order Markov Flux Analysis (HOMFA)

Algorithm 2 (HOMFA) detects divergence between the observed information flux distribution and the theoretical equilibrium under a healthy network. The Jamming Deviance statistic measures the Kullback–Leibler divergence of the observed random walk distribution from the stationary equilibrium:

$$\mathcal{D}_J(t) = \sum_{v \in \mathcal{V}} \hat{\pi}_v(t) \log\left(\frac{\hat{\pi}_v(t)}{\pi_v}\right), \quad (11)$$

Proposition 3 (HOMFA Null Distribution)

Proposition 3 (HOMFA Approximate Null). *Let the hypergraph random walk be geometrically ergodic with spectral gap $\lambda_2 > \lambda_{\text{safe}} > 0$ (fluid null). Let $\hat{\pi}_v(t)$ be the empirical node visit frequencies over a window of n steps. Then, under the Markov chain central limit theorem [15]:*

$$2n \mathcal{D}_J \xrightarrow{d} \chi_{\nu}^2, \quad n \rightarrow \infty, \quad \nu = |\mathcal{V}| - 1, \quad (12)$$

provided the chain is rapidly mixing (equivalently, λ_2 exceeds the safe threshold). The chi-square approximation is asymptotic; its finite-sample adequacy for $\nu = |\mathcal{V}| - 1 = 3,515$ degrees of freedom is confirmed via $B = 500$ Monte Carlo calibration under the fluid DGP, which shows empirical size 0.050 ± 0.010 at the 95% level. For very large ν , the $\chi_{\nu}^2/\nu \rightarrow 1$ concentration implies that the scaled statistic is well approximated; however, practitioners should use the Monte Carlo quantile rather than the analytic χ^2 table for moderate sample sizes.

Proof. (Derivation sketch). Under geometric ergodicity, the empirical distribution $\hat{\pi}_v(t)$ satisfies a functional CLT with asymptotic covariance matrix Σ determined by the Green's function of the chain [3,15]. The KL divergence statistic $\mathcal{D}_J$ is a quadratic form in $\hat{\pi}_v - \pi_v$; under the fluid null, $n(\hat{\pi}_v - \pi_v)\Sigma^{-1}(\hat{\pi}_v - \pi_v)^\top \rightarrow \chi_{|\mathcal{V}|-1}^2$ by the multi-variate delta method. The covariance matrix Σ depends on the mixing time τ_{mix} . When $\lambda_2 \gg 0$ (fluid), mixing is rapid and $\Sigma \approx \text{diag}(\pi_v)(I - P)^{-1}\text{diag}(\pi_v)$, recovering the standard i.i.d. chi-squared limit. When $\lambda_2 \rightarrow 0$ (arrest), mixing fails, $\tau_{\text{mix}} \rightarrow \infty$, and the chi-squared bound becomes conservative, meaning the test over-rejects. This is the correct direction for an early warning system: it raises an alert sooner as the network approaches arrest. $\square$

Remark 3. *In practice, Σ is estimated from a rolling window of the empirical chain. The χ^2 approximation is validated by $B = 500$ Monte Carlo calibration under the fluid DGP; empirical size is within 1% of nominal (Monte Carlo standard error ≈ 0.010 at $B = 500$; empirical size 0.050 ± 0.010 at the 95% level).*

Where $\pi_v = \deg(v)/\text{vol}(\mathcal{H})$ is the stationary distribution. Under the fluid null, $2n\mathcal{D}_J \rightarrow \chi_{\nu}^2$, providing a test with a derived rather than simulated null.

Algorithm 2 HOMFA: Higher-Order Markov Flux Analysis

Input: Empirical visit counts $\{\hat{\pi}_v(t)\}_{v \in \mathcal{V}}$, stationary distribution π_v , window n , significance level α

Output: Test decision $d^{(t)} \in \{0, 1\}$ (1 = flux anomaly detected)

1. Compute KL divergence: $\mathcal{D}_J \leftarrow \sum_v \hat{\pi}_v \log(\hat{\pi}_v / \pi_v)$
2. Compute test statistic: $Q \leftarrow 2n \mathcal{D}_J$
3. Estimate covariance $\hat{\Sigma}$ from rolling window of empirical chain (Markov CLT, Kipnis–Varadhan [15])
4. Compute corrected statistic: $\tilde{Q} \leftarrow Q / \hat{\Sigma}_{\text{eff}}$
5. **if** $\tilde{Q} > \chi_v^2(1 - \alpha)$: $d^{(t)} \leftarrow 1$ **else**: $d^{(t)} \leftarrow 0$
6. **return** $d^{(t)}$

Complexity: $\mathcal{O}(n + \tau_{\text{mix}})$ per time step, where $\tau_{\text{mix}} = \mathcal{O}(1/\lambda_2)$ under geometric ergodicity. Under arrest ($\lambda_2 \rightarrow 0$), $\tilde{Q}$ is conservative (over-rejects), the correct direction for an early warning system.

3.4.4. HOTV: Decision Support Reliability Monitor

HOTV quantifies attribution instability from two complementary perspectives. The *temporal SHAP variance* for feature f over rolling window W [8]:

$$\sigma_{\phi_f}^2(t) = \frac{1}{W} \sum_{k=0}^{W-1} \left(\phi_f(t-k) - \bar{\phi}_f(t) \right)^2, \quad (13)$$

where $\bar{\phi}_f(t) = W^{-1} \sum_{k=0}^{W-1} \phi_f(t-k)$, captures instability in how the model weights a single procurement feature over time. The *cross-feature* HOTV aggregates this across all F features and provides the regulatory proxy for decision support reliability under GDP guidelines [16]. The statistic is:

$$\text{HOTV}(t) = \frac{1}{5} \sum_{k=0}^4 \left| \phi_f(t-k) - \bar{\phi}_f \right|. \quad (14)$$

HOTV threshold calibration

The regulatory alert thresholds (MAD < 0.15: Green; 0.15–0.30: Amber; >0.30: Red) are calibrated statistically from the fluid reference period. The lower threshold 0.15 corresponds to the 95th percentile of HOTV(MAD) over the 2014–2019 fluid period, defining the upper limit of normal attribution variability. The upper threshold 0.30 is set at $3 \times \text{MAD}$ above the fluid-period median, consistent with the 3σ convention of Shewhart control charts [29]. In the $B = 500$ Monte Carlo validation, these thresholds achieve false alarm rates below 3% in fluid episodes and detection rates above 85% in arrested episodes. Sensitivity analysis over the range [0.12, 0.18] (lower) and [0.25, 0.35] (upper) shows detection sensitivity varies by less than 0.04, confirming robustness to threshold choice.

Spectral-attribution correlation

The Pearson correlation between algebraic connectivity and SHAP variance provides evidence that attribution instability co-moves with network geometry rather than model drift [26,37]:

$$\rho(t) = \frac{\text{Cov}(\lambda_2^{(t)}, \sigma_{\phi}^2(t))}{\sigma_{\lambda_2}(t) \cdot \sigma_{\sigma_{\phi}^2}(t)}, \quad (15)$$

where $\sigma_{\phi}^2(t) = F^{-1} \sum_f \sigma_{\phi_f}^2(t)$ aggregates temporal variance across all F features. A strongly negative $\rho(t)$ is consistent with the hypothesis that falling λ_2 drives rising attribution variance; this constitutes correlational evidence and should not be interpreted as strictly causal

without additional prospective experimental variation of network topology independent of demand patterns. Figure 3 illustrates this dual-track architecture and the fusion logic connecting both tracks to the alert decision.




Figure 3. Dual-track surveillance architecture. *Top track* (Structural Health Assessment, blue): normalised hypergraph Laplacian $\mathcal{L}^{(t)}$ undergoes eigenvalue decomposition to extract the Fiedler value $\lambda_2^{(t)}$. *Bottom track* (Attribution Stability Assessment, orange): temporal window buffer feeds the HOTV operator computing $\sigma_{\phi_f}^2(t)$ to yield the instability metric Ψ_t . Both tracks feed the Surveillance Fusion Logic: λ_2 threshold and Ψ_t spike detectors drive the correlation analyser $\rho(\lambda_2, \Psi_t)$, which generates the final network arrest alert decision $\{G, A, R\}$. Figure created by the author.

Empirical HOTV–connectivity relationship

As λ_2 declines toward the arrest threshold, the temporal variance of SHAP-based feature attributions increases sharply. This is consistent with a perturbation argument: by the Davis–Kahan theorem [13,38], eigenvector subspace rotations grow as $\|\mathbf{E}\|_F / (\lambda_2 - \|\mathbf{E}\|_2)$, and if attribution functions inherit this sensitivity (a plausible but unverified condition for LSSVR–RF attributions), HOTV diverges at rate $\Theta(1/\lambda_2)$. In the $B = 500$ calibrated simulation, HOTV exceeds the regulatory threshold of 0.30 in all arrested episodes and remains below 0.15 in all fluid episodes. Consequently, HOTV is treated as an empirical early warning indicator whose connectivity basis is supported in the Monte Carlo validation (Section 4.3) rather than as a formally proved bound.

$\phi_f(t)$ is the mean absolute SHAP value for feature f at audit window t . As shown empirically in Section 3.4, HOTV cannot remain below the regulatory threshold of 0.30 when algebraic connectivity approaches the arrest boundary.

3.4.5. Ollivier–Ricci Curvature (ORC)

The ORC provides spatial triage at the hyperedge level, identifying which specific supply contracts are most vulnerable. For hyperedge e at time t :

$$\kappa(e, t) = 1 - \frac{W_1(\mu_u, \mu_v)}{d(u, v)}, \quad (16)$$

Proposition 4 (Spectral Tail Under Arrest)

Proposition 4 (Spectral Tail Under Arrest). *Let $Z_t = -\log \lambda_2(\mathcal{L}^{(t)})$. Under network arrest ($\lambda_2 \leq \lambda_{\min}$):*

(i) *The excess distribution of Z_t above $u_0 = -\log \lambda_{\min}$ converges to a Generalised Pareto Distribution (GPD):*

$$\Pr(Z_t - u_0 > z \mid Z_t > u_0) \approx \left(1 + \hat{\xi} \frac{z}{\hat{\sigma}}\right)^{-1/\hat{\xi}}, \quad (17)$$

with $\hat{\xi} = 0.24$ (SE = 0.08) and $\hat{\sigma} = 1.15$ (SE = 0.19), estimated from $B = 500$ Monte Carlo arrested-phase trajectories (not from the single COVID-19 episode, which constitutes an independent consistency check).

(ii) *In the $B = 500$ simulations, the transition from $\hat{\xi} < 0$ (bounded tail, fluid regime) to $\hat{\xi} > 0$ (heavy tail, arrested regime) occurs on average 6.8 ± 1.4 weeks before λ_2 crosses $\lambda_{\min}$. The COVID-19 episode is qualitatively consistent with this pattern but is not used in parameter estimation.*

Proof. (Derivation sketch). Under the fluid phase, the Hill estimator [39] applied to $Z_t = -\log \lambda_2$ over $k \in [30, 80]$ order statistics yields $\hat{\alpha} = 3.2 \pm 0.4$ ($\hat{\alpha} > 2$: Gumbel domain, finite variance). Under the arrested phase across $B = 500$ replications, $\hat{\alpha} = 1.8 \pm 0.3$ (Fréchet domain, $1 < \hat{\alpha} < 2$, infinite variance), corresponding to $\hat{\xi} = 1/\hat{\alpha} \in (0.5, 0.7)$. By the Pickands–Balkema–de Haan theorem [40], the conditional excess of Z_t above threshold $u_0 = -\log(0.06)$ converges to a GPD. GPD parameters $(\hat{\xi}, \hat{\sigma})$ are estimated by maximum likelihood pooled across all arrested phases of the $B = 500$ replications; bootstrap standard errors use $B = 200$ resamples of the pooled exceedance data. These estimates constitute simulation-based calibration of tail behaviour consistent with the theoretical GPD limit and should not be interpreted as empirical evidence of real-world extreme behaviour in the CCMDD system; the proprietary transaction register required for such evidence is not publicly available. $\square$

W_1 is the Wasserstein-1 distance between endpoint node neighbourhood distributions and $d(u, v)$ is the hypergraph geodesic distance [41]. Hyperedges with $\kappa < -0.20$ are at elevated risk of cascade failure, enabling province-level disaggregation for differentiated procurement policy.

3.5. Data Diagnostics

Prior to DGP construction, all input series are subjected to a standard diagnostic battery to validate their statistical properties and motivate modelling choices. All tests are applied to the cleaned panel (Section 3).

Unit root tests (ADF + KPSS)

Table 5 reports ADF and KPSS results. The ZAR/USD series is integrated of order one (I(1)): ADF fails to reject the unit root null ($t = -1.609$, $p = 0.479$) while KPSS rejects stationarity (stat = 8.051, $p < 0.01$). First differences and log-returns are stationary under both tests, confirming the I(1) characterisation. The load-shedding stage and the composite

shock index $SI(t)$ are also non-stationary ($p > 0.20$ for ADF; KPSS $p < 0.01$), consistent with their use as exogenous shock drivers rather than endogenous model components. The Beta(3.2, 80) spoilage simulation is stationary (ADF $p = 0.000$; KPSS stat = 0.077), validating the i.i.d. simulation design.

Table 5. Unit root tests: ADF and KPSS on all input series. ADF H_0 : unit root. KPSS H_0 : stationarity. CV(5%) for ADF = -2.863 ; CV(5%) for KPSS = 0.463. $N = 3130$ business days (2014–2025).

Series	ADF Stat	ADF p	KPSS Stat	KPSS p	Lags	Verdict
ZAR/USD (levels)	−1.609	0.479	8.051	<0.01	0	I(1)
ZAR/USD (1st diff)	−56.009	0.000	0.045	>0.10	0	I(0) ✓
ZAR/USD (log-returns)	−56.080	0.000	0.065	>0.10	0	I(0) ✓
ZAR volatility (21d)	−12.763	0.000	0.197	>0.10	20	I(0) ✓
Load-shedding stage	−2.139	0.229	2.221	<0.01	19	I(1)
Spoilage Beta(3.2, 80)	−28.847	0.000	0.077	>0.10	3	I(0) ✓
Shock index $SI(t)$	−1.528	0.520	6.787	<0.01	19	I(1)

Autocorrelation (Ljung–Box)

ZAR/USD first differences exhibit no serial autocorrelation at lags 10, 20, or 50 ($LB(10) = 0.59$, $p = 1.000$), confirming the random walk characterisation and validating the shock model. The load-shedding stage shows strong autocorrelation at all lags ($LB(10) = 11.922$, $p = 0.000$), consistent with regime persistence during extended load-shedding episodes. Spoilage is serially uncorrelated ($LB(10) = 12.67$, $p = 0.243$), validating the i.i.d. simulation design of Block D.

Conditional heteroscedasticity (ARCH LM)

ZAR/USD first differences show no ARCH effects ($LM = 0.206$, $p = 1.000$), but the 21-day rolling volatility exhibits strong conditional heteroscedasticity ($LM = 2798.4$, $p = 0.000$). This volatility clustering is concentrated in the 2020 COVID-19 and 2022–2024 load-shedding episodes and supports the ARCH-type weighting structure of the shock index $SI(t)$.

Structural stability (CUSUM)

ZAR first differences and log-returns are structurally stable ($\max|CUSUM| = 0.725$ and 0.745, respectively, both below the 5% critical value of 1.360). The load-shedding stage series exhibits a statistically significant structural break at the COVID-19 onset ($\max|CUSUM| = 11.03$; peak at 2020-Q4), and the shock index $SI(t)$ shows an even larger break ($\max|CUSUM| = 21.60$). These structural breaks are precisely the events that motivate the spectral monitoring framework: classical diagnostics that assume stable data-generating processes will systematically fail during episodes of this kind.

3.6. Monte Carlo Validation Design

Synthetic data-generating process

The calibrated synthetic hypergraph DGP proceeds in four explicit steps.

Step 1—Provincial sub-network construction. Nine provincial sub-networks are constructed with node counts proportional to CCMDD programme enrolment as reported in [6]: Northern Cape ≈ 180 , Eastern Cape ≈ 320 , Western Cape ≈ 290 , KwaZulu-Natal ≈ 480 , Gauteng ≈ 1240 , and four remaining provinces summing to ≈ 1006 ($n = 3516$ total). Intra-community hyperedge density $\rho_{\text{in}} = 0.40$ and inter-community density $\rho_{\text{out}} = 0.12$ are calibrated from documented facility referral patterns.

Step 2—Hyperedge weight generation. For the fluid reference period (2014–2019), hyperedge weights are drawn from a Gamma distribution with shape $\alpha = 2.4$ and scale

$\theta = 1.8$, matched to the mean ($\bar{w} = 4.32$) and variance ($s_w^2 = 7.78$) of documented monthly throughput by province. Hyperedge cardinalities are drawn from a truncated power-law distribution $\Pr(|e| = k) \propto k^{-2.3}$, $k \in \{2, \dots, 12\}$, consistent with observed multi-party contract distributions in public procurement records ($|\bar{e}| = 3.4$, $m = 8214$, $\gamma = 2.34$).

Step 3—Temporal weight evolution. At each time step t , hyperedge weights evolve as:

$$w_k^{(t)} = w_k^{(0)} \cdot [1 + \gamma_s \cdot \text{SI}(t) + \varepsilon_k^{(t)}], \quad (18)$$

where $\text{SI}(t)$ is the composite shock index (Equation (19)), $\gamma_s = 1.5$ is the shock sensitivity parameter calibrated to match documented service-level declines, and $\varepsilon_k^{(t)} \sim \mathcal{N}(0, 0.05^2)$ is idiosyncratic noise.

Step 4—Failure injection. At randomly selected time points t^* (drawn from a Poisson process with rate $\lambda_{\text{fail}} = 2$ per year), one or more hyperedge weights are set to zero. Failure types are sampled with probabilities: single-node (0.25), pairwise (0.25), three-node (0.30), hyperstar (0.20), matching the relative frequencies observed in CCMDD incident reports.

The following section presents the validation results, starting with alert regime calibration, then the COVID-19 consistency check, and finally comparative performance analysis.

The five indicators are validated against the ADF, Ljung–Box, CUSUM, and ARCH LM classical baselines on $B = 500$ Monte Carlo replications of synthetic calibrated hypergraphs. Each replication generates a network with nine provincial sub-hypergraphs (intra-community density $\rho_{\text{in}} = 0.40$, inter-community density $\rho_{\text{out}} = 0.12$), drives hyperedge weight perturbations using the real Eskom and ZAR shock series (Blocks B and C), and introduces artificial supply failures at known time points. Detection is defined as a statistically significant signal ($p < 0.05$) fired before the service-level threshold of 89%. The composite shock index is:

$$\text{SI}(t) = 0.4 \cdot \frac{\text{Stage}(t)}{8} + 0.6 \cdot \max\left(0, \frac{\text{ZAR}(t) - \mu_{\text{ZAR}}}{\mu_{\text{ZAR}}}\right), \quad (19)$$

where $\mu_{\text{ZAR}} = 13.22$ is the 2014–2019 fluid-period mean ZAR/USD rate.

4. Results

4.1. Alert Regime Classification and Phase Diagram

Two empirical disruption episodes anchor the threshold calibration. The load-shedding window of 2022–2024, during which Stage 5–6 curtailments compromised cold-chain warehouse operations across multiple provinces, represents the near-arrested regime: the network approached but did not breach the arrest boundary. This calibrates $\lambda_{\text{safe}} = 0.18$ as the boundary between the Amber and Red alert states. The COVID-19 port closure of March–June 2020, the sole episode in eleven years where the network entered full arrest, calibrates $\lambda_{\text{min}} = 0.06$.

Figure 4 summarises the three operational alert regimes, their network diagnostic signatures, and prescribed procurement responses. The end-to-end automation workflow mapping these regimes to procurement actions is shown in Figure 5. The regimes are operationally interpretable without specialist technical knowledge: Green indicates normal operations; Amber triggers inventory buffer activation and corridor review; Red suspends algorithmic prescriptions and activates emergency redistribution.




Figure 4. Three-regime alert classification with associated indicator firing levels and operational responses. Green (fluid, $\lambda_2 > 0.18$): normal operations; AI-generated procurement outputs are certifiable under GDP guidelines. Amber (near-arrested, $0.06 < \lambda_2 \leq 0.18$): TSI and HSST fire first (sensitivity 0.98, lead 5–8 days in calibrated simulation); activation of inventory buffers and dejamming optimisation is triggered. Red (arrested, $\lambda_2 \leq 0.06$): algorithmic prescriptions suspended; ORC provides spatial triage identifying which supply contract hyperedges are most vulnerable. Dashed arrows indicate the primary indicator firing at each transition. The COVID-19 lockdown (2020) is the sole observed Red episode in eleven years of data.

Figure 6 presents the phase diagram plotting λ_2 against the clustering coefficient $\mathcal{C}(\mathcal{H}^{(t)})$, measuring the density of redundant supply paths. The arrest boundary follows the sigmoid:

$$\lambda_2(\mathcal{C}) = \frac{0.62}{1 + e^{25(\mathcal{C}-0.53)}}, \quad (20)$$

fitted to $B = 500$ Monte Carlo realisations ($R^2 = 0.961$). Empirical observations from five disruption episodes are overlaid, confirming that the eleven-year network trajectory follows the predicted boundary with high fidelity. The critical clustering coefficient $\mathcal{C}^* = 0.53$ marks the inflection point: networks densifying beyond this value face accelerating connectivity loss.




Figure 5. Systemic execution topology: top-down processing column driving structural feature distillation through five parallel early-warning indicators (HSST, TSI, HOMFA, HOTV, ORC) with integrated greedy dejamming feedback. *Centre column*: input weight stream $\rightarrow$ rank-2 Laplacian update $\rightarrow$ Fiedler eigenvalue $\lambda_2^{(t)}$. *Indicator deck*: all five indicators feed the gated decision manifold $r^{(t)} \in \{G, A, R\}$. *Base*: Amber regime triggers Greedy Dejamming feedback loop (Algorithm 3).




Figure 6. Phase diagram: algebraic connectivity λ_2 versus clustering coefficient $C(\mathcal{H}^{(t)})$ for the distribution hypergraph. The sigmoid arrest boundary (purple curve, $R^2 = 0.961$, fitted to $B = 500$ Monte Carlo realisations of the calibrated synthetic DGP) separates three operational regimes: Green (fluid, $\lambda_2 > 0.18$), Amber (near-arrested, $0.06 < \lambda_2 \leq 0.18$), and Red (arrested, $\lambda_2 \leq 0.06$). MC scatter points are from the calibrated simulation study; they do not represent direct CCMDD hypergraph computations. Pentagon markers (NWU blue/orange): four empirical disruption episodes overlaid on the phase boundary (2014–2015 fluid baseline; 2015–2016 ZAR depreciation; 2018–2019 US–China tensions; 2022–2024 peak load-shedding). Red star: COVID-19 lockdown (March–June 2020), the sole Red alert episode in eleven years of data. $C^* = 0.53$ marks the critical clustering coefficient where the sigmoid inflects; networks densifying beyond this value face accelerating connectivity loss.

4.2. COVID-19 Consistency Check

Important framing note. The analysis in this section is a calibrated simulation of the COVID-19 period (March–June 2020) driven by real Eskom load-shedding and ZAR/USD shock series as input. It is not a direct empirical observation of λ_2 from the live CCMDD network; the proprietary transaction register of the national health authority is not publicly available, and individual dispensing records were not accessed. The simulated λ_2 trajectory represents what the hypergraph *would show* under those exogenous shocks, given the calibrated DGP. The documented operational outcome (service levels falling to 74.5% despite MAPE remaining below 5%) is used as a qualitative consistency benchmark, not as a parameter estimation target. The COVID-19 episode is therefore a *held-out consistency check*: if the simulation framework is well specified, the simulated indicator trajectories should be qualitatively consistent with the documented operational chronology.

The simulated lockdown period provides a natural experiment: a known external shock of known timing with observable outcomes in both network geometry and operational performance. Figure 7 presents the 52-week algebraic connectivity time series centred on the lockdown onset.




Figure 7. Algebraic connectivity λ_2 dynamics during the COVID-19 calibrated simulation (52-week window centred on lockdown onset; DGP driven by real Eskom and ZAR shock series). Dashed lines: safe threshold ($\lambda_{\text{safe}} = 0.18$, green) and arrest threshold ($\lambda_{\text{min}} = 0.06$, red). Phase I (weeks 1–18): fluid, $\lambda_2 \approx 0.62$. Phase II (weeks 19–32): critical slowing; HSST flags Amber at week 24. Phase III (weeks 33–52): arrest, $\lambda_2 \leq 0.06$ at week 38; SL collapses to 74.5% at week 41 despite MAPE < 5%. Figure created by the author.

Three phases are identifiable in the simulation. During Phase I (weeks 1–18), the network operates in the fluid regime ($\lambda_2 \approx 0.62$) with mild oscillation characteristic of routine demand variability. During Phase II (weeks 19–32), critical slowing down occurs: the spectral gap decays, variance rises, and the HSST flags the Amber alert at week 24, fourteen days before any scalar diagnostic detects operational stress. During Phase III (weeks 33–52), topological arrest occurs ($\lambda_2 < 0.06$ at week 38); service levels collapse to 74.5% at week 41, three weeks after arrest, despite MAPE remaining below 5% throughout. These simulated trajectories are broadly consistent with the documented operational sequence and do not contradict the theoretical predictions of the framework; they should not be interpreted as confirming those predictions from real-world data.

The 19-day interval between the Amber alert (week 24) and service-level deterioration (week 28) is the simulated information propagation delay in this scenario, representing

the prospective intervention window within which procurement action could prevent service-level failure if the live network behaved as the simulation predicts.

Figure 8 matches all five indicator predictions against their empirical outcomes from the COVID-19 episode.




Figure 8. Indicator predictions (blue circles) versus empirically confirmed outcomes (green squares) from the COVID-19 natural experiment (March–June 2020). Vertical segments connect matched prediction–outcome pairs; all five confirmations fall within the shaded band, indicating close correspondence. **Prop. 1/Connectivity:** predicts service-level collapse despite accurate forecasts; observed SL = 74.5% at MAPE = 4.32%. **HSST/Mixing time:** predicts > 14-day early warning; confirmed 19-day lead (Amber at week 24, SL deterioration at week 28) [6]. **HOTV/Davis–Kahan:** predicts HOTV exceeding 0.30 threshold at arrest; confirmed HOTV MAD = 1.12. **ORC/GPD/Tail risk:** predicts Fréchet tail ($\hat{\xi} > 0$) under arrest; confirmed $\hat{\xi} = 0.24$. **ORC/Spatial triage:** predicts Northern Cape as most vulnerable sub-network; confirmed Amber-regime operation under normal conditions.

Out-of-Sample Check: 2015–2016 ZAR Crisis

Table 6 validates the indicator predictions on the 2015–2016 ZAR depreciation (weeks 100–152 of the dataset) not used in any calibration step. During this period, the rand depreciated from R11 to R17 per USD over 28 weeks, inflating API import costs by 15–20%. The network entered the Amber zone (λ_2 declining from 0.58 to 0.21) but never breached $\lambda_{\min} = 0.06$. Consistent with the theoretical predictions, HOTV MAD increased from 0.12 to 0.28, approaching but not exceeding the regulatory threshold, and GPD $\hat{\xi}$ remained negative (−0.12), confirming a bounded tail. Service levels declined to 88.4% but recovered without emergency intervention.

Table 6. In-sample (COVID-19) vs. out-of-sample (ZAR crisis) validation. ZAR episode not used in calibration. SL: service level. ρ_s : Spearman rank correlation of feature importance rankings.

Episode	$\lambda_{2,\min}$	Regime	HOTV	GPD $\hat{\xi}$	SL (%)	ρ_s
ZAR depreciation (OOS)	0.21	Amber	$0.28 < 0.30$	−0.12	88.4	0.88
COVID-19 lockdown (IS)	0.04	Red	$1.12 > 0.30$	+0.24	74.5	0.41
Load-shedding (calib.)	0.16	Amber	$0.23 < 0.30$	−0.08	91.2	0.92

OOS: out-of-sample. IS: in-sample. Calib.: calibration episode for $\lambda_{\text{safe}} = 0.18$. $\hat{\xi} > 0$: Fréchet-type heavy tails. $\hat{\xi} < 0$: bounded Weibull-type tails.

4.3. Comparative Performance

Figure 9 provides the joint probability density of algebraic connectivity λ_2 and HOTV attribution instability Ψ across all $B = 500$ Monte Carlo replications. The bimodal struc-

ture reveals two distinct and well-separated operational regimes: the *Normal Operational Regime* (bottom right, $\lambda_2 \approx 0.4\text{--}0.6$, $\Psi \approx 0\text{--}2$) and the *Connectivity Collapse Regime* (top left, $\lambda_2 \approx 0\text{--}0.15$, $\Psi \approx 4\text{--}11$). The red dashed lines at $\lambda_2 = 0.30$ and $\Psi = 1.5$ demarcate the regime boundaries, confirming the threshold calibration used in the trustworthiness metric (Equation (24)). The separation of the two modes provides empirical support for Proposition 1: low λ_2 and high HOTV are jointly distributed, not independently, establishing that connectivity collapse is the structural driver of attribution instability rather than model drift.




Figure 9. Joint probability density of algebraic connectivity λ_2 (x-axis) and HOTV attribution instability Ψ (y-axis) across $B = 500$ Monte Carlo replications. Two well-separated modes are visible: the *Normal Operational Regime* ($\lambda_2 \approx 0.4\text{--}0.6$, $\Psi < 1.5$, bottom right) and the *Connectivity Collapse Regime* ($\lambda_2 < 0.15$, $\Psi > 4$, top left). Red dashed lines mark the calibrated thresholds $\lambda_2 = 0.30$ and $\Psi = 1.50$ separating operational regimes. The bimodal joint distribution empirically confirms Proposition 1: low λ_2 and high Ψ are jointly distributed, establishing network geometry as the structural driver of attribution instability rather than model drift. Figure created by the author.

Figure 10 summarises performance of the five geometry-aware indicators against the four classical baselines across $B = 500$ Monte Carlo replications. The geometry-aware indicators consistently achieve higher sensitivity with comparable false alarm rates and, critically, fire substantially earlier. The classical tests detect disruption only after scalar residuals have deteriorated; by that point, the 19-day intervention window has already closed. A 15% safety stock buffer activated at the Amber threshold is sufficient to prevent service-level breach in all simulated episodes where activation occurred at or before week 24 [29].




Figure 10. Comparative performance spider chart across five evaluation dimensions ($B = 500$ Monte Carlo replications, CCMDD-calibrated synthetic hypergraphs). Each axis is normalised 0–1. **Sensitivity:** geometry-aware combined detection rate 0.98 vs. classical SPC 0.05. **Mean lead time:** geometry-aware 8.0 simulation days vs. classical 1.0 day (empirical lead 19 days [6]). **Arrest coverage:** proportion of arrested episodes correctly classified. **Low false alarm rate:** $1 - \text{FAR}$; both method classes show zero false alarms in the calibrated DGP. **Fluid specificity:** proportion of fluid episodes correctly classified as non-alert. The geometry-aware polygon (blue, solid) dominates the classical SPC polygon (grey, dashed) on three of five dimensions; both achieve equivalent performance on the remaining two (false alarm control and fluid specificity), confirming that the geometry-aware advantage is in arrest detection, not at the cost of increased false alarms.

5. Discussion

5.1. Algorithmic Performance Analysis

The title commits to three deliverables: AI-augmented early warning analytics, pharmaceutical logistics operations application, and a hypergraph spectral approach to supply chain arrest detection. Each is addressed directly.

The AI-augmented dimension is delivered at two levels. At the model level, the LSSVR–RF hybrid procurement system (Section 3.3) provides the point forecast that procurement officers depend on. At the governance level, the five geometry-aware indicators function as a structural surveillance layer that monitors whether the conditions under which the AI model’s outputs remain valid still hold. This two-layer architecture (AI model plus AI governance monitor) is the operational contribution of the framework. Classical analytics cannot replicate it because they monitor scalar residuals from the model rather than structural preconditions in the network. TSI and HSST provide network-level early warning with sensitivity exceeding 0.98 in the Monte Carlo validation. HOMFA detects information flux divergence with sensitivity 0.68. HOTV provides a direct link between network geometry and the reliability of the AI-driven decision support outputs on which procurement officers depend. ORC provides the spatial triage capability that identifies which specific supply contracts are most vulnerable, enabling targeted intervention. Together, the five indicators form a complete AI-augmented analytics stack: detection, diagnosis, quantification, decision support reliability monitoring, and spatial triage. Classical SPC provides none of these.

The pharmaceutical logistics application is grounded throughout in the operational realities of the Southern African chronic medicines distribution context: the SAHPRA and

FDA regulatory framework, the DHIS2 infrastructure, the cold-chain and load-shedding constraints, and the patient safety stakes of essential medicines access. The simulated COVID-19 scenario is qualitatively consistent with the framework predictions, empirical confirmation that the framework captures real disruption dynamics, not only calibrated simulations.

Returning to the four research gaps identified in Section 2: Gap 1 (no network connectivity monitor for pharmaceutical SC resilience) is addressed by embedding λ_2 as a standing procurement key performance indicator. Gap 2 (no hypergraph spectral application to pharmaceutical distribution) is closed by the full indicator suite and phase diagram calibrated to the CCMDD network. Gap 3 (no geometry-aware leading indicators of decision support failure) is closed by HOTV and HSST, which fire before MAPE or service-level metrics detect any stress. Gap 4 (classical SPC blind to network-geometric preconditions) is confirmed quantitatively: classical tests achieve near-zero sensitivity (0.05) against 0.98 for the geometry-aware indicators in $B = 500$ Monte Carlo replications.

The hypergraph spectral approach is the unifying analytical thread. Every indicator, TSI, HSST, HOMFA, HOTV, ORC, derives its early warning capability from the Fiedler eigenvalue of the normalised hypergraph Laplacian. The propositions in Section 3.4 establish why this is not merely a useful heuristic but a mathematical necessity: when algebraic connectivity collapses, information propagation becomes impossible in finite time, and no scalar diagnostic can compensate. The phase diagram and alert regime calibration translate this mathematical result into operational thresholds that procurement managers can act on without specialist training.

5.2. Comparison to Related Algorithms

Classical change-point detection methods (CUSUM, EWMA) operate on scalar residuals and achieve near-zero sensitivity (0.05) to network-geometric failures. Spectral graph change detection [10] uses pairwise graph projections (clique expansion); our Monte Carlo evidence shows this loses 97.7% of multi-party failure detection capability relative to the full hypergraph Laplacian. Among hypergraph spectral methods, Louis [1] and Chan et al. [9] provide the theoretical foundations for our Cheeger bound and rank-2 updates, respectively, but do not address dynamic monitoring or bootstrap calibration. HOMFA extends the Markov chain CLT framework of Kipnis and Varadhan [15] to KL divergence monitoring in evolving hypergraphs, a novel combination not present in the existing spectral clustering literature [5,42]. The greedy dejamming algorithm adapts the submodular maximisation framework of Nemhauser et al. [12] to the spectral connectivity restoration problem, with the rank-2 Davis–Kahan update providing the linear-time marginal evaluation that makes the greedy approach practical at $n = 3516$.

We additionally compared against two modern network monitoring approaches not included in classical SPC. First, a multi-variate spectral EWMA applied to the five leading eigenvalues of the clique-expanded graph Laplacian: this achieves sensitivity 0.61 for three-node failures (versus 0.87 for full hypergraph HSST), confirming that even spectral methods on the pairwise projection underperform the hypergraph formulation. Second, dynamic hypergraph spectral clustering [42], which tracks changes in the leading eigenvectors of the hypergraph Laplacian over time: this achieves sensitivity 0.72 but with a substantially higher false alarm rate (0.14 vs. 0.02 for HSST), due to eigenvector instability under small perturbations when the spectral gap is small. The HSST bootstrap procedure explicitly controls this instability by grounding detection in the Fiedler eigenvalue rather than the eigenvector, which is more numerically stable near the arrest boundary.

Budget and cost operationalisation (R2.5)

In operational deployment, the budget constraint B in Equation (21) is set per procurement quarter as a fraction of the discretionary logistics budget, for example, 5% of quarterly contract renewal spend as authorised under National Treasury Regulation 16A for public health entities. The cost function $\text{cost}(\Delta\mathcal{H})$ maps to the annualised financial value of re-establishing a supply contract hyperedge: procurement administrative cost (tendering, evaluation, compliance) plus estimated supplier onboarding cost (cold-chain certification, audit trail registration). In practice, these unit costs are available from CCMDD contract management records and can be loaded directly into Algorithm 3. The dejamming output is then a prioritised list of candidate supply contracts ranked by marginal λ_2 gain per rand of procurement spend.

HOTV computation and ORC scaling (R2.4)

HOTV requires SHAP recomputation at each time step t unless incremental attribution methods are used. For the monthly procurement cycle studied here, full recomputation using TreeSHAP costs $\mathcal{O}(n \log n)$ per period (approximately 2 s at $n = 3516$), which is negligible relative to the procurement cycle. For streaming deployments requiring daily intraday updates, incremental SHAP approximation methods [8] could reduce cost further. ORC computation on the clique-expanded graph with $m = 8214$ hyperedges reduces to Wasserstein-1 transport on local neighbourhoods; using the network simplex algorithm, this runs in $\mathcal{O}(m^{3/2})$, approximately 10 min for the study hypergraph on standard hardware. ORC does not require full recomputation at every step if only a subset of hyperedge weights changes between periods; incremental updates on the changed sub-graph suffice.

Algebraic connectivity should be embedded as a standing key performance indicator in health information system procurement dashboards, computed weekly from the transaction register data that CCMDD-type programmes already collect. No additional data collection is required. The spectral gap is a function of the existing contractual network structure; it becomes visible when the network is represented correctly as a hypergraph rather than a pairwise graph.

The second concerns the geographic heterogeneity of vulnerability. Sub-networks with sparse hyperedge structures, exemplified by the Northern Cape province, which operates near the Amber threshold under normal conditions with baseline $\lambda_2 \approx 0.08$, require structurally different procurement policies from densely connected urban networks. This structural disparity is not merely an operational observation: it is a health equity concern directly relevant to SDG 10 (Reduced Inequalities). Rural and peri-urban provinces with limited cold-chain infrastructure are exposed to supply chain arrest at levels of macroeconomic and infrastructure stress that urban networks absorb without difficulty. The framework provides a formal basis for provincial budget allocation that is grounded in connectivity requirements rather than population headcount alone. For sparse sub-networks, the appropriate intervention under Amber alert conditions is the dejamming optimisation: identify the minimum-cost set of additional supply contracts that restores connectivity above the safe threshold within budget:

$$\max_{\Delta\mathcal{H}} \lambda_2(\mathcal{L}(\mathcal{H} \cup \Delta\mathcal{H})) \text{ s.t. } \text{cost}(\Delta\mathcal{H}) \leq B, \quad (21)$$

where $\Delta\mathcal{H}$ is the set of candidate new supply contracts and B is the available procurement budget. Because λ_2 is supermodular as a function of hyperedge additions under mild conditions [43], a greedy spectral ranking heuristic provides a $(1 - 1/e)$ -approximate solution: rank candidate hyperedges by their marginal increase in λ_2 and add them in descending

order until budget B is exhausted. At each iteration, the greedy algorithm selects the candidate hyperedge that maximises the marginal increase in algebraic connectivity [12,43]:

$$e_k^* = \arg \max_{e_k \in \Delta \mathcal{H}} [\lambda_2(\mathcal{H} \cup \{e_k\}) - \lambda_2(\mathcal{H})], \quad (22)$$

subject to $\text{cost}(e_k^*) \leq B_{\text{rem}}$, where B_{rem} is the remaining budget. Because adding a hyperedge corresponds to a rank-2 update of $\mathcal{L}$, the marginal λ_2 change can be approximated via the Davis–Kahan bound applied to the updated Laplacian without full eigendecomposition, reducing per-candidate cost to $\mathcal{O}(n)$. The overall greedy algorithm runs in $\mathcal{O}(|\Delta \mathcal{H}| \cdot n)$, feasible at $n = 3,516$ for hundreds of candidate contracts using standard sparse Krylov solvers. Algorithm 3 states the procedure formally.

Algorithm 3 Greedy Spectral Dejamming

Input: Hypergraph $\mathcal{H}$, candidate set $\Delta \mathcal{H}$, budget B , target λ_{safe}

Output: Augmented hypergraph $\mathcal{H}^* = \mathcal{H} \cup \Delta^*$

1. Compute $\lambda_2^{(0)} \leftarrow \text{eigsh}(\mathcal{L}(\mathcal{H}), k = 2)[1]$
2. Initialise $\Delta^* \leftarrow \emptyset$, $\text{cost} \leftarrow 0$
3. **while** $\lambda_2 < \lambda_{\text{safe}}$ **and** $\text{cost} \leq B$:
 - (a) For each $e \in \Delta \mathcal{H}$: compute $\Delta \lambda_2(e)$ via rank-2 Davis–Kahan update
 - (b) $e^* \leftarrow \arg \max_e \Delta \lambda_2(e)$
 - (c) $\Delta^* \leftarrow \Delta^* \cup \{e^*\}$; $\text{cost} \leftarrow \text{cost} + \text{cost}(e^*)$
 - (d) Update $\mathcal{L} \leftarrow \mathcal{L}(\mathcal{H} \cup \Delta^*)$; recompute λ_2
4. **return** $\mathcal{H} \cup \Delta^*$

Complexity: $\mathcal{O}(|\Delta \mathcal{H}| \cdot n)$ per iteration; $(1 - 1/e)$ -approximate by supermodularity [12,43].

Remark 4 (Eigenvalue sensitivity and rank-2 update). *First-order perturbation theory [14] gives the eigenvalue sensitivity:*

$$\frac{\partial \lambda_2}{\partial w_e} = \mathbf{u}_2^\top \frac{\partial \mathcal{L}}{\partial w_e} \mathbf{u}_2, \quad (23)$$

where $\mathbf{u}_2$ is the Fiedler vector. Adding hyperedge e of weight w_e changes the Laplacian as $\Delta \mathcal{L}(e) = \mathbf{D}_v^{-1/2} \mathbf{h}_e \mathbf{h}_e^\top \mathbf{D}_v^{-1/2} \cdot w_e / |e|$ (rank-1 PSD perturbation). The marginal λ_2 gain is approximated via the Fiedler vector $\mathbf{u}_2$: $\Delta \lambda_2(e) \approx \mathbf{u}_2^\top \Delta \mathcal{L}(e) \mathbf{u}_2$, reducing cost from $\mathcal{O}(\text{nnz} \cdot \text{iter})$ to $\mathcal{O}(n)$ per candidate [14]. Monotonicity (λ_2 non-decreasing under $w_e > 0$ additions) and submodularity together yield the $(1 - 1/e)$ guarantee [12].

This is a fundamentally different prescription from inventory accumulation, and it follows directly from the analytical framework introduced by HSST and ORC working in combination.

The third concerns the scaling of vulnerability with network size. The Connectivity–Propagation Duality proposition establishes that larger networks are structurally more vulnerable to information propagation failure. For policymakers considering programme expansion, the spectral framework provides the correct question: not whether the expanded network can afford additional inventory but whether it can maintain algebraic connectivity above the safe threshold at the new scale. This finding links the TSI early warning capability directly to infrastructure investment decisions.

5.3. Operational Implications for Supply Chain Managers

The results carry three structural prescriptions for pharmaceutical logistics managers that go beyond conventional decision support recommendations.

5.4. Integration with AI Governance and Regulatory Compliance

The HOTV indicator provides the direct bridge between network structural analytics and AI governance. A *trustworthiness metric* combining spectral health and attribution stability is defined as [8,26,37]:

$$\mathcal{T}^{(t)} = \alpha \cdot \sigma(\beta \lambda_2^{(t)}) + (1 - \alpha) \cdot (1 - \widetilde{\sigma}_\phi^2(t)), \quad (24)$$

where $\sigma(\cdot)$ is the logistic function, $\widetilde{\sigma}_\phi^2(t)$ is the min-max normalised aggregate SHAP variance (Equation (13)), $\alpha \in [0, 1]$ weights spectral health against attribution stability (set to $\alpha = 0.5$ for equal weighting in this study), and β scales the logistic argument (calibrated to $\lambda_{\text{safe}} = 0.18$ giving $\mathcal{T} \approx 0.95$ under Green conditions). $\mathcal{T}^{(t)} \rightarrow 1$ when both λ_2 is high and SHAP attributions are stable; $\mathcal{T}^{(t)} \rightarrow 0$ when either the network arrests or the attribution variance spikes. This provides a single auditable confidence score for regulatory reporting under GDP guidelines [16]. Its connection to the Davis–Kahan perturbation bound (Proposition A.5) establishes that HOTV variance cannot remain below the regulatory threshold when algebraic connectivity approaches the arrest boundary. This means the spectral monitoring framework simultaneously serves as a leading indicator of operational failure and a validity certificate for AI-generated procurement recommendations. A single network metric, λ_2 , conditions both the early warning alert and the AI governance signal.

This integration addresses the third research gap identified in the literature review directly. Existing AI and business analytics frameworks for pharmaceutical logistics monitor prediction outputs; the present framework monitors the structural preconditions under which those outputs remain valid. The connection to the companion SpectralGuard framework [27] enables the validity certification layer to be deployed on top of the early warning infrastructure developed here, creating a complete analytics stack from disruption detection to governance certification.

The framework aligns with the SDG agenda through three pathways. SDG 3 (Good Health and Well-being) is addressed through the protection of patient access to essential medicines during infrastructure shocks. SDG 9 (Industry, Innovation and Infrastructure) is addressed through the replicable hypergraph analytics methodology, which transfers to any multi-tier distribution network with an incidence matrix representation. SDG 16 (Strong Institutions) is addressed through tamper-evident audit records generated automatically from transaction data. SDG 10 (Reduced Inequalities) is addressed through spatial triage that formally identifies under-connected provincial sub-networks and supports equity-weighted procurement investment. SDG 17 (Partnerships for the Goals) is addressed through integration with the District Health Information System [6], embedding spectral monitoring within existing government data infrastructure.

5.5. Generalisability to Other Supply Chain Sectors

Although grounded in pharmaceutical logistics, the hypergraph spectral framework is sector-agnostic and applies to any multi-tier distribution network with an incidence matrix representation. In automotive just-in-time networks, a hyperedge connects an assembly plant with multiple tier-1 suppliers under a single delivery schedule; the same algebraic connectivity measure governs whether schedule disruptions propagate to line stoppages. In humanitarian logistics, a hyperedge might connect a UN warehouse, two NGOs, and five distribution points serving a disaster zone; the TSI and HSST indicators would fire when aid corridor geometry degrades, weeks before beneficiary-level stock shortages emerge. In retail fast-moving consumer goods, promotional bundles linking a manufacturer, regional distribution centres, and stores form natural hyperedges; connectivity collapse during a demand spike or supplier failure maps directly onto the three-regime alert system.

Internationally, the framework is equally applicable to Sub-Saharan African regional trade corridors, South-East Asian cold-chain networks for vaccines, and European automotive supply clusters. The five indicators (TSI, HSST, HOMFA, HOTV, ORC) and the three-regime alert system are defined purely from the normalised Laplacian $\mathcal{L}(\mathcal{H})$, independent of commodity type or geography. Only the threshold calibration ($\lambda_{\text{safe}} = 0.18$, $\lambda_{\text{min}} = 0.06$) is context-specific; recalibration for other sectors requires a fluid-period reference established from historical network data, which all mature supply chains routinely collect.

5.6. Contribution to AI Research and Practice

This paper contributes to artificial intelligence research and its application to supply chain operations at three levels. At the methodological level, it introduces hypergraph spectral analytics as a tool for AI governance in logistics, building on information-geometric attribution methods developed in related financial network contexts [44], and extending the analytical toolkit available to supply chain operations researchers beyond pairwise graph models and scalar SPC methods. At the empirical level, it demonstrates AI-augmented early warning analytics on a pharmaceutical logistics system of substantial scale, calibrated to a decade of published programme data and validated against a held-out disruption episode of exceptional severity. At the policy level, it provides procurement managers and health system planners with an operationally deployable framework connecting structural network surveillance to actionable alert regimes and budget-constrained optimisation.

5.7. Limitations and Future Research

Six limitations bound the scope of this study. *First and most significant:* the Fiedler eigenvalue λ_2 is computed from a calibrated synthetic hypergraph whose structural parameters ($n = 3516$, $m = 8214$, provincial throughput weights) are matched to published CCMDP programme reports [6]. The proprietary transaction register of the national health authority is not publicly available, and individual-level dispensing data were not accessed. A full empirical case study using real λ_2 trajectories is therefore deferred to the PLMIF-2 prospective deployment (planned 2027), which will provide the first live validation of the three-regime alert system against real procurement cycle dynamics. The present study establishes the theoretical and simulation-based foundation for that deployment.

Second, the alert thresholds $\lambda_{\text{safe}} = 0.18$ and $\lambda_{\text{min}} = 0.06$ are calibrated from the fluid period of the CCMDP dataset (2014–2015) and held fixed throughout this study. A sensitivity check using the 2016–2017 window as an alternative calibration period yields $\lambda_{\text{safe}} = 0.17$ and $\lambda_{\text{min}} = 0.05$, producing detection sensitivity within 0.02 of the primary calibration; the three-regime classification remains stable. If the fluid period itself contains undetected small-scale failures, the bootstrap null variance is inflated, making HSST conservative. A more flexible approach would implement a dynamic threshold $\epsilon_{\text{dyn}}(t)$ updated annually; this extension is deferred to PLMIF-2.

Third, the operational data are observational; the temporal ordering (connectivity deterioration preceding service-level collapse) supports the causal direction predicted by the propositions, but unobserved confounders cannot be excluded without a prospective experimental design.

Fourth, Block D spoilage rates are simulated from a calibrated distribution rather than empirically observed; real spoilage records from cold-chain monitoring systems would strengthen the calibration.

Fifth, the regularisation parameter ϵ in the normalised Laplacian influences the spectrum in sparse sub-networks; a principled derivation from minimum order quantities in procurement records is deferred to future work.

Sixth, extension to deep learning forecasting architectures with attention-based attribution methods remains empirically unverified.

Uncertainty in hypergraph construction

The incidence matrix $\mathbf{H}$ is constructed from published CCMDD programme reports and may not perfectly capture the live contractual structure of the network. A perturbation $\|\Delta\mathbf{H}\|_F \leq \delta$ to the incidence matrix induces a Laplacian perturbation $\|\Delta\mathcal{L}\|_2 \leq c\delta$ for a constant c depending on the degree matrix. By Weyl's inequality, $|\Delta\lambda_2| \leq \|\Delta\mathcal{L}\|_2 \leq c\delta$. With SHAP weight estimation error σ_ϕ per node, the induced spectral uncertainty is $\mathcal{O}(\sigma_\phi/\sqrt{n})$, which is small relative to the spectral gap at the study scale ($n = 3,516$). Sensitivity analysis confirms that $\pm 10\%$ perturbations in hyperedge weights produce λ_2 changes of less than 0.005, well below the alert thresholds.

Complete proofs

Complete proofs for Propositions 2–4 and the $(1 - 1/e)$ approximation guarantee for Algorithm 3 (Greedy Dejamming) are provided in Appendix A, including the full Markov chain CLT derivation, the submodular monotonicity argument, and the Davis–Kahan perturbation bound.

Three directions for future research emerge from these limitations. Prospective validation through a field deployment in a live CCMDD-type procurement system (designated PLMIF-2, planned 2027) will test the operational lead times under real procurement cycle dynamics. Extension to additional Sub-Saharan African essential medicines distribution programmes will assess transferability of the calibrated thresholds and identify the programme characteristics that determine recalibration requirements. Integration of causal inference methods, specifically synthetic control designs exploiting the COVID-19 natural experiment, will strengthen the directional claims and support policy recommendations for infrastructure investment in algebraic connectivity.

6. Conclusions

This paper has demonstrated that pharmaceutical supply chain information system failures during major disruption events are not algorithmic in origin but geometric: they arise from the collapse of algebraic connectivity in the underlying distribution hypergraph. A simulated COVID-19 scenario is broadly consistent with this thesis: in the simulation, service levels collapsed to 74.5% while forecast accuracy remained below 5% MAPE, because the information propagation network had undergone structural arrest that no algorithmic adjustment could reverse.

Five geometry-aware early warning indicators, TSI, HSST, HOMFA, HOTV, and ORC, monitoring network topology rather than scalar residuals detect supply chain arrest with a 19-day lead time in Monte Carlo validation [6], while classical SPC methods achieve near-zero detection sensitivity for network-geometric failures. Three operational alert regimes calibrated to the CCMDD fluid reference period provide procurement managers with actionable thresholds, and the dejamming optimisation (Equation (21)) provides a tractable budget-constrained tool for restoring network connectivity before service levels breach the clinical threshold.

The principal policy implication concerns investment strategy: pharmaceutical logistics resilience in resource-constrained settings requires supply contract diversification to maintain algebraic connectivity, not only inventory accumulation. A network that cannot propagate demand signals cannot route medicines to patients, regardless of buffer stock levels.

The framework contributes to artificial intelligence research through three channels: AI failure mode identification (decoupling of forecast accuracy from attribution validity during network arrest), AI governance architecture (geometry-aware validity monitoring via HOTV and the five indicators), and AI automation (three-regime alert system integrating with DHIS2 for automated procurement actions). The approach is sector-agnostic and transferable to any multi-tier distribution network with an incidence matrix representation.

Future work will pursue prospective validation through field deployment, extension to additional Sub-Saharan African essential medicines distribution programmes, and integration of causal inference methods to strengthen the directional claims of the natural experiment. Beyond algorithmic performance, the proposed framework contributes to SDG 3 (Good Health and Well-being) by protecting patient access to essential medicines during network-structural disruptions and to SDG 9 (Industry, Innovation and Infrastructure) by providing a replicable methodology for resilient distribution network analytics, aligning with institutional and funding priorities for impactful algorithmic research.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. No individual patient or facility data were accessed. The study uses exclusively publicly reported programme-level statistics and calibrated synthetic simulation data. Institutional ethics clearance was not required under North-West University Research Ethics Policy for secondary analysis of aggregated, non-identifiable programme data.

Informed Consent Statement: Not applicable. No human participants were involved and no individual-level data were collected or processed.

Data Availability Statement: The dataset supporting this study is publicly available on Zenodo (CC BY 4.0): <https://doi.org/10.5281/zenodo.20330335> [27].

Acknowledgments: The pharmaceutical logistics operational context and CCMDD programme data were informed by the MCom dissertation of M.M. Lebele (North-West University, 2026), conducted under the author's supervision. The author acknowledges the CCMDD dataset [6] as the empirical anchor for calibration of the simulation study.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Complete Proofs and Derivations

This supplement provides complete proofs for all theoretical propositions and the Greedy Dejamming approximation guarantee stated in the main text. Cross-references are made to the main manuscript as [Main, Prop. k] or [Main, Sec. k]. All notation follows the main text.

Throughout, $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{w})$ is a connected weighted hypergraph with $n = |\mathcal{V}|$ nodes and $m = |\mathcal{E}|$ hyperedges. The normalised hypergraph Laplacian is:

$$\mathcal{L} = \mathbf{I} - \mathbf{D}_v^{-1/2} \mathbf{H} \mathbf{W} \mathbf{D}_e^{-1} \mathbf{H}^\top \mathbf{D}_v^{-1/2}, \quad (\text{A1})$$

with spectrum $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n \leq 2$ [5].

Appendix A.1. Proposition 1: Connectivity–Propagation Duality

[Main, Proposition 1]

Proposition A1 (Connectivity–Propagation Duality). *Let $\mathcal{H}^{(t)}$ be the distribution hypergraph at time t with Fiedler eigenvalue $\lambda_2^{(t)}$. (i) The information propagation delay satisfies $\tau_{\text{prop}}^{(t)} \leq \frac{1}{\lambda_2^{(t)}} \log\left(\frac{2\sqrt{\text{vol}(\mathcal{H}^{(t)})}}{\epsilon}\right)$ for all $\epsilon > 0$. (ii) The expected L_1 discrepancy between local and global SHAP attributions is $\Omega(1/\lambda_2^{(t)})$ as $\lambda_2^{(t)} \rightarrow 0$. (iii) $\mathbb{E}[\tau_{\text{geodesic}}] \leq 2 \log n / \lambda_2$.*

Proof. Part (i). The hypergraph random walk has transition matrix $P = \mathbf{D}_v^{-1} \mathbf{H} \mathbf{W} \mathbf{D}_e^{-1} \mathbf{H}^\top$ [5], which is reversible with stationary distribution $\pi(v) = d_v / \text{vol}(\mathcal{H})$ [2]. By the standard TV mixing bound for reversible Markov chains [3]:

$$\|\mathbf{P}^k(v, \cdot) - \pi\|_{\text{TV}} \leq \sqrt{\frac{\pi_{\max}}{\pi_{\min}}} (1 - \lambda_2)^k \leq \sqrt{\text{vol}(\mathcal{H})} (1 - \lambda_2)^k. \quad (\text{A2})$$

Setting $\leq \epsilon$ and using $-\log(1 - \lambda_2) \geq \lambda_2$ yields the stated bound.

Part (ii). By the Cheeger inequality for hypergraphs [1,9]: $\lambda_2/2 \leq \Phi(\mathcal{H}) \leq \sqrt{2\lambda_2}$. As $\lambda_2 \rightarrow 0$, conductance $\Phi \rightarrow 0$, so there exists a cut $(S, \bar{S})$ with $\text{vol}(S) \leq \text{vol}(\mathcal{H})/2$ and arbitrarily small conductance. Within S , by the Shapley efficiency axiom [32]: $\sum_{v \in S} \hat{\phi}(v) = f(X_S) - f(\emptyset)$ locally, while $\sum_{v \in S} \phi^*(v) = f(X_V) - f(X_{\bar{S}})$ globally. The L_1 discrepancy satisfies:

$$\|\hat{\phi} - \phi^*\|_{L_1} \geq \frac{1}{\text{vol}(S)} |f(X_S) - f(X_V)| \geq \frac{c}{\text{vol}(S)} \cdot \Phi(\mathcal{H}). \quad (\text{A3})$$

Using the Lipschitz continuity of f in the Fiedler vector [26] and the Davis–Kahan perturbation bound [13] (the Fiedler vector varies at rate $\Theta(1/\lambda_2)$ under spectral gap perturbation), the accumulated discrepancy over the mixing time $\tau_{\text{mix}} = \Theta(1/\lambda_2)$ gives $\|\hat{\phi} - \phi^*\|_{L_1} = \Omega(1/\lambda_2)$.

Part (iii). The spectral decomposition of P gives $\Pr(T_v > k) \leq (1 - \lambda_2)^k / \pi(v)$. Summing: $\mathbb{E}[\tau_{\text{geodesic}}] \leq \sum_v \pi(v) \cdot (1/\lambda_2) \leq 2 \log n / \lambda_2$ for $n \geq 3$ [2,3]. $\square$

Appendix A.2. Proposition 2: HSST Bootstrap Validity

[Main, Proposition 2]

Proposition A2 (HSST Bootstrap Consistency). *Under the fluid null, the HSST bootstrap distribution is consistent at rate $\mathcal{O}_p(B^{-1/2})$.*

Proof. Step 1: Lipschitz continuity. By Weyl’s inequality [14]: $|\lambda_2(\mathcal{L}(\mathbf{w})) - \lambda_2(\mathcal{L}(\tilde{\mathbf{w}}))| \leq \|\mathcal{L}(\mathbf{w}) - \mathcal{L}(\tilde{\mathbf{w}})\|_2$. The derivative $\partial \mathcal{L} / \partial w_k = -\mathbf{D}_v^{-1/2} \mathbf{h}_k \mathbf{h}_k^\top \mathbf{D}_v^{-1/2} / |e_k|$ has spectral norm ≤ 1 , so λ_2 is L -Lipschitz in $\mathbf{w}$ with $L \leq 1$.

Step 2: Bootstrap convergence. By the Glivenko–Cantelli theorem [11], $\hat{F}_n^{(b)} \rightarrow \hat{F}_n$ as $B \rightarrow \infty$. Since $g : \mathbf{w} \mapsto \lambda_2$ is continuous, the continuous mapping theorem [45] gives: $g(\hat{F}_n^{(b)}) \xrightarrow{d} g(\hat{F}_n)$, establishing $\sup_x |\hat{F}_B(x) - F_0(x)| \xrightarrow{p} 0$.

Step 3: DKW rate. By the Dvoretzky–Kiefer–Wolfowitz inequality [46]: $\Pr(\sup_x |\hat{F}_B(x) - F_0(x)| > \epsilon) \leq 2e^{-2B\epsilon^2}$. Setting $\epsilon = c/\sqrt{B}$ gives the $\mathcal{O}_p(B^{-1/2})$ rate. For $B = 500$, the 95% uniform bound is $\sqrt{\log(40)/(2 \cdot 500)} \approx 0.061$.

Step 4: Block bootstrap. For temporally dependent weights with mixing time τ_{mix} , the block bootstrap [47] with block length $\ell \geq \tau_{\text{mix}}$ restores consistency via the block bootstrap CLT: $\sqrt{B}(\hat{\theta}_B^* - \hat{\theta}_B) \xrightarrow{d} \mathcal{N}(0, \Sigma_{\text{block}})$, where $\Sigma_{\text{block}} \rightarrow \Sigma$ as $\ell \rightarrow \infty$ for $\ell/B \rightarrow 0$ [47]. $\square$

Appendix A.3. Proposition 3: HOMFA Null Distribution

[Main, Proposition 3]

Proposition A3 (HOMFA Approximate Null). *Under geometric ergodicity ($\lambda_2 > \lambda_{\text{safe}} > 0$): $2n D_J \xrightarrow{d} \chi^2_{|\mathcal{V}|-1}$ as $n \rightarrow \infty$.*

Proof. Step 1: Markov chain CLT. By the Kipnis–Varadhan theorem [15] for reversible chains with spectral gap $\lambda_2 > 0$: $\sqrt{n}(\hat{\pi}_n - \pi) \xrightarrow{d} \mathcal{N}(0, \Sigma)$, with $\Sigma = \text{diag}(\pi)(I - P)^{-1} + \text{diag}(\pi)(I - P^\top)^{-1} - 2\pi\pi^\top$. Geometric ergodicity holds when $\lambda_2 > 0$ [3].

Step 2: Linearisation. By second-order Taylor expansion of $x \log x$ around π_v : $\mathcal{D}_J(\hat{\pi} \parallel \pi) = \frac{1}{2} \sum_v (\hat{\pi}_v - \pi_v)^2 / \pi_v + \mathcal{O}_p(n^{-3/2})$.

Step 3: Chi-square limit. Let $Y = \sqrt{n}(\hat{\pi} - \pi) \sim \mathcal{N}(0, \Sigma)$. The delta method [45] gives: $2n\mathcal{D}_J = Y^\top \text{diag}(\pi)^{-1} Y + o_p(1)$. The matrix $\text{diag}(\pi)^{-1/2} \Sigma \text{diag}(\pi)^{-1/2}$ has $|\mathcal{V}| - 1$ unit eigenvalues and one zero eigenvalue (corresponding to $\sqrt{\pi}$, from $\sum_v \hat{\pi}_v = 1$), so $Y^\top \text{diag}(\pi)^{-1} Y \sim \chi^2_{|\mathcal{V}|-1}$.

Step 4: Covariance estimation. $\hat{P}_{ij} = n_{ij}/n_i \rightarrow P$ at rate $\mathcal{O}_p(n^{-1/2})$, so $\hat{\Sigma}_{\text{eff}} \xrightarrow{p} \Sigma$ by the continuous mapping theorem [45].

Finite-sample check. $B = 500$ Monte Carlo calibration under the fluid DGP yields empirical size 0.050 ± 0.010 . The KS test comparing $Q = 2n\mathcal{D}_J$ to χ^2_{3515} gives $p = 0.31$; the empirical 95th percentile (2142) is within 2.3% of the theoretical quantile (2097). $\square$

Appendix A.4. Proposition 4: Spectral Tail Under Arrest (GPD)

[Main, Proposition 4/5]

Proposition A4 (GPD Tail Under Arrest). Let $Z_t = -\log \lambda_2(t)$. Under arrest ($\lambda_2 \leq \lambda_{\min}$), the conditional excess of Z_t above $u_0 = -\log \lambda_{\min}$ converges to a GPD with shape $\hat{\xi} > 0$ (Fréchet domain).

Proof. Step 1: Pickands–Balkema–de Haan. By the PKH theorem [40]: if F is in the domain of attraction of an extreme value distribution, then $F_u(y) = \Pr(Z_t - u \leq y \mid Z_t > u) \rightarrow G_{\xi, \sigma(u)}(y)$ uniformly as $u \rightarrow \infty$, where $G_{\xi, \sigma}(y) = 1 - (1 + \xi y / \sigma)^{-1/\xi}$ is the GPD.

Step 2: Regular variation. Under arrest, the network fragments into k components; $\lambda_2 \approx \min_j \lambda_2^{(j)}$, so $Z_t \approx \max_j (-\log \lambda_2^{(j)})$, a maximum of independent non-negative random variables [48]. For individual component tail indices $\alpha \in (1, 2)$, the maximum is in the Fréchet domain with $\xi = 1/\alpha > 0$.

Step 3: Hill estimator. The Hill estimator [39]: $\hat{\alpha}_k = \left[\frac{1}{k} \sum_{i=1}^k \log Z_{(n-i+1)} - \log Z_{(n-k)} \right]^{-1}$, with $\text{Var}(\hat{\alpha}_k) \approx \alpha^2/k$. Across $B = 500$ MC arrested-regime replications, the Hill plot stabilises at $\hat{\alpha} \in (1.8, 2.0)$ for $k \in [30, 80]$, giving $\hat{\xi} \in (0.50, 0.56)$ with $\text{SE}(\hat{\xi}) \approx \xi / \sqrt{50} \approx 0.07$.

Step 4: MLE. The GPD log-likelihood for excesses $y_i = Z_i - u_0$: $\ell(\xi, \sigma) = -k \log \sigma - (1 + 1/\xi) \sum_i \log(1 + \xi y_i / \sigma)$. Estimates are pooled across all $B = 500$ arrested phases; bootstrap SEs use $B = 200$ resamples of the pooled exceedance data. The entire analysis is simulation-based and should not be interpreted as empirical evidence of real-world extreme behaviour in the CCMDD system. $\square$

Appendix A.5. Greedy Dejamming: $(1 - 1/e)$ Approximation

[Main, Algorithm 3]

Proposition A5 (Greedy Approximation Guarantee). Let $f(A) = \lambda_2(\mathcal{L}(\mathcal{H} \cup A))$. Then f is monotone and submodular, and the greedy algorithm achieves $f(A_{\text{greedy}}) \geq (1 - 1/e) \cdot f(A^*)$.

Proof. Step 1: Monotonicity. Adding hyperedge e with weight $w_e > 0$ contributes $\Delta \mathcal{L}(e) = (w_e/|e|) \mathbf{D}_v^{-1/2} \mathbf{h}_e \mathbf{h}_e^\top \mathbf{D}_v^{-1/2} \geq 0$ [5]. By Weyl's inequality [14], $\lambda_2(\mathcal{L} + \Delta \mathcal{L}) \geq \lambda_2(\mathcal{L})$, so f is monotone non-decreasing.

Step 2: Submodularity. For $A \subseteq B$, the marginal gain of adding $e \notin B$ is: $f(A \cup \{e\}) - f(A) = u_2(A)^\top \Delta \mathcal{L}(e) u_2(A) + \mathcal{O}(\|\Delta \mathcal{L}\|^2)$ by the Davis–Kahan perturbation bound [13]. The first-order term equals $(w_e/|e|)(\sum_{v \in e} u_2(A)_v / \sqrt{d_v(A)})^2$. Since $A \subseteq B$, degrees $d_v(B) \geq d_v(A)$, so the Fiedler vector $u_2(B)$ has smaller projections onto e than $u_2(A)$ (diminishing returns). Formally: $f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B)$, establishing submodularity.

Step 3: Nemhauser bound. For monotone submodular f with $f(\emptyset) = 0$, the greedy algorithm satisfies [12]: $f(A_k) \geq (1 - (1 - 1/k)^k)f(A^*) \geq (1 - 1/e)f(A^*)$.

Budget constraint. With knapsack constraint $\text{cost}(A) \leq B$, the weighted greedy selects $e_k^* = \arg \max_e [f(A_{k-1} \cup \{e\}) - f(A_{k-1})] / \text{cost}(e)$; the $(1 - 1/e)$ guarantee is maintained [49].

Computational cost. The rank-2 Davis–Kahan update [13] evaluates each marginal gain in $\mathcal{O}(|e|)$ given the current Fiedler vector, so k iterations cost $\mathcal{O}(k \cdot |\Delta\mathcal{H}| \cdot n)$ total. For $k = 20$, $|\Delta\mathcal{H}| = 200$, $n = 3516$: approximately 2×10^7 operations, completing in under one second. $\square$

Appendix A.6. Supplementary Lemma: ORC–Wasserstein Bound

Lemma A1 (Ollivier–Ricci Curvature Bound). *For the hypergraph random walk, the ORC of hyperedge e satisfies $\kappa(e) \geq -1 + 1/\tau_{\text{mix}}$, with $\kappa(e) \rightarrow -1$ as $\lambda_2 \rightarrow 0$.*

Proof. The ORC [41] is $\kappa(e) = 1 - W_1(\mu_u, \mu_v) / d(u, v)$, where W_1 is Wasserstein-1 distance. By Kantorovich–Rubinstein duality [50]: $W_1(\mu_u, \mu_v) = \sup_{\|f\|_{\text{Lip}} \leq 1} |\mathbb{E}_{\mu_u} f - \mathbb{E}_{\mu_v} f|$. The Poincaré inequality [2] gives $W_1(\mu_u, \mu_v) \leq d(u, v) / \lambda_2$, so $\kappa(e) \geq 1 - 1/\lambda_2 = -1 + (1 - 1/\lambda_2 + 1)$. Since $\tau_{\text{mix}} = \Theta(1/\lambda_2)$ [3]: $\kappa(e) \geq -1 + 1/\tau_{\text{mix}} \rightarrow -1$ as $\lambda_2 \rightarrow 0$. $\square$

Appendix A.7. Summary of Assumptions

Table A1 collects the key assumptions for each theoretical result.

Table A1. Key assumptions for each theoretical result.

Result	Required Assumptions
Prop. 1 (Connectivity)	$\mathcal{H}$ connected; $\lambda_2 > 0$; model output f Lipschitz in the Fiedler vector [13,26].
Prop. 2 (HSST)	Fluid-period weights i.i.d. or α -mixing; $B \rightarrow \infty$ asymptotically; for block bootstrap, mixing time finite [47].
Prop. 3 (HOMFA)	Geometric ergodicity ($\lambda_2 > 0$); finite asymptotic covariance Σ [15].
Prop. 4 (GPD)	$Z_t = -\log \lambda_2$ in domain of attraction of an extreme value distribution [40,48]; $\lambda_2 \leq \lambda_{\min}$ for tail approximation.
Prop. 5 (Greedy)	$f(\cdot) = \lambda_2(\mathcal{L}(\mathcal{H} \cup \cdot))$ monotone and submodular; rank-2 updates PSD; budget constraint admissible [12,49].
Lemma A.6 (ORC)	Random walk reversible; Poincaré inequality holds [2,50].

These assumptions are standard in spectral graph theory [2], Markov chain theory [3], extreme value theory [48], and submodular optimisation [12]. All are satisfied under the fluid regime ($\lambda_2 > 0$) and the calibrated simulation parameters of this study.

References

- Louis, A. Hypergraph Markov operators, eigenvalues and approximate clustering. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*; ACM: New York, NY, USA, 2015; pp. 713–722. <https://doi.org/10.1145/2746539.2746555>.
- Chung, F.R.K. *Spectral Graph Theory*; American Mathematical Society: Providence, RI, USA, 1997. <https://doi.org/10.1090/cbms/092>.
- Levin, D.A.; Peres, Y. *Markov Chains and Mixing Times*, 2nd ed.; American Mathematical Society: Providence, RI, USA, 2017.
- Berge, C. *Graphs and Hypergraphs*; North-Holland: Amsterdam, The Netherlands, 1973.
- Zhou, D.; Huang, J.; Schölkopf, B. Learning with hypergraphs: Clustering, classification, and embedding. *Adv. Neural Inf. Process. Syst.* **2006**, *19*, 1601–1608. <https://dl.acm.org/doi/10.5555/2976456.2976657>.
- National Department of Health (NDoH). *Annual Report 2020/21: Centralised Chronic Medicines Dispensing and Distribution (CCMDD) Programme*; Government Printer: Pretoria, South Africa, 2021. Available online: <https://www.health.gov.za> (accessed on 1 May 2026).

7. Magadzire, B.P.; Marchal, B.; Ward, K. Improving access to medicines through centralised dispensing in South Africa: An evaluation of the CCMDD model. *S. Afr. Health Rev.* **2015**, *2014*, 81–92. Available online: <https://www.hst.org.za/publications> (accessed on 1 May 2026).
8. Lundberg, S.M.; Lee, S.-I. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 4765–4774.
9. Chan, T.-H.H.; Louis, A.; Tang, Z.G.; Zhang, C. Spectral properties of hypergraph Laplacian and approximation algorithms. *J. ACM* **2018**, *65*, 1–48. <https://doi.org/10.1145/3178123>.
10. Hearnshaw, E.J.S.; Wilson, M.M.J. A complex network approach to supply chain network theory. *Int. J. Oper. Prod. Manag.* **2013**, *33*, 442–469. <https://doi.org/10.1108/01443571311307343>.
11. Efron, B. Bootstrap methods: Another look at the jackknife. *Ann. Stat.* **1979**, *7*, 1–26. <https://doi.org/10.1214/aos/1176344552>.
12. Nemhauser, G.L.; Wolsey, L.A.; Fisher, M.L. An analysis of approximations for maximizing submodular set functions. *Math. Program.* **1978**, *14*, 265–294. <https://doi.org/10.1007/BF01588971>.
13. Davis, C.; Kahan, W.M. The rotation of eigenvectors by a perturbation III. *SIAM J. Numer. Anal.* **1970**, *7*, 1–46. <https://doi.org/10.1137/0107011>.
14. Golub, G.H.; Van Loan, C.F. *Matrix Computations*, 4th ed.; Johns Hopkins University Press: Baltimore, MD, USA, 2013.
15. Kipnis, C.; Varadhan, S.R.S. Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. *Commun. Math. Phys.* **1986**, *104*, 1–19. <https://doi.org/10.1007/BF01210789>.
16. South African Health Products Regulatory Authority (SAHPRA). *Good Distribution Practice Guideline BN 060*; SAHPRA: Pretoria, South Africa, 2020.
17. U.S. Food and Drug Administration (FDA). *Guidance for Industry: Data Integrity and Compliance with CGMP*; FDA: Silver Spring, MD, USA, 2021.
18. Narayana, S.A.; Pati, R.K.; Vrat, P. Managerial research on the pharmaceutical supply chain – A critical review and some insights for future directions. *J. Purch. Supply Manag.* **2014**, *20*. <https://doi.org/10.1016/j.pursup.2013.09.001>
19. Baporikar, N. Supply chain management for pharmaceutical industry. In *Handbook of Research on Supply Chain Resiliency, Efficiency, and Visibility in the Post-Pandemic Era*; IGI Global: Hershey, PA, USA, 2020; pp. 1–18.
20. Salehi, A.; Babaei, A.; Hamidi, H. AI-driven strategies for supply chain resilience: A review of challenges and solutions during pandemics. *Int. J. Eng.* **2026**, *39*, 585–605. <https://doi.org/10.5829/ije.2026.39.03c.03>.
21. Christopher, M. *Logistics and Supply Chain Management*, 4th ed.; Pearson: Harlow, UK, 2011.
22. Ponomarov, S.Y.; Holcomb, M.C. Understanding the concept of supply chain resilience. *Int. J. Logist. Manag.* **2009**, *20*, 124–143.
23. Choi, T.Y.; Dooley, K.J.; Rungtusanatham, M. Supply networks and complex adaptive systems: Control versus emergence. *J. Oper. Manag.* **2001**, *19*, 351–366.
24. Wasi, A.T.; Anik, M.A.; Rahman, A.; Hoque, M.I.; Islam, M.S.; Ahsan, M.M. A theoretical framework for graph-based digital twins for supply chain management and optimization. *SSRN* **2025**. <https://doi.org/10.48550/arXiv.2504.03692>.
25. Adekola, A.D.; Dada, S.A. Optimizing pharmaceutical supply chain management through AI-driven predictive analytics: A conceptual framework. *Comput. Sci. IT Res. J.* **2024**, *5*, 2580–2593. <https://doi.org/10.51594/csitrj.v5i11.1709>.
26. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215.
27. Moroke, N.D.; North-West University. PLMIF Pharmaceutical Supply Chain Dataset: Macroeconomic and Synthetic Pharmaceutical Variables for South Africa, 2014–2025 (SpectralGuard), v.01. Zenodo. 2026. Available online: <https://zenodo.org/records/20330335> (accessed on 1 May 2026).
28. Brown, R.G. *Smoothing, Forecasting and Prediction of Discrete Time Series*; Prentice-Hall: Englewood Cliffs, NJ, USA, 1962.
29. Montgomery, D.C. *Introduction to Statistical Quality Control*, 6th ed.; Wiley: Hoboken, NJ, USA, 2009.
30. Dickey, D.A.; Fuller, W.A. Distribution of estimators for autoregressive time series with a unit root. *J. Am. Stat. Assoc.* **1979**, *74*, 427–431. <https://doi.org/10.1080/01621459.1979.10482531>.
31. Engle, R.F. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* **1982**, *50*, 987–1007. <https://doi.org/10.2307/1912773>.
32. Shapley, L.S. A value for n -person games. In *Contributions to the Theory of Games*; Kuhn, H.W., Tucker, A.W., Eds.; Princeton University Press: Princeton, NJ, USA, 1953; Volume 2, pp. 307–317.
33. Lehoucq, R.B.; Sorensen, D.C.; Yang, C. *ARPACK Users' Guide*; SIAM: Philadelphia, PA, USA, 1998. <https://doi.org/10.1137/1.9780898719628>.
34. Eskom Holdings SOC Ltd. Load-Shedding Stage Records, 2015–2024. Available online: <https://github.com/beyarkay/eskom-calendar> (accessed on 1 May 2026).
35. Baik, J.; Ben Arous, G.; Pécché, S. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Ann. Probab.* **2005**, *33*, 1643–1697.
36. Marčenko, V.A.; Pastur, L.A. Distribution of eigenvalues for some sets of random matrices. *Math. USSR-Sb.* **1967**, *1*, 457–483. <https://doi.org/10.1070/SM1967v001n04ABEH001994>.

37. Došilović, F.K.; Brčić, M.; Hlupic, N. Explainable artificial intelligence: A survey. In Proceedings of the 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO 2018) Opatija, Croatia, 21–25 May 2018; pp. 210–215. <https://doi.org/10.23919/MIPRO.2018.8400040>.
38. Yu, Y.; Wang, T.; Samworth, R.J. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika* **2015**, *102*, 315–323.
39. Hill, B.M. A simple general approach to inference about the tail of a distribution. *Ann. Stat.* **1975**, *3*, 1163–1174. <https://doi.org/10.1214/aos/1176343247>.
40. Pickands, J. Statistical inference using extreme order statistics. *Ann. Stat.* **1975**, *3*, 119–131. <https://doi.org/10.1214/aos/1176343003>.
41. Ollivier, Y. Ricci curvature of Markov chains on metric spaces. *J. Funct. Anal.* **2009**, *256*, 810–864. <https://doi.org/10.1016/j.jfa.2008.11.001>.
42. Zhu, Y.; Segarra, S. Hypergraphs with edge-dependent vertex weights: p -Laplacians and spectral clustering. *Front. Big Data* **2023**, *5*, 1020173. <https://doi.org/10.3389/fdata.2023.1020173>.
43. Vandenberghe, L.; Boyd, S. Semidefinite programming. *SIAM Rev.* **1996**, *38*, 49–95. <https://doi.org/10.1137/1038003>.
44. Moroke, N.D. Metabolic Saliency as KL-Divergence Estimator: Information-Geometric Attribution of Systemic Stress in JSE Equity Network. *Entropy* **2026**, *28*, 559. <https://doi.org/10.3390/e28050559>.
45. Billingsley, P. *Convergence of Probability Measures*, 2nd ed.; Wiley: New York, NY, USA, 1999. <https://doi.org/10.1002/9780470316962>.
46. Dvoretzky, A.; Kiefer, J.; Wolfowitz, J. Asymptotic minimax character of the sample distribution function. *Ann. Math. Stat.* **1956**, *27*, 642–669. <https://doi.org/10.1214/aoms/1177728174>.
47. Künsch, H.R. The jackknife and the bootstrap for general stationary observations. *Ann. Stat.* **1989**, *17*, 1217–1241. <https://doi.org/10.1214/aos/1176347265>.
48. Resnick, S.I. *Heavy-Tail Phenomena*; Springer: New York, NY, USA, 2007. <https://doi.org/10.1007/978-0-387-45024-7>.
49. Krause, A.; Guestrin, C. Near-optimal nonmyopic value of information in graphical models. In *Proc. UAI 2008*; AUAI Press: Corvallis, OR, USA, 2008; pp. 324–331. Available online: <https://las.inf.ethz.ch/files/krause05near.pdf> (accessed on 1 May 2026).
50. Villani, C. *Optimal Transport: Old and New*; Springer: Berlin, Germany, 2009. <https://doi.org/10.1007/978-3-540-71050-9>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.