

Article

Eskom-Induced Metabolic Arrest in JSE Financial Hypergraphs: A Physics-Informed Entropy Protocol for Systemic Risk Governance

Ntebogang Dinah Moroke 

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

Abstract

Physical infrastructure failure induces topological phase transitions in financial networks that existing systemic risk models cannot detect, attribute, or govern. This study introduces CASCADent (Cascading Systemic-Entropy-Detecting Network), a physics-informed entropy protocol for infrastructure-coupled financial hypergraphs. The model integrates three components: (i) a Landauer-motivated entropy threshold ($S^* = 2.852$ nats) that detects imminent metabolic arrest before it manifests as market stress; (ii) Gradient-Boosted Integrated Gradients (GB-IG) metabolic centrality that attributes systemic collapse to the specific apex nodes driving it; and (iii) a Lyapunov-constrained Hamilton–Jacobi–Bellman protocol that governs macroprudential intervention with deterministic stability tendency under a linear control assumption (Theorem 1); stochastic extensions remain future work. The key variables are daily equity returns and hypergraph topology for 19 JSE Top40 securities, integrated with Eskom Energy Availability Factor (EAF) telemetry and CBOE VIX ($T = 2838$ trading days, January 2015–April 2026). CASCADent achieves $F1 = 0.643$ at a 96 h early-warning lead time, outperforming all advance-warning baselines by 23.6 percentage points, with a false alarm rate of zero across 553 out-of-distribution days. Twelve apex nodes concentrate 85.5% of total system entropy, providing a governance target for macroprudential capital buffer design in infrastructure-dependent emerging economies.

Keywords: systemic risk; financial hypergraph; Von Neumann entropy; physics-informed machine learning; macroprudential governance; infrastructure risk; SDG 8; SDG 9; SDG 10; SDG 17; Eskom load-shedding; JSE financial networks; metabolic arrest; optimal control

1. Introduction

Physical infrastructure failure has emerged as a first-order systemic risk driver in emerging economies. The JSE–Eskom relationship provides a decade-long natural experiment demonstrating three empirical regularities: (i) network density rises while structural diversity simultaneously collapses under load-shedding stress, the densification paradox; (ii) this collapse follows a thermodynamically predictable entropy trajectory; and (iii) the collapse concentrates in a small number of identifiable apex nodes. Existing systemic risk frameworks treat infrastructure failure as an exogenous shock rather than a structural feature of the financial network, leaving these regularities undetected and ungoverned. More critically, the literature frames systemic risk monitoring as a classification problem: detect a crisis and issue a binary alert. This framing is operationally incomplete. A regulator who receives a crisis alert does not automatically know which institutions are driving the collapse, whether the signal will persist, or what specific intervention to execute at what cost. Without attribution and a governance protocol, an alert is a confirmation of failure,



Academic Editor: Anatoliy Swishchuk

Received: 29 May 2026

Revised: 27 June 2026

Accepted: 2 July 2026

Published: 10 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.

not a pre-emptive tool. CASCADent addresses this operational gap through a unified three-stage pipeline: (i) detect metabolic arrest before it manifests as observable market stress; (ii) attribute the collapse to the specific apex nodes driving it; and (iii) govern the intervention through a cost-optimised protocol with mathematically established stability properties. The contribution of this paper is not a marginal improvement in classification accuracy but is among the first empirically validated pipelines from physical infrastructure telemetry to macroprudential capital buffer design in an emerging market. The paper is organised as follows. Section 2 reviews the literature. Section 3 presents the CASCADent architecture. Section 4 introduces GB-IG attribution. Section 5 reports experimental results. Section 6 presents the HJB governance protocol. Section 7 discusses limitations and generalisability. Section 8 concludes the paper.

1.1. Conceptual Framework

The following two subsections establish the theoretical and architectural foundations of CASCADent. Section 1.1.1 formalises metabolic arrest as a topological phase transition. Section 1.1.2 identifies the two deficiencies in existing architectures that CASCADent addresses.

1.1.1. The Physics Problem: Metabolic Arrest as a Topological Phase Transition

The stability of modern financial systems is increasingly intertwined with the resilience of critical physical infrastructure, including electricity grids, telecommunications networks, and settlement platforms. Infrastructure failure does not merely slow financial markets, it induces a structural phase transition in which topology reorganises catastrophically. The Johannesburg Stock Exchange (JSE), coupled to South Africa's Eskom electricity grid, provides a decade-long natural laboratory for observing this phenomenon. Eskom load-shedding stages are fully exogenous to financial microstructure, offering uncontaminated variation for isolating the causal effect of physical constraints on network topology. Traditional systemic risk frameworks model contagion through bilateral financial exposures (Acemoglu et al. 2015) and cannot capture the higher-order dependencies that characterise real-world market coordination.

The signature of metabolic arrest is the densification paradox: under severe infrastructure stress, network density ρ rises while Von Neumann structural diversity H simultaneously collapses. Similar to jamming in granular physics (Stauffer and Aharony 1994) and trophic collapse in ecological networks (May 1972), this counterintuitive inversion is a topological phase shift that is caused by exogenous infrastructure deterioration rather than endogenous financial dynamics. The multilayer coupling between physical infrastructure and financial network topology, formalised in the foundational framework of (Kivelä et al. 2014) and reviewed comprehensively by (Aleta et al. 2026), means that infrastructure shocks propagate as structural phase transitions rather than smooth performance degradations. South Africa's 2015–2025 electricity crisis provides a natural laboratory for financial market microstructure, offering uncontaminated variation for isolating the effect of physical constraints on network topology.

Financial transactions are information-processing events. Each one requires a minimum energy expenditure of $k_B T \ln 2$ joules per bit erased, a bound established by Landauer's irreversibility principle (Landauer 1961). When physical infrastructure degrades below the energy level needed to sustain k -adic coordination among three or more institutions, the network's capacity to coordinate collapses. The critical entropy floor at which this collapse becomes self-reinforcing takes the following form:

$$S^* = k_B \ln 2 + k_B \gamma + \frac{k_B \ln k}{4} \quad (1)$$

where $\gamma \approx 0.5772$ is the Euler–Mascheroni constant (Havil 2003). For a k -adic hyper-edge whose member institutions are spatially distributed across the settlement network, maintaining phase-lock requires coordination overhead scaling as the harmonic sum $H_k = \sum_{i=1}^k 1/i \rightarrow \ln k + \gamma$ (Havil 2003), yielding the spatial coherence contribution $k_B(\ln k + \gamma)$; the term $k_B \ln k/4$ captures Byzantine consensus overhead. Equation (1) provides a theoretical lower bound that motivates anchoring a governance signal to Von Neumann entropy rather than to volatility or correlation measures. Equation (1) is motivated by Landauer’s irreversibility principle and provides a thermodynamically consistent lower bound for the entropy floor; it is not a strict derivation from first principles but, rather, a motivating analogy that guides the functional form and calibration direction of S^* . The analogy is consistent with established practice in physics-informed machine learning applied to complex systems (Kissas et al. 2020; Raissi et al. 2019), wherein physical laws motivate architectural choices without requiring strict derivation in the target domain. The empirical calibration of $S^* = 2.852$ nats is determined from the observed bimodal distribution of network entropy in the JSE dataset, as described in Section 3; the Landauer equation supplies the functional form and motivates why Von Neumann entropy is the appropriate order parameter rather than volatility or correlation.

CASCADent (Cascading Systemic-Entropy-Detecting Network) addresses this gap through a unified framework that integrates thermodynamic principles with hypergraph theory. The core innovation is the Landauer floor: a physically grounded entropy threshold S^* , derived from Landauer’s irreversibility principle (Landauer 1961) and calibrated to the specific hypergraph topology of the JSE, below which metabolic arrest becomes self-reinforcing. A hypergraph Laplacian (Zhou et al. 2006) coupled with Von Neumann entropy (Minello et al. 2019) quantifies structural degradation of the financial network in real time, providing an order parameter that moves beyond traditional volatility and correlation metrics. An algorithmic early warning system based on Gradient-Boosted Integrated Gradients (GB-IG) metabolic centrality then identifies the specific apex nodes driving the collapse, enabling targeted rather than system-wide intervention. The governance mechanism is formalised through a Hamilton–Jacobi–Bellman protocol (Fleming and Soner 2006) stabilised by Lyapunov functions (LaSalle and Lefschetz 1961), ensuring that corrective interventions maintain system stability during critical-infrastructure-induced metabolic disruptions. This combination constitutes, to the best of our knowledge, one of the first physics-informed hypergraph frameworks that simultaneously detects, attributes, and governs infrastructure-driven systemic risk in an emerging financial market. Throughout this paper, a distinction is maintained between two types of contribution: (i) *formal analytical results* (Propositions and Theorems), which are mathematically proved within the stated framework; and (ii) *empirical findings*, which are validated on the JSE dataset and are explicitly not claimed to generalise to other markets without cross-market replication. Claims about generalisation are framed as conjectures and directions for future work.

1.1.2. Existing Architectures, Their Failure Mode, and CASCADent

Graph neural network (GNN) applications to systemic risk monitoring have developed across three methodological generations (Battiston et al. 2016, 2012; Benson et al. 2016; Gai and Kapadia 2010; Hamilton et al. 2017; Veličković et al. 2018). Despite architectural advances from static bilateral models through temporal GNNs to hypergraph methods, all three generations share an identical failure mode: each architecture learns crisis signatures empirically, encoding volatility patterns, correlation structures, and contagion sequences observed during training. (Kamiński et al. 2024) demonstrated that even modularity-based hypergraph methods require external structural priors to avoid overfitting regime-specific community organisation. The black-box opacity of GNNs applied to complex

network classification tasks, documented by (Cvetković and Mitrović Dankulov 2026), compounds this failure mode: architectures that cannot explain their own decisions cannot be trusted to generalise across regime changes. CASCADent constitutes a fourth generation, distinguished by thermodynamic priors that are physically invariant across regime changes rather than learned from regime-specific training data.

This creates systematic vulnerability under a distribution shift. Architectures calibrated to one infrastructure regime are not guaranteed to transfer to another, because the statistical patterns they encode are regime-specific rather than physically invariant. Two specific deficiencies motivate CASCADent. First, existing architectures treat the crisis threshold as a latent pattern to discover from training data rather than a thermodynamically grounded constant to enforce through design. Second, classification-based architectures provide alerts without intervention protocols.

CASCADent addresses both deficiencies through three architectural innovations, each described in Section 3 and quantified in Section 5.

Innovation 1: Conservative entropy gate. A decision floor anchored at an empirically calibrated entropy operating point S^* prevents the classifier from issuing low-alert predictions when Von Neumann entropy approaches the stability threshold, regardless of what the data-driven component indicates.

Innovation 2: Hardware-in-the-loop inhibitory gate. The inhibitory gate $\gamma_L(\mathcal{I}_t)$ maps real-time Eskom Energy Availability Factor (EAF) telemetry directly to feature generation, suppressing messages from hyperedges that are operationally stressed by infrastructure failure, grounded in Landauer’s principle (Landauer 1961).

Innovation 3: Metabolic time-warping kernel. The kernel $\mathcal{K}(\tau_t)$ replaces calendar time with effective metabolic time $\tau_t = \int_0^t \nu(s) ds$, where $\nu(s)$ is the grid’s Energy Availability Factor (EAF), maintaining phase-lock between temporal features and the market’s physical coordination rate (Brunton and Kutz 2022).

1.1.3. Scope and Boundaries of This Study

This study is designed as an emerging-market proof-of-concept, not a general-purpose classifier contest. The framework is evaluated on a single infrastructure-coupled financial market (JSE–Eskom) over one decade. Three types of claims are explicitly distinguished:

- (1) **Formal analytical results** (Proposition 1, Empirical Finding 1, Empirical Finding 2 and Theorem 1) are mathematically proved within the stated framework and hold regardless of the empirical data.
- (2) **Empirical findings** (Section 5) are validated on the JSE dataset and are explicitly not claimed to generalise to other markets without cross-market replication.
- (3) **Generalisation conjectures** (Section 7.6) are presented as directions for future work, not established results.

The term *metabolic arrest* is used throughout as a conceptual label for the entropy collapse observed in the JSE hypergraph under infrastructure stress. The densification paradox (simultaneous rise in network density and fall in entropy) was a motivating observation but is not statistically robust above a permutation null (Section 5.2). The empirical signature of metabolic arrest in this study is the statistically robust compression of structural diversity and the persistence of hyperedge structure (both reported in Section 5), not the densification paradox, per se. The term is retained for conceptual clarity and consistency with the thermodynamic analogy but should not be interpreted as a claim that the densification paradox is empirically validated.

The following claims are explicitly outside the scope of this paper: that $S^* = 2.852$ nats is a universal thermodynamic constant; that CASCADent achieves the highest possible F1 score among all conceivable models; that simple infrastructure-threshold classifiers (e.g.,

alert when Eskom stage ≥ 4) cannot achieve competitive contemporaneous F1 scores; they can, but they provide zero advance warning, no attribution, and no governance protocol; and that the HJB protocol has been validated in a live regulatory environment. Researchers interested in pure classification benchmarks are referred to the open-source code and dataset archived at Zenodo (Moroke 2026) for comparative studies beyond the scope of this paper.

2. Related Work and Literature Synthesis

2.1. Generation 1: Static Network Contagion Models

Early systemic risk work modelled contagion through balance-sheet exposure channels on static weighted graphs. The canonical framework treats interbank liabilities as a weighted directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{W})$, where the entry $w_{ij} \geq 0$ represents institution i 's exposure to institution j . (Allen and Gale 2000) showed that highly connected networks are simultaneously robust to small shocks and are fragile under large ones, a duality that continues to shape the field. The clearing vector $\mathbf{p}^*$ of (Eisenberg and Noe 2001) is the fixed point of

$$p_i^* = \min\left(\bar{p}_i, e_i + \sum_j \frac{w_{ji}}{\bar{p}_j} p_j^*\right), \quad i = 1, \dots, N, \quad (2)$$

where $\bar{p}_i$ is the total nominal liability and e_i is the external asset value of institution i . (Gai and Kapadia 2010) demonstrated that the same density ρ that buffers idiosyncratic shocks amplifies systemic ones, establishing that contagion probability is non-monotone in connectivity. DebtRank (Battiston et al. 2012) quantifies systemic importance as the fraction of total economic value that would be lost following distress at node i :

$$R_i = \frac{1}{E_{\text{tot}}} \sum_{j \neq i} v_j h_{ji}, \quad (3)$$

where v_j is the economic value of institution j and h_{ji} measures the impact propagation from i to j . (Battiston et al. 2016) established that complexity theory requires interdisciplinary network analysis extending beyond standard regulatory models, while (Cont et al. 2013) localised default cascades within high-connectivity nodes, providing the empirical basis for the apex node concept formalised in Section 4.

2.2. Generation 2: Temporal Graph Neural Networks

Graph convolution (Kipf and Welling 2017) introduced the spectral message-passing update

$$\mathbf{H}^{(l+1)} = \sigma\left(\tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2} \mathbf{H}^{(l)} \mathbf{W}^{(l)}\right), \quad (4)$$

where $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ is the adjacency matrix with self-loops, $\tilde{\mathbf{D}}$ is the corresponding degree matrix, and $\mathbf{W}^{(l)}$ is a trainable weight matrix at layer l . Graph attention (Veličković et al. 2018) extended this with learnable attention coefficients α_{ij} , enabling GNN-based prediction of systemic risk in financial systems using deep graph learning (Balmaseda et al. 2023). Temporal architectures (Cui et al. 2023) extended detection lead times by 31%, while continuous-time frameworks (Rossi et al. 2020) enabled fine-grained surveillance of evolving network topology. GraphSAGE (Hamilton et al. 2017) variants exhibit significant F1 degradation during out-of-distribution stress periods (Fan et al. 2024), a vulnerability whose structural cause is overfitting to regime-specific correlations rather than learning invariant physical laws. This motivates the thermodynamic prior of CASCADent. Continual learning (Febrinanto et al. 2023) reduces catastrophic forgetting but remains limited by unconstrained hypothesis spaces that admit physically impossible network states. (Cvetković and Mitrović Dankulov 2026) demonstrated that standard GNN attribution methods cannot

resolve explanations at the hyperedge level, directly motivating the GB-IG framework of Section 4.

2.3. Generation 3: Hypergraph and Higher-Order Methods

A hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ generalises a graph by allowing each hyperedge $e \in \mathcal{E}$ to connect $k \geq 2$ nodes simultaneously. The associated Laplacian is

$$\mathbf{L}_{\mathcal{H}} = \mathbf{D}_v^{-1/2} \mathbf{B} \mathbf{W} \mathbf{D}_e^{-1} \mathbf{B}^{\top} \mathbf{D}_v^{-1/2}, \quad (5)$$

where $\mathbf{B} \in \{0, 1\}^{N \times |\mathcal{E}|}$ is the incidence matrix, $\mathbf{W}$ is a diagonal edge-weight matrix, and $\mathbf{D}_v$, $\mathbf{D}_e$ are the node and edge degree matrices, respectively. (Benson et al. 2016) proved that k -adic interactions generate contagion dynamics qualitatively distinct from pairwise approximations, establishing the theoretical necessity of this representation. (Kamiński et al. 2024) demonstrated that modularity-based community detection in hypergraphs reveals organisational structure invisible to pairwise algorithms. (Majhi et al. 2022) confirmed that k -adic interactions produce qualitatively distinct propagation dynamics that pairwise GNNs cannot replicate. (David et al. 2024) showed that functional hypergraphs of stock markets recover Von Neumann entropy and Forman–Ricci curvature signals invisible to pairwise correlation networks, directly anticipating the thermodynamic monitoring approach of CASCADent. (Zhang et al. 2023) proved that higher-order interactions shape collective dynamics differently in hypergraphs and simplicial complexes, with implications for contagion modelling under infrastructure stress. State-of-the-art baselines HTHGN (Liu et al. 2025) and industry-driven hypergraph attention (Han et al. 2023) exhibit F1 degradation from 0.787 to 0.621 during recovery regimes, confirming that higher-order architectures fail without thermodynamic anchors.

CASCADent constitutes a fourth generation of systemic risk model, distinguished from Generations 1–3 by thermodynamic priors that are physically invariant across regime changes rather than learned from regime-specific training data. This generation is evaluated empirically in Sections 3–6.

2.4. Physics-Informed Neural Networks

The PINN paradigm of (Raissi et al. 2019) encodes physical laws as residual constraints in the loss function:

$$\mathcal{L}_{\text{PINN}} = \mathcal{L}_{\text{data}} + \lambda \mathcal{L}_{\text{physics}}, \quad (6)$$

where $\mathcal{L}_{\text{physics}}$ penalises violations of the governing equation and $\lambda > 0$ balances the two objectives. This reduces effective hypothesis dimensionality more efficiently than additional training data alone. (Kissas et al. 2020) demonstrated that physical conservation laws improve out-of-distribution accuracy in data-sparse regimes, directly relevant to the OOD evaluation of Section 5. (Cai et al. 2021) showed that PINNs substitute physical priors effectively when governing equations are known but labels are scarce. CASCADent applies a thermodynamic operating point derived from Landauer’s principle (Landauer 1961) and the Euler–Mascheroni constant (Havil 2003) as an architectural prior. Unlike hybrid adaptive PINNs (Luo et al. 2025) that achieve soft penalty sampling, CASCADent applies a conservative entropy gate acting as a floor on predicted crisis probability rather than a purely learned threshold.

2.5. Entropy, Infrastructure Constraints, and Cyber-Physical Finance

(Bianconi 2009) formalised the Von Neumann entropy of graph Laplacians as a physically meaningful observable of network coordination capacity. For a normalised Laplacian $\hat{\mathbf{L}} = \mathbf{L}/\text{tr}(\mathbf{L})$ with eigenvalues $\hat{\lambda}_1, \dots, \hat{\lambda}_N$, the entropy is

$$S = -\text{tr}(\hat{\mathbf{L}} \log \hat{\mathbf{L}}) = -\sum_{i=1}^N \hat{\lambda}_i \log \hat{\lambda}_i. \quad (7)$$

(Minello et al. 2019) proved the additive decomposition of S into node contributions (Proposition 1), providing the spectral foundation for metabolic centrality. (Ye et al. 2018) demonstrated that entropy component analysis surfaces structural roles systematically missed by degree-based centrality. The critical threshold phenomenon appears across independent domains: EEG entropy collapse before hypoxic brain injury (Hofmeijer and van Putten 2016); structural diversity declining in financial networks during market stress (Billio et al. 2016); and critical slowing following mean-field universality class scaling (May 1972).

Infrastructure coupling.

The Eskom crisis, the 2021 Texas winter storm (Busby et al. 2021), and European gas disruptions established physical infrastructure as a systemic financial risk orthogonal to Basel III frameworks (Battiston et al. 2017). The multilayer network framework of (Aleta et al. 2026) argues that interconnected systems require multilayer representations to capture their full structural dynamics, grounding the inhibitory gate design of Section 3. (Busby et al. 2021) documented an early empirical instance of density–diversity decoupling during infrastructure failure. (Lee 2008) provides the cyber-physical co-design framework for the inhibitory gate.

2.6. Optimal Control and HJB Methods in Macroprudential Policy

Stochastic optimal control via HJB equations provides the rigorous foundation for value iteration in regulatory optimisation (Fleming and Soner 2006; Merton 1971). For a controlled diffusion $dS = \alpha(S, u) dt + \sigma dW_t$, the value function $V(S)$ satisfies

$$\delta V(S) = \min_{u \in \mathcal{U}} \left[c(u) + \alpha(S, u) V'(S) + \frac{1}{2} \sigma^2 V''(S) \right], \quad (8)$$

where $\delta > 0$ is the discount rate and $c(u)$ is the intervention cost (Fleming and Soner 2006). Entropy-based indicators provide superior HJB inputs over conventional financial metrics (Billio et al. 2016). CASCADent diverges from existing HJB macroprudential models (Brunnermeier and Sannikov 2014; Fleming and Soner 2006) by employing a GNN as the state-transition function, and from communication network controls (Amos and Kolter 2017) by operating on financial hypergraphs with thermodynamic priors. (Brunton and Kutz 2022) established the data-driven framework for embedding physical laws in machine learning architectures, providing the methodological precedent for CASCADent's physics-first design.

The full 34-study synthesis table is provided in Appendix A.

3. CASCADent Architecture

Section overview. This section describes the complete CASCADent methodology: hypergraph construction (Section 3.2), entropy threshold calibration (Section 3.3), the hybrid architecture (Section 3.4), and the inhibitory gate (Section 3.5). It is followed by Section 4, which introduces the GB-IG attribution component. Together, these two sections constitute the complete technical framework for detecting, attributing, and governing infrastructure-induced systemic risk.

3.1. Overview and Design Philosophy

Figure 1 illustrates the three-stage CASCADent pipeline: Stage I (Detect) constructs the hypergraph and triggers the entropy gate; Stage II (Attribute) applies GB-IG centrality; Stage III (Govern) solves the HJB control problem.




Figure 1. Internal architecture of CASCADent (Cascading Systemic-Entropy-Detecting Network). The pipeline is organised into three explicit stages: **Stage I** (Detect) constructs the financial hypergraph, computes Von Neumann entropy, and triggers the entropy gate when $S(t) < S^*$; **Stage II** (Attribute) applies GB-IG metabolic centrality to identify the apex nodes concentrating systemic fragility; **Stage III** (Govern) solves the Hamilton–Jacobi–Bellman control problem and validates Lyapunov stability of the intervention. Filled circles denote securities (nodes): gold = Tier 1 apex, teal = Tier 2, grey = non-apex; connecting lines indicate active hyperedges.

Figure 2 shows the JSE hypergraph trajectory in the entropy-derivative phase space $(S(t), \dot{S}(t))$, illustrating entropy collapse under Eskom Stage ≥ 4 stress and asymptotic recovery under the HJB governance protocol.




Figure 2. Evolution of the JSE financial hypergraph state in the entropy-derivative phase space $(S(t), \dot{S}(t))$, demonstrating activation of the entropy gate and asymptotic convergence to stability under optimal control. The shaded critical region ($S < S^* = 2.852$ nats) demarcates the metabolic arrest zone. The solid trajectory traces entropy collapse under Eskom Stage 4+ stress; the dashed trajectory shows recovery under the HJB governance protocol (Theorem 1). The vertical dashed line marks $S^* = 2.852$ nats; the filled square (■) marks the shock onset point.

CASCADent (Cascading Systemic-Entropy-Detecting Network) is a physics-inspired hypergraph classifier designed for binary crisis detection in infrastructure-conditioned financial markets. The architecture departs from the three-generation trajectory reviewed in Section 2 by anchoring an empirically calibrated entropy operating point as a conservative decision gate rather than learning the crisis threshold from data alone. Figure 3 provides a complete schematic of the six-component pipeline.




Figure 3. CASCADent architecture: six-panel pipeline schematic. (A) (Hypergraph construction): 19 JSE Top40 securities as nodes; Tier 1 apex nodes (gold), Tier 2 (purple), non-apex (grey); hyperedges formed as maximal cliques where $|\rho_{ij}| \geq 0.3$. (B) (Von Neumann entropy): Normalised Laplacian eigenspectrum and entropy time series $S(t)$ with empirical threshold $S^* = 2.852$ nats (dashed line); crisis shading (coral) below S^* . (C) (Thermodynamic gates): Entropy gate $\hat{p}_t = \max(\hat{p}_{\text{clt}}, 1[S(t) < S^*] \cdot 0.75)$, Eskom inhibitory gate $\gamma_L(\mathcal{I}_{t-5})$, and metabolic time-warping kernel $\tau_t = \int_0^t \nu(s) ds$. Arrows indicate the direction of information flow between pipeline components. (D) (Feature vector): 20-dimensional physics-informed feature vector $\mathbf{x}_t \in \mathbb{R}^{20}$ comprising entropy (7), density (5), Eskom (2), VIX (3), and volatility (3) blocks. (E) (GBT decision layer): Gradient-boosted ensemble benchmark table; CASCADent row highlighted (F1 = 0.643 at 96 h lead time; FAR = 0.000; +23.6 pp over best advance-warning baseline). $T = 2838$ trading days; crisis rate 7.8%. (F) (HJB governance): Three-regime decision-support protocol (Normal/Alert/Crisis) grounded in the linear stability result $\dot{\mathcal{L}} = -2\kappa(S - S^*)^2 \leq 0$. Green = Normal regime (no intervention); amber = Alert regime (pre-emptive buffers); red = Crisis regime (full intervention).

The pipeline operates in two layers. The physics-informed feature layer (panels A–C) encodes four physical constraints: the Landauer entropy floor, Eskom infrastructure telemetry, the harmonic-series coordination overhead, and the additive entropy decomposition of Proposition 1. The output is the 20-dimensional feature vector $\mathbf{x}_t$ (panel D). The decision layer (panel E) is a gradient-boosted tree ensemble whose interpretable feature importances support GB-IG attribution (Section 4) and regulatory transparency. The *governance layer* (panel F) translates classifier output into a regime-specific intervention protocol. The three-layer separation is deliberate: it isolates the contribution of physics-informed features from architectural complexity, enabling clean ablation attribution.

Three design decisions distinguish CASCADent from existing approaches.

First, the crisis operating point S^* is motivated by analogy with Landauer’s irreversibility principle (Landauer 1961) and calibrated empirically via Youden index optimisation on the training set, yielding $S^* = 2.852$ nats for the 19-node JSE network (bootstrap 95% CI: [2.797, 2.855]). This value is network-size-dependent and not transferable across markets without recalibration. Second, a hardware-aware inhibitory gate maps lagged Eskom load-shedding stage to the decision gate, suppressing low-crisis-probability predictions under infrastructure stress. Third, a metabolic time-warping kernel replaces calendar time with

effective metabolic time τ_t , maintaining temporal alignment with the market's physical coordination rate.

The input layer receives three data streams: (i) a rolling 30-day log-return matrix $\mathbf{R} \in \mathbb{R}^{T \times N}$ for $N = 19$ JSE Top40 securities, winsorised at $\pm 3\sigma$ per ticker to remove data discontinuities from unadjusted corporate actions; (ii) the lagged Eskom national load-shedding stage $\mathcal{I}_{t-5}$, offset by five trading days to prevent label leakage; and (iii) the log-transformed CBOE VIX index $\log(\text{VIX}_t)$, validated as a complementary systemic stress signal with SHAP attribution scores of 0.092–0.093 across all regime classes (David et al. 2024).

3.2. Hypergraph Construction

From each 30-day rolling window $\mathbf{R}_t$, the hypergraph $\mathcal{H}_t = (\mathcal{V}, \mathcal{E}_t)$ is constructed as follows. The node set $\mathcal{V}$ contains the $N = 19$ securities as permanent vertices. The hyperedge set $\mathcal{E}_t$ is derived from the Pearson correlation matrix $\mathbf{C}_t = [\rho_{ij}]$ by thresholding at $|\rho_{ij}| \geq 0.3$: securities whose pairwise correlations exceed this threshold jointly participate in a hyperedge. Formally, a hyperedge $e \in \mathcal{E}_t$ is a maximal clique in the thresholded adjacency graph. The incidence matrix $\mathbf{H} \in \{0, 1\}^{N \times |\mathcal{E}_t|}$ encodes hyperedge membership (Kamiński et al. 2024).

Limitation of the maximal-clique construction. The maximal-clique approach requires that *every pair* of securities within a hyperedge exceeds the correlation threshold $|\rho_{ij}| \geq 0.3$. This is a conservative choice: it encodes strong bilateral dependency across all members, whereas more general hyperedge constructions (e.g., star-shaped or arbitrary set-covering hyperedges) do not require all pairwise correlations to be elevated. The clique construction is appropriate for JSE equity data because high bilateral correlation in liquid markets reliably indicates joint factor exposure rather than coincidental co-movement. However, it may miss higher-order dependencies that are not reflected in all pairwise correlations, for example, a set of securities that collectively respond to Eskom stage but whose bilateral correlations are individually moderate. A hyperedge community-detection approach such as that of Kamiński et al. (Kamiński et al. 2024) would relax this assumption and is a direction for future work.

Two quantities are computed from $\mathcal{H}_t$ at each time step t :

$$S(t) = - \sum_{i=1}^N \lambda_i(\mathbf{L}_t) \log \lambda_i(\mathbf{L}_t), \quad (9)$$

$$\rho(t) = \frac{\sum_{i < j} \mathbf{A}_{ij}}{N(N-1)}, \quad (10)$$

where $\mathbf{L}_t = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A}_t \mathbf{D}^{-1/2}$ is the normalised Laplacian of the thresholded adjacency matrix $\mathbf{A}_t$, and λ_i denotes its i -th eigenvalue after normalisation to a unit probability vector.¹ Equation (9) gives the Von Neumann entropy of the hypergraph, formalised by (Bianconi 2009) and extended to hypergraph Laplacians by (Minello et al. 2019). Equation (10) gives the fraction of node pairs connected above the threshold. A key property of this entropy measure is that it decomposes into additive node contributions, which is what makes GB-IG attribution possible. We state this formally because the decomposition underpins both the thermodynamic gate design and the apex node identification in Section 4.

Proposition 1 (Additive Entropy Decomposition). *Let $\mathbf{L}$ be the normalised Laplacian of the JSE hypergraph with eigenvalues $\lambda_1 \leq \dots \leq \lambda_N$. The Von Neumann entropy $S(\mathcal{H}) = -\text{tr}(\hat{\mathbf{L}} \log \hat{\mathbf{L}})$, where $\hat{\mathbf{L}} = \mathbf{L} / \text{tr}(\mathbf{L})$, decomposes additively into node contributions $S(\mathcal{H}) = \sum_{i=1}^N \mu_i(\mathcal{H})$, where $\mu_i = -\sum_{j=1}^N \hat{\lambda}_j u_{ji}^2 \log \hat{\lambda}_j$ and u_{ji} is the i -th component of the j -th eigenvector.*

Proof. This follows directly from spectral decomposition $\hat{\mathbf{L}} = \sum_j \hat{\lambda}_j \mathbf{u}_j \mathbf{u}_j^\top$ and linearity of the trace. The general result is given in (Minello et al. 2019). $\square$

Figure 4 presents the 19-node JSE correlation network and the Von Neumann entropy time series $S(t)$ over the full 2015–2026 sample, with Eskom Stage ≥ 4 crisis episodes shaded.




Figure 4. JSE financial hypergraph topology and Von Neumann entropy dynamics (January 2015–April 2026). (a): Correlation network for the 19-node JSE panel (recovery regime, 2025); gold nodes = Tier 1 apex nodes, teal = Tier 2, grey = non-apex; edge colour denotes sign (positive = coral, negative = violet); edge width proportional to $|\rho_{ij}|$. (b): Von Neumann entropy $S(t)$ time series; dashed line = calibrated threshold $S^* = 2.852$ nats; red shading = Eskom Stage ≥ 4 crisis episodes.

Figure 5 shows the Pearson correlation matrix for the 19-node panel and the Von Neumann entropy distribution stratified by market regime.




Figure 5. Hypergraph adjacency structure and entropy regime distribution. (a): Pearson correlation matrix for the 19-node JSE panel; gold lines highlight Tier 1 apex node rows and columns (PRX, NPN, IMP, SBK). (b): Von Neumann entropy distribution stratified by regime: baseline 2015–2018, early crisis 2019–2021, peak crisis 2022–2023, and recovery 2024–2026; vertical dashed line = $S^* = 2.852$ nats; shaded bands = interquartile range of entropy distribution per regime. Overlapping distributions reflect genuine clustering of observations near the entropy threshold and do not impair scientific interpretation.

We observe that the densification paradox is the empirical signature of metabolic arrest: $\rho(t)$ increases simultaneously with a collapse in $S(t)$. This pattern is observed in some rolling windows but is not a statistically robust signal above a permutation null

model (see Section 5 for the full surrogate test); its role here is motivational rather than definitional. The structural finding that is statistically robust is hyperedge persistence, presented in Section 5.

The value of the hypergraph representation lies not in the Laplacian entropy measure, per se. The weighted hypergraph and binary pairwise Laplacians are highly correlated ($r = 0.981$), and the difference in entropy AUC between representations is modest (0.477 vs. 0.434). Rather, its value lies in the multiway attribution that the hypergraph representation enables. GB-IG metabolic centrality (Section 4) requires entropy perturbation at the node level, which demands that hyperedges encode joint dependencies among three or more nodes simultaneously. A pairwise graph reduces every multiway relationship to a set of bilateral edges, destroying the joint structure that GB-IG attribution exploits. The hypergraph representation is, therefore, motivated by interpretability rather than predictive entropy alone.

Correlation threshold sensitivity.

The default hyperedge activation threshold $\rho = 0.30$ follows the established literature on correlation-filtered financial networks (David et al. 2024; Kamiński et al. 2024). A sensitivity analysis varying $\rho \in \{0.20, 0.30, 0.40, 0.50, 0.60\}$ is reported in Table 1.

Table 1. Sensitivity of CASCADent performance to correlation threshold ρ . F1 and FAR on in-distribution test set (2022–2023 crisis episodes). S^* crossing rate: fraction of labelled crisis days on which $S(t) < S^*$. Default $\rho = 0.30$ (bold) maximises F1. Note: computed on the 17-security panel (PRX.JO and OMU.JO excluded, not listed January 2015; $T = 2819$ trading days). The full CASCADent pipeline uses all 19 securities ($T = 2838$) and achieves $S^* = 2.852$ nats; the values in this table reflect entropy-gate calibration per ρ setting.

ρ	Mean Degree	Hyperedges	$\bar{S}$ (nats)	S^* Crossing (%)	F1	FAR
0.20	10.9	93	2.756	90.8	0.678	0.985
0.30	8.6	73	2.732	100.0	0.757	0.824
0.40	6.4	54	2.692	98.9	0.751	0.824
0.50	4.5	38	2.655	100.0	0.744	0.882
0.60	2.8	24	2.645	65.5	0.708	0.250

Rolling window sensitivity.

Table 2 reports CASCADent performance across five rolling window lengths ($w \in \{15, 20, 30, 45, 60\}$ days), holding $\rho = 0.30$ fixed.

Table 2. Sensitivity of CASCADent performance to rolling window length w (days), $\rho = 0.30$ fixed. Computed on the 17-security panel; default $w = 30$ days (bold) maximises F1 for the entropy gate component.

Window w (Days)	Mean Degree	$\bar{S}$ (Nats)	F1	FAR
15	9.4	2.745	0.673	0.891
20	9.0	2.740	0.746	0.968
30	8.6	2.732	0.757	0.976
45	8.2	2.723	0.731	0.985
60	8.0	2.717	0.719	0.998

3.3. Empirical Entropy Threshold and Gate

Drawing an analogy with Landauer’s irreversibility principle (Landauer 1961), financial transactions are information-processing events that require physical infrastructure to execute. This motivates the use of Von Neumann entropy as the governance signal. This is an analogy rather than a strict physical derivation; it motivates the form and direction of

the threshold but does not constitute a thermodynamic proof. This motivates a real-time decision threshold S^* below which the gate elevates the predicted crisis probability.

Calibration protocol.

The entropy threshold S^* is calibrated using the following four-step protocol applied to the training sample (January 2015–December 2022):

- Step 1. Compute the rolling Von Neumann entropy $S(t) = -\text{tr}(\hat{\mathbf{L}} \log \hat{\mathbf{L}})$ daily on the 30-day rolling hypergraph Laplacian for all T_{train} trading days.
- Step 2. Fit a two-component Gaussian mixture model (GMM) to the empirical distribution of $S(t)$, identifying a bimodal structure corresponding to normal-coordination and stress-regime entropy levels.
- Step 3. Set S^* at the inter-modal inflection point of the fitted GMM, yielding $S^* = 2.852$ nats for the 19-security JSE panel.
- Step 4. Validate: confirm that $S(t) < S^*$ is observed in 100% of the labelled crisis episodes in the held-out validation set (2023).

Three limitations of this calibration are noted. First, S^* is network-size-dependent and must be recalibrated for any panel of different size or composition. Second, the GMM assumes bimodality of the entropy distribution, which may not hold in markets with smoother stress transitions. Third, the calibration requires a minimum of five years of continuous data covering at least one full infrastructure stress cycle.

The threshold is determined empirically. Let $\mathbf{S} = \{S(t)\}_{t=1}^{T_{\text{train}}}$ be the entropy series over the training period and $y_t \in \{0, 1\}$ the crisis label. The operating point is

$$S^* = \arg \max_{\theta} [\text{TPR}(\theta) - \text{FPR}(\theta)], \quad (11)$$

where J is the Youden index. For the 19-node JSE hypergraph this yields $S^* = 2.852$ nats (bootstrap 95% CI: [2.797, 2.855]). The confidence interval overlaps the mean crisis-period entropy ($\bar{S}_{\text{crisis}} = 2.848$ nats), confirming that S^* should be interpreted as a practical operating point rather than a sharp natural boundary. Sensitivity of S^* and the network diversity index to rolling window (20–60 days) and correlation threshold (0.2–0.4) is reported in the robustness subsection of Section 5; both are stable across all configurations tested. Its contribution is as a conservative decision anchor: it consistently elevates sensitivity at the cost of some precision, a trade-off quantified by the ablation study in Section 5.

The Landauer derivation (Landauer 1961) and its harmonic-series asymptotic (Havil 2003) provide the theoretical motivation for anchoring the gate to an entropy quantity rather than a volatility or correlation measure, and suggest the form that S^* should take for networks of different sizes. For the 19-node JSE network, the empirically calibrated value is network-size-dependent and not a universal thermodynamic constant.

The thermodynamic gate applies the operating point as follows:

$$\hat{p}_t^{\text{crisis}} = \max(\hat{p}_t^{\text{clf}}, \mathbf{1}[S(t) < S^*] \cdot 0.75), \quad (12)$$

where $\hat{p}_t^{\text{clf}}$ is the raw classifier probability from the gradient-boosted backbone. Equation (12) acts as a conservative floor: when entropy falls below the operating point, the model is prevented from issuing a low-crisis-probability prediction, regardless of what the data-driven component indicates. The ablation study indicates that this gate is the single largest contributor to recall and to OOD generalisation (Section 5).

3.4. Hybrid Architecture: Physics-Informed Feature Layer and Decision Layer

CASCADent implements a deliberate two-layer hybrid architecture consistent with the physics-informed machine learning paradigm of (Raissi et al. 2019): physical laws and domain knowledge are embedded in the feature engineering layer, while a stable out-of-distribution classifier operates in the decision layer.

The *physics-informed feature layer* comprises hypergraph construction, Laplacian eigen-decomposition, rolling Von Neumann entropy computation, thermodynamic gate activation, inhibitory gate mapping, and metabolic time-warping. This layer encodes the four physical constraints that distinguish CASCADent from purely data-driven architectures: the Landauer entropy floor, Eskom infrastructure telemetry, harmonic-series coordination overhead, and the additive entropy decomposition of Proposition 1. The output is the 20-dimensional physics-grounded feature vector $\mathbf{x}_t$ defined in Equation (13).

The *decision layer* is a gradient-boosted tree ensemble (300 estimators, maximum depth 4, learning rate 0.05, subsample fraction 0.8). This choice is deliberate and principled rather than a concession to tractability. Gradient-boosted trees provide monotone feature importance scores that are directly interpretable in terms of the physical quantities in $\mathbf{x}_t$; this interpretability is essential for the GB-IG attribution of Section 4 and for regulatory credibility. Full GNN implementations with hypergraph message passing are a natural extension explored in Section 7; the present two-layer design isolates the contribution of the physics-informed features from architectural complexity, enabling cleaner ablation attribution.

The ensemble is trained with sample weights that upweight observations near the entropy boundary.

A 20-dimensional feature vector is constructed at each time step t from five groups of rolling metrics and exogenous stress signals. The entropy group comprises $[S_t, \bar{S}_t, \dot{S}_t, \sigma_t^S, \bar{S}_t, \mathbf{1}_{S < S^*}, \mathbf{1}_{S \approx S^*}]$ (7 features); the density group $[\rho_t, \dot{\rho}_t, \delta_t, \bar{\rho}_t^c, \dot{\rho}_t^c]$ (5 features); the Eskom group $[\mathcal{I}_{t-5}, \bar{\mathcal{I}}_t]$ (2 features); the VIX group $[\log \text{VIX}_t, \log \bar{\text{VIX}}_t, \log \dot{\text{VIX}}_t]$ (3 features); and the volatility group $[\sigma_t^{\text{cs}}, \sigma_t^{\text{ts}}, \text{DD}_t]$ (3 features), giving

$$\mathbf{x}_t = [\mathbf{e}_t \parallel \boldsymbol{\rho}_t \parallel \mathbf{i}_t \parallel \mathbf{v}_t \parallel \boldsymbol{\sigma}_t] \in \mathbb{R}^{20}, \quad (13)$$

where $\parallel$ denotes vector concatenation, $\mathbf{e}_t \in \mathbb{R}^7$ is the entropy block, $\boldsymbol{\rho}_t \in \mathbb{R}^5$ is the density block, $\mathbf{i}_t \in \mathbb{R}^2$ is the Eskom block, $\mathbf{v}_t \in \mathbb{R}^3$ is the VIX block, and $\boldsymbol{\sigma}_t \in \mathbb{R}^3$ is the volatility block. Dot notation denotes 10-day rolling slopes; δ_t is the binary densification paradox indicator ($\dot{\rho} > 0$ and $\dot{S} < 0$ simultaneously); DD_t is mean drawdown over the rolling window. The gate indices $\mathbf{1}_{S < S^*}$ and $\mathbf{1}_{S \approx S^*}$ (within 0.15 nats of S^*) operationalise the conservative entropy floor.

The backbone ensemble is trained with sample weights:

$$w_t = \left(1 + 3 \exp\left[-\frac{(S_t - S^*)^2}{0.04}\right]\right) \cdot 3 \mathbb{I}[y_t = 1] + \mathbb{I}[y_t = 0], \quad (14)$$

where $\mathbb{I}[\cdot]$ is the indicator function, so that both crisis examples ($y_t = 1$) and boundary examples near S^* receive elevated influence during training, motivated by the finding of (Cvetković and Mitrović Dankulov 2026) that GNN-based classifiers systematically underweight structurally critical observations near the classification boundary.

3.5. Inhibitory Gate and Metabolic Time Kernel

The inhibitory gate $\gamma_L(\mathcal{I}_t)$ maps the lagged Eskom load-shedding stage to a suppression coefficient applied to the message-passing aggregation. Under Stage 4 or higher ($\mathcal{I}_{t-5} \geq 4$), the gate coefficient is reduced proportionally to the stage level, reflecting Landauer's principle (Landauer 1961): higher infrastructure stress reduces the energy available

for k -adic coordination, effectively removing high-order hyperedges from the network's active topology. During the 2022–2023 crisis period, which constitutes 49.9% of the test set, the inhibitory gate was active on a substantial fraction of trading days, contributing to CASCADent's recall advantage over static baselines.

The metabolic time-warping kernel replaces calendar time t with effective metabolic time $\tau_t = \int_0^t v(s) ds$, where the metabolic velocity $v(s)$ is proportional to the grid's Energy Availability Factor. This rescaling ensures that temporal features computed over rolling windows of length w in calendar time correspond to windows of consistent informational content in metabolic time, maintaining phase-lock with the market's physical coordination rate (Rossi et al. 2020).

3.6. Data and Variables

The empirical analysis uses three data sources, each justified by the literature on infrastructure-coupled financial networks.

JSE equity returns. Daily closing prices for 19 JSE Top40 securities are used to construct the financial hypergraph. The Top40 constituents represent approximately 80% of JSE market capitalisation by value; using this near-census of apex institutions rather than a random sample is standard practice in systemic risk studies where contagion concentration is the object of interest (Battiston et al. 2012; Cont et al. 2013). Pairwise Pearson correlations of log-returns are computed on a 30-day rolling window to produce the hypergraph incidence matrix, following (David et al. 2024) and the broader correlation-based network literature (Kamiński et al. 2024). Sensitivity to window length (15, 20, 30, 45, 60 days) is reported in Table 2.

Eskom Energy Availability Factor (EAF). Daily load-shedding stage data (Stage 0–8) are obtained from the EskomSePush public archive (March 2014–May 2025), extended to April 2026 using SARB energy monitors. EAF is the primary inhibitory gate input because Eskom stage is fully exogenous to JSE microstructure, providing uncontaminated variation for causal identification of infrastructure effects on network topology (Battiston et al. 2017; Busby et al. 2021).

CBOE VIX. The daily VIX index is included as an international stress covariate to disentangle global risk-off episodes from Eskom-specific structural effects (Billio et al. 2016). Including VIX ensures that the entropy-gate signal reflects domestic infrastructure stress rather than global sentiment shocks.

The full dataset spans $T = 2838$ trading days (January 2015–April 2026) and is publicly archived at Zenodo (Moroke 2026) (DOI: 10.5281/zenodo.20008530). The training/validation/test split follows an 80/10/10 temporal partition: training (2015–2022), validation (2023), and out-of-distribution test (2024–2026).

Research design justification.

Three design choices require explicit justification. **(D1) Security selection:** The 19 JSE Top40 constituents are the highest-capitalisation securities with complete daily price records over the full 2015–2026 window and continuous listing on the Top40 index. They represent approximately 80% of JSE market capitalisation and constitute a near-census of the apex institutions in this market rather than a statistical sample. Expanding to the full Top40 introduces survivorship bias: securities delisted or added mid-sample have incomplete price histories that require forward-filling or window truncation, both of which introduce differential data quality across nodes and distort the Laplacian eigendecomposition that underpins the Von Neumann entropy computation. (Minello et al. 2019) demonstrate that Von Neumann entropy computed from the normalised graph Laplacian is sensitive to the structural configuration of the node set; artefactual heterogeneity introduced by

data gaps would distort the entropy estimate in ways indistinguishable from genuine infrastructure-driven structural changes, undermining the calibration of S^* (Kamiński et al. 2024). The 19-node panel avoids these artefacts entirely: all nodes have identical, uninterrupted price histories across the full 11-year window. The panel retains sectoral diversity across 8 of the 11 GICS sectors represented in the Top40, covering the dominant financial, mining, technology, consumer, and industrial sectors relevant to systemic risk governance. **(D2) Study period (2015–2026):** This window was selected to span the full Eskom load-shedding crisis, from its onset in Stage 1 events (2015) through the most severe Stage 6–8 episodes (2022–2023), capturing at least two complete infrastructure stress cycles necessary for reliable entropy calibration. **(D3) Rolling window (30 days):** The 30-day window (approximately one trading month) balances responsiveness to structural change against estimation noise in the Pearson correlation matrix. Sensitivity to this choice across five window lengths (15, 20, 30, 45, 60 days) is reported in Table 2, confirming that F1 is maximised at $w = 30$ and degrades monotonically with longer windows. Note: sensitivity analyses use a 17-security panel (PRX.JO listed September 2019; OMU.JO listed June 2018—both excluded for the 2015–2026 window), giving $T = 2819$ trading days. All primary CASCADent results use the full 19-security panel ($T = 2838$).

3.7. Experimental Setup and Implementation Details

The JSE hypergraph dataset comprises daily closing prices for 19 JSE Top40 securities spanning January 2015 to April 2026 ($T = 2838$ trading days), sourced from Yahoo Finance (Yahoo Inc., Sunnyvale, CA, USA). The 19-node panel represents the liquid core of the JSE Top40 with continuous OHLCV data across the full 2015–2026 window; expanding to the full Top40 introduces survivorship bias (securities delisted or added mid-sample) and missing-data hyperedges that distort Laplacian eigendecomposition. The 19-node panel retains sectoral diversity across 8 of the 11 GICS sectors represented in the Top40, mitigating concerns about sectoral concentration despite the reduced sample size. Log returns were winsorised at $\pm 3\sigma$ per ticker prior to all hypergraph computations; the three securities exhibiting extreme kurtosis ($\kappa > 1000$) due to unadjusted corporate action events in 2025 (SBK.JO, SLM.JO, VOD.JO) each required only two values to be winsorised. The CBOE VIX (Chicago Board Options Exchange, Chicago, IL, USA) was sourced from Yahoo Finance (Yahoo Inc., Sunnyvale, CA, USA) over the same period (range: 9.1–82.7). Eskom national load-shedding stages were compiled from the public EskomSePush historical archive (EskomSePush, Cape Town, South Africa), covering March 2014 to May 2025, and extrapolated at Stage 0 through April 2026 consistent with the publicly documented end of continuous load-shedding in March 2024 (Eskom Holdings SOC Ltd. 2025).

A stratified chronological split was applied to ensure that both crisis and non-crisis observations appear in training: training spans January 2015 to June 2022 ($n = 1861$, crisis rate 1.3%); the test set spans July 2022 to December 2023 ($n = 375$, crisis rate 49.9%, corresponding to the peak Stage 5–6 crisis period); and the out-of-distribution (OOD) evaluation set spans January 2024 to April 2026 ($n = 562$, crisis rate 1.6%), covering the grid recovery period. The crisis label is defined as Eskom national stage ≥ 4 , producing 220 crisis days (7.8%) over the full sample and 9 distinct crisis episodes of mean duration 24.4 trading days. The Eskom stage variable is lagged by five trading days before inclusion as a feature, eliminating any possibility of contemporaneous label leakage.

Data split justification. The 49.9% crisis rate in the test set reflects the documented reality of the 2022–2023 JSE period, during which Eskom operated at Stage 4+ for the majority of trading days. This is an empirical property of the crisis episode, not a sampling artefact. The OOD period (1.6% crisis rate) has a qualitatively different Eskom stage distribution, confirming that the test-to-OOD shift is genuine: the test set evaluates detection

performance during sustained high-stage crisis, while the OOD set evaluates false alarm behaviour during recovery. These two objectives are complementary.

Hyperparameter selection. All GBT hyperparameters (300 estimators, depth 4, learning rate 0.05, subsample 0.8) were selected by 5-fold time-series cross-validation on the training set (January 2015–June 2022), with the Youden index as the optimisation objective. The entropy gate floor (0.75) and the boundary weight multiplier (3 in Equation (14)) were cross-validated on the same folds. A complete hyperparameter search log is provided in the Zenodo repository alongside the analysis code.

Computational cost. The full pipeline—hypergraph construction, entropy computation, feature engineering, training, and evaluation—runs in approximately 4.2 min on a standard laptop (Apple M2, 16 GB RAM, Python 3.12). The dominant cost is Laplacian eigendecomposition over 2838 rolling windows (≈ 3.8 min). Runtime scales as $\mathcal{O}(T \cdot N^3)$ with the number of securities N , making the framework tractable for panels up to approximately $N = 50$ without parallelisation.

Using `scipy` 1.12, `scikit-learn` 1.4, `pandas` 2.2, and `networkx` 3.2, the entire pipeline hypergraph building, entropy computation, feature engineering, model training, and evaluation is implemented in Python 3.12, with all random seeds fixed at 42. The processed dataset and full analytic code are archived at Zenodo (Moroke 2026) (DOI: 10.5281/zenodo.20008530, Version 2).

4. Graph-Based Integrated Gradients and Metabolic Centrality

Section overview. While Section 3 provides the early-warning signal (a binary alert issued 96 h ahead of systemic stress), that alert alone is insufficient for governance. A regulator who receives an alert needs to know *which institutions* are driving the entropy collapse in order to apply targeted capital buffer adjustments rather than system-wide interventions. This section introduces Gradient-Boosted Integrated Gradients (GB-IG) metabolic centrality, which attributes the CASCADent alert to the specific apex nodes concentrating systemic entropy. The method moves from black-box prediction to interpretable attribution, directly supporting the HJB governance protocol in Section 6.

4.1. Motivation: From Black-Box Alert to Attributed Node

Crisis detection without attribution is operationally insufficient for macroprudential governance. A regulator receiving a binary alert must know *which* institutions are anchoring system entropy above S^* and *how much* each contributes to the network's coordination capacity. Standard degree-based centrality measures are inadequate for this purpose: they characterise pairwise connectivity but are insensitive to the higher-order, entropy-mediated coordination capacity that the CASCADent framework identifies as the binding constraint on financial stability. (Cvetković and Mitrović Dankulov 2026) demonstrated that GNNs applied to complex network classification systematically fail to surface the structural roles that matter most for network integrity, precisely because their attribution methods operate in the feature space of the model rather than the structural space of the network.

This restriction is overcome by the Graph-Based Integrated Gradients (GB-IG) technique, which calculates node importance as the anticipated change in Von Neumann entropy brought about by eliminating each node from the hypergraph. The approach draws on (Minello et al. 2019), who showed that Von Neumann entropy decomposes additively into node and edge contributions across network topologies. It extends the axiomatic Integrated Gradients framework of (Sundararajan et al. 2017) to the hypergraph setting by using entropy change, rather than gradient magnitude, as the attribution signal.

4.2. Formal Definition of Metabolic Centrality

Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be the JSE financial hypergraph with $N = 19$ nodes and let $S(\mathcal{H})$ denote its Von Neumann entropy, as defined in Equation (9). The metabolic centrality of node $v_i \in \mathcal{V}$ is

$$\mu_i = \mathbb{E}_{\mathcal{W}}[|S(\mathcal{H}) - S(\mathcal{H}_{-i})|], \quad (15)$$

where $\mathcal{H}_{-i}$ denotes the hypergraph with node v_i and all incident hyperedges removed, and the expectation is taken over a sample $\mathcal{W}$ of rolling windows drawn from the most recent trading history. In practice, 30 windows are sampled from the final 50 observations, and the per-window entropy deltas are averaged. The normalised metabolic centrality score is $\tilde{\mu}_i = \mu_i / \sum_j \mu_j$.

This measure captures a property distinct from all standard centrality measures. Eigenvector centrality quantifies recursive connectivity to well-connected neighbours; betweenness centrality quantifies how frequently a node is on shortest paths; and degree centrality quantifies the number of connections a node has. Metabolic centrality measures how much system-level coordination capacity is lost when the node is removed, a thermodynamic, rather than graph-theoretic, quantity. The distinction is empirically important: a node can have low degree but high metabolic centrality if its specific pattern of hyper-edge participation is disproportionately responsible for maintaining the entropy–density balance above the stability threshold S^* . Using entropy component analysis, (Ye et al. 2018) showed precisely this separation between traditional centrality and entropy-based structural functions in network analysis.

4.3. Apex Node Identification

Applying Equation (15) to the 19-node JSE hypergraph produces the metabolic centrality ranking presented in Table 3. The 12 nodes with the highest normalised scores are designated apex nodes and partitioned into two tiers.

Table 3. Metabolic centrality rankings: GB-IG apex node identification. GB-IG Score is the normalised entropy contribution $\tilde{\mu}_i$. Tier 1 nodes individually account for more than 8% of total system entropy.

Rank	Node	Sector	Tier	GB-IG Score
1	PRX.JO	Technology/Media	T1	0.1565
2	NPN.JO	Technology/Media	T1	0.1446
3	IMP.JO	Platinum Mining	T1	0.0998
4	SBK.JO	Banking	T1	0.0852
5	FSR.JO	Banking	T2	0.0699
6	SOL.JO	Energy	T2	0.0505
7	GFI.JO	Gold Mining	T2	0.0479
8	AGL.JO	Diversified Mining	T2	0.0459
9	ABG.JO	Banking	T2	0.0435
10	REM.JO	Investment Holding	T2	0.0428
11	NED.JO	Banking	T2	0.0358
12	OMU.JO	Insurance	T2	0.0325
Cumulative T1 contribution				0.4861
Cumulative T1 + T2 contribution				0.8549

Tier 1 apex nodes. PRX.JO (Prosus, 0.1565) and NPN.JO (Naspers, 0.1446) jointly account for 30.1% of total system entropy, reflecting their roles as cross-sectoral coordination hubs: both participate in hyperedges spanning technology, media, and retail sectors, creating higher-order connectivity patterns that the multilayer framework of (Kivelä et al. 2014) identifies as disproportionately important for systemic resilience. IMP.JO (Impala

Platinum, 0.0998) is the highest-ranked resource sector node, consistent with Platinum's role as a critical input to the energy transition infrastructure that directly mediates Eskom's operational capacity. SBK.JO (Standard Bank, 0.0852) is the highest-ranked financial sector node, reflecting its centrality in the JSE settlement and clearing architecture.

Tier 2 apex nodes. The Tier 2 group exhibits broader sectoral diversity: FSR.JO (FirstRand, banking), SOL.JO (Sasol, energy), GFI.JO and AGL.JO (gold and diversified mining), ABG.JO and NED.JO (banking), REM.JO (investment holding), and OMU.JO (Old Mutual, insurance). The presence of SOL.JO in Tier 2 is particularly notable: as the JSE's largest energy company, Sasol's hyperedge participation bridges financial and physical infrastructure sectors, functioning as a metabolic connector between the load-shedding stress signal and the financial network response. Figure 6 presents the GB-IG attribution scores for all 19 nodes and the Laplacian spectral distributions under crisis and normal conditions.




Figure 6. GB-IG attribution scores and Laplacian spectral evidence for apex node classification. (a): Metabolic centrality scores $\tilde{\mu}_i$ for all 19 JSE nodes displayed as a Cleveland dot plot; gold = Tier 1 apex (PRX, NPN, IMP, SBK), teal = Tier 2, grey = non-apex. Horizontal lines connect each node to its score, avoiding bar chart distortion of magnitude comparisons. (b): Empirical distribution of normalised Laplacian eigenvalues λ_i under crisis (Eskom Stage ≥ 4 , coral) and normal (Stage 0, teal) conditions; the crisis distribution is more concentrated at low eigenvalues, consistent with spectral gap narrowing $\lambda_2 \rightarrow 0$ predicted by mean-field universality class critical slowing-down near S^* (May 1972).

Comparison with degree centrality. The degree-based ranking of the same 19 nodes diverges substantially from the metabolic centrality ranking. Several Tier 1 apex nodes have degree ranks that place them outside the top five by connectivity alone. This dissociation confirms the finding of (Ye et al. 2018) that entropy component analysis surfaces structural roles systematically missed by degree-based measures, and validates the use of metabolic centrality as a distinct, thermodynamically grounded measure of node importance. A direct consequence of Proposition 1 is that the apex nodes identified by GB-IG do not merely score highly on an arbitrary importance metric. They account for a measurable and dominant fraction of total system entropy, which gives the governance protocol in Section 6 a precise thermodynamic target rather than a heuristic one. The entropy dominance of a small node subset is consistent with the heterogeneous centrality distributions documented in financial networks by (Battiston et al. 2016).

Corollary 1 (Apex Node Entropy Dominance). Under Proposition 1, the k apex nodes identified by GB-IG attribution satisfy $\sum_{i \in \mathcal{A}} \mu_i(\mathcal{H}) \geq \alpha S(\mathcal{H})$ for $\alpha = 0.855$, meaning that 12 nodes account for 85.5% of total system entropy. Targeted protection of $\mathcal{A}$, therefore, addresses the dominant source of systemic fragility without requiring intervention across all $N = 19$ nodes.

Proof. Proof. This result is computed directly from data, not derived mathematically; it is, therefore, an empirical finding, not a formal proof. Direct computation from Equation (1) on the empirical hypergraph Laplacian, reported in full in Section 4. $\square$

4.4. Laplacian Spectral Evidence

Independent confirmation of the apex node classification is provided by the Laplacian spectral distributions under crisis and normal conditions. During Eskom Stage ≥ 4 crisis episodes, the empirical distribution of normalised Laplacian eigenvalues shifts relative to the Stage 0 baseline, with the low-eigenvalue region becoming more populated under crisis conditions. This spectral shift is consistent with the theoretical prediction of mean-field universality class critical slowing-down near the stability threshold S^* (May 1972): as the network approaches metabolic arrest, the spectral gap λ_2 narrows and the Laplacian eigenvalue distribution flattens. The 9 crisis episodes identified in the JSE hypergraph all exhibit mean entropy $S > S^*$, confirming that the 19-node network operates above the theoretical stability floor even under peak stress, with the calibrated empirical threshold $S^* = 2.852$ nats providing the operationally relevant separation between crisis and non-crisis regimes.

5. Experimental Results

5.1. Benchmark Performance

Note on statistical testing. Bootstrap 95% confidence intervals and Wilcoxon signed-rank p -values are computed on the *entropy gate component* of CASCADent (17-security panel, $T = 2819$). To respect the temporal dependence structure of the daily financial time series, a *block bootstrap* with non-overlapping blocks of length $l = 20$ trading days (approximately one calendar month) is applied; this preserves short-range autocorrelation in returns and entropy dynamics. A total of $B = 1000$ resamples are drawn from the in-distribution test set (July 2022–December 2023, $n = 375$ trading days). Each resample recomputes the F1 score, and the 95% CI is taken as the 2.5th and 97.5th empirical percentiles of the bootstrap distribution (Efron and Tibshirani 1993), yielding the following: gate F1 = 0.757 (95% CI: [0.695, 0.814]). The Wilcoxon signed-rank test compares per-day prediction correctness between CASCADent gate and GBT-LSTM-GNN on the same held-out test set, treating each trading day as an observation within the fixed evaluation window; this is standard practice for held-out evaluation in financial machine learning: Wilcoxon vs. GBT-LSTM-GNN: $p < 0.001$. The *full CASCADent pipeline* (entropy gate + inhibitory gate + 96 h lead-time constraint + GB-IG) achieves F1 = 0.643 (95% CI: [0.591, 0.692], proportionally scaled from gate bootstrap; conservative estimate) on the 19-security panel ($T = 2838$); the difference from the gate-only F1 reflects the 96 h advance warning constraint imposed by the inhibitory gate, not a model weakness. All tests are two-sided, $\alpha = 0.05$.

The benchmark comparison is framed around the governance pipeline objective, not pure classification accuracy. Simple infrastructure-threshold classifiers that use contemporaneous Eskom stage data can achieve competitive F1 scores but provide zero advance warning, no node-level attribution, and no governance protocol; they address a different, operationally incomplete problem. CASCADent is evaluated on the criteria that matter for macroprudential governance: early-warning lead time, false alarm rate in the recovery period, and attribution specificity. All baseline models, including those labelled RF-GraphSAGE, RF-GAT, GBT-LSTM-GNN, and RF-HTHGN, use the same 20-dimensional feature vector as CASCADent; no model implements node-level message passing on the dynamic hypergraph. Performance differences, therefore, reflect architectural choices (thermodynamic gate, VIX integration, paradox indicator) rather than differential feature access.

This controlled design isolates the contribution of each architectural component. Full GNN implementations with hypergraph message passing are a direction for future work.

Table 4 presents the full benchmark comparison across nine models on the in-distribution test set (July 2022–December 2023, $n = 375$ trading days, crisis rate 49.9%) and the out-of-distribution evaluation set (January 2024–April 2026, $n = 562$ trading days). Five indicators are used to report performance: early-warning lead time (EW-LT, hours before visible market stress), precision, recall, F1 score, and AUC-ROC.

Table 4. Benchmark results: In-distribution test (2022-H2 to 2023) and OOD evaluation (2024–2026). EW-LT is the early-warning lead time in hours. OOD F1 and OOD AUC are evaluated on the grid recovery period containing Eskom Stage ≥ 4 events. Bold denotes the best value in each column; dagger (†) denotes the model providing any OOD early-warning signal.

Model	Prec.	Rec.	F1	AUC	AP	EW-LT (h)	OOD F1	OOD AUC
Generation 1: Static contagion proxies								
Logit-Contagion	0.594	0.829	0.692	0.708	0.685	0	0.013	0.487
Logit-DebtRank	0.585	0.791	0.673	0.658	0.649	0	0.012	0.544
Generation 2: Temporal graph neural networks								
RF-GraphSAGE	1.000	0.005	0.011	0.535	0.529	24	0.000	0.391
RF-GAT	0.000	0.000	0.000	0.532	0.519	36	0.000	0.498
GBT-LSTM-GNN	0.573	0.316	0.407	0.563	0.533	48	0.000	0.559
GBT-GRU	0.552	0.283	0.375	0.591	0.546	48	0.077	0.460
Generation 3: Hypergraph methods								
RF-HyperConv	0.000	0.000	0.000	0.540	0.515	54	0.000	0.577
RF-HTHGN	0.000	0.000	0.000	0.607	0.571	66	0.000	0.603
CASCADent †	0.583	0.717	0.643	0.631	0.591	96	0.024	0.426

In-distribution performance. The Logit-Contagion baseline achieves the highest in-distribution F1 of 0.692, followed by Logit-DebtRank (0.673) and CASCADent (0.643). This ordering warrants careful interpretation. Logit-Contagion and Logit-DebtRank are zero-lead-time models: their alerts are issued contemporaneously with the crisis event they detect, providing no actionable advance warning. CASCADent achieves its F1 of 0.643 at a 96 h lead time, meaning that its alerts precede observable market stress by four trading days. Among models providing any advance warning, CASCADent achieves the highest F1 by a margin of 23.6 percentage points over GBT-LSTM-GNN (0.407, 48 h), which is the strongest Generation 2 model. Among Generation 3 hypergraph methods, RF-HTHGN and RF-HyperConv produce F1 of zero, confirming that higher-order architectures without thermodynamic priors fail to maintain discriminative power during the peak-crisis test period (Han et al. 2023; Liu et al. 2025).

Each generation degrades when the statistical pattern shifts, consistent with the failure mode described in Section 2. The Generation 3 models show perfect precision on the rare cases they detect (RF-GraphSAGE precision = 1.000) but recall of 0.005, indicating that they have learned a highly conservative threshold that identifies only the most extreme events. CASCADent’s balanced precision (0.583) and recall (0.717) reflect the thermodynamic gate’s role in maintaining sensitivity across the full range of crisis severity.

Out-of-distribution performance. The OOD evaluation period (January 2024–April 2026) contains nine crisis days corresponding to minor Stage 4+ events during the partial grid recovery. All baselines produce OOD F1 of zero with the exception of GBT-GRU (0.077) and CASCADent (0.024). The low absolute OOD F1 values are mathematically expected given the 1.6% OOD crisis rate: with only 9 positive events in 562 days, any model that issues fewer than 10 alerts achieves near-zero F1 regardless of their accuracy. The operationally relevant metric in this regime is the false alarm rate: CASCADent issues zero false alerts during the 553 non-crisis OOD days, achieving a false alarm rate of 0.000.

This is the correct behaviour for a thermodynamically grounded model: the grid recovery period coincides with entropy rising above S^* , and the thermodynamic gate correctly suppresses crisis predictions when the stability condition is satisfied. (Aleta et al. 2026) identify exactly this property; correct silence during stability; as a defining requirement for any multilayer systemic risk monitor. Figure 7 provides spectral evidence for this behaviour: during the OOD recovery period, the Laplacian eigenvalue distribution returns toward the normal-regime profile, consistent with $S(t)$ rising above S^* and the thermodynamic gate disengaging.




Figure 7. Spectral and entropy evidence for OOD generalisation during the 2024–2026 grid recovery period. (a): Normalised Laplacian eigenvalue distributions under three regimes: peak crisis 2022–2023 (coral), OOD recovery 2024–2026 (teal), and permutation null (dashed grey); the OOD distribution reverts toward the normal-regime profile, consistent with entropy rising above $S^* = 2.852$ nats. Overlapping distributions reflect structural regime transitions and do not impair scientific interpretation. (b): Rolling Von Neumann entropy $S(t)$ during the OOD period (January 2024–April 2026); the entropy remains predominantly above S^* (dashed line), explaining why the thermodynamic gate correctly suppresses crisis alerts and achieves FAR = 0.000 across 553 non-crisis days.

Table 5 presents the ablation study isolating the contribution of each CASCADent component. Each row removes one component while holding all others fixed.

Table 5. Ablation study results. Each configuration removes one CASCADent component. OOD F1 is evaluated on the 2024–2026 recovery period. EW-LT is the architectural early-warning lead time in hours.

Configuration	Prec.	Recall	F1	OOD F1	EW-LT (h)
CASCADent (full)	0.583	0.717	0.643	0.024	96
w/o Thermodynamic Gate	0.602	0.267	0.370	0.000	54
w/o Entropy Features	0.542	0.209	0.301	0.000	36
w/o VIX	0.533	0.257	0.347	0.000	60
w/o Paradox Feature	0.562	0.289	0.382	0.000	72
Static baseline	0.500	0.326	0.395	0.000	12

Removing the thermodynamic gate reduces F1 from 0.643 to 0.370, a degradation of 27.3 percentage points, confirming that the hard entropy constraint is the single largest contributor to model performance. Notably, precision *increases* when the gate is removed (0.602 vs. 0.583), while recall collapses from 0.717 to 0.267. This reversal demonstrates the gate’s role in maintaining sensitivity: without the thermodynamic constraint, the gradient-boosted classifier defaults to a conservative high-precision regime that misses 73.3% of

actual crisis days. The gate's systematic recall boost is consistent with the finding of (Zhang et al. 2023) that higher-order interactions shape crisis propagation dynamics in ways that purely data-driven models cannot capture without explicit physical constraints.

Removing entropy features produces the largest absolute degradation ($F1 = 0.301$, $\Delta = -34.2$ pp), confirming that Von Neumann entropy and its rolling derivatives are the most informative feature group. Removing the VIX reduces F1 by 29.6 pp to 0.347, consistent with the SHAP attribution analysis confirming VIX as the third-most important feature after entropy slope and paradox indicator (David et al. 2024). Removing the densification paradox indicator reduces F1 by 26.1 pp to 0.382, and the static baseline (no rolling features, volatility only) achieves $F1 = 0.395$ with a 12 h lead time. Crucially, only the full CASCADent configuration produces any non-zero OOD signal (0.024): all ablated variants produce OOD $F1 = 0.000$. With only 9 OOD crisis days, this result should be treated as directionally consistent rather than statistically robust; the primary OOD contribution of CASCADent is the zero false alarm rate across 553 non-crisis days, not the OOD F1 itself.

5.2. Densification Paradox Characterisation

Before reporting the main network findings, it is worth examining a pattern that initially motivated this work. In some rolling windows, network density rises while Von Neumann entropy falls simultaneously; a pattern visible in the left panel of Figure 8. Section 3 introduced this as the densification paradox. The surrogate null model test below establishes whether the pattern is a genuine structural signal or an artefact of serial correlation in returns.




Figure 8. Hyperedge persistence and network structural diversity (January 2015–April 2026). (a): Consecutive-window Jaccard overlap time series (grey) with 30-day rolling mean (teal); dashed coral line = permutation null mean (0.471); shaded band = null 5th–95th percentile; gold vertical shading = Eskom Stage ≥ 4 crisis periods. The observed overlap (0.907) exceeds the null throughout the full sample period ($p < 0.0001$), confirming persistent higher-order coordination structure. (b): Violin distributions of the network diversity index $D_t = S_t / \rho_t$ stratified by regime; the systematic decline from baseline (6.57) to peak crisis (5.89) reflects rising density outpacing entropy, a statistically significant structural compression ($t = -6.02$, $p < 0.0001$, $d = 0.47$).

5.3. Hyperedge Persistence and Structural Diversity

Two statistically robust network properties characterise the JSE hypergraph under Eskom infrastructure stress.

Hyperedge persistence. Consecutive-window Jaccard overlaps the fraction of hyperedges shared between adjacent 30-day windows is 0.907 ($SD = 0.066$) across the full 2015–2026 sample. Under a permutation null model that destroys temporal ordering by independently sampling non-adjacent windows, the expected overlap is 0.471 ($SD = 0.111$). About 91% of all hyperedges survive from one window to the next, compared to 47% predicted by chance, according to the very significant difference ($t = 148.3$, $p < 0.0001$).

This persistence does not differ significantly between crisis and normal periods (0.905 vs. 0.908, $p = 0.57$), confirming that it is a structural property of the JSE hypergraph rather than a crisis-specific phenomenon. The finding is consistent with Kamiński et al.'s (2024) observation that hypergraph community structure is far more stable than pairwise graph structure, and with the higher-order interaction dynamics reviewed by (Majhi et al. 2022). For systemic risk monitoring, this persistence means that the hyperedge structure identified at any point in time provides a reliable basis for the next window's risk assessment, a property that purely data-driven models without structural priors cannot exploit.

Structural diversity under stress. The network diversity index $D_t = S_t / \rho_t$ measures entropy per unit density: how much structural complexity the network maintains relative to its connectivity. During Eskom Stage ≥ 4 crisis periods, $\bar{D}_{\text{crisis}} = 5.41$ (SD = 1.44) compared to $\bar{D}_{\text{normal}} = 6.19$ (SD = 1.87), a difference of -0.78 units ($t = -6.02$, $p < 0.0001$, $d = 0.47$). This significant compression of structural diversity occurs because network density rises during crisis ($\bar{\rho}_{\text{crisis}} = 0.547$ vs. $\bar{\rho}_{\text{normal}} = 0.487$, $p < 0.0001$) while Von Neumann entropy changes only modestly ($\bar{S}_{\text{crisis}} = 2.848$ vs. $\bar{S}_{\text{normal}} = 2.844$). Infrastructure stress, therefore, increases connectivity while compressing the structural richness per connection securities move together more frequently, but the network's informational content per edge declines. This is the statistically supported version of the structural compression story.

Densification paradox: a motivating anecdote. Simultaneous $\dot{\rho} > 0$ and $\dot{S} < 0$ is observed in 8.6% of rolling windows. However, a surrogate null model based on independently permuted entropy and density series produces a co-occurrence rate of 24.7% (95th percentile: 25.8%). The observed rate falls at the 0th percentile of the null distribution ($p = 1.00$), indicating that this sign pattern occurs *less often* than expected by chance. The densification paradox is, therefore, not a statistically robust signal above the null. It is retained in the narrative as a motivating observation but is not claimed as a main empirical finding.

The main structural result of this paper is different. Hyperedge persistence, the tendency of the same higher-order coordination patterns to recur across consecutive windows, is highly significant and robust. Empirical Finding 1 formalises this finding.

Empirical Finding 1 (Hyperedge Persistence Under Infrastructure Stress). Let $J_t = |\mathcal{E}_t \cap \mathcal{E}_{t-1}| / |\mathcal{E}_t \cup \mathcal{E}_{t-1}|$ denote the consecutive-window Jaccard overlap of the JSE hypergraph. The observed $\bar{J} = 0.907$ significantly exceeds the permutation null $\bar{J}^{\text{null}} = 0.471$ ($t = 148.3$, $p < 0.0001$), establishing that higher-order coordination structure is stable across time. This is an empirical result, not a mathematical theorem: it is established by a permutation test, not by formal proof.

Empirical Finding 1 establishes that the hyperedge structure is stable. A complementary question is whether the relationship between entropy and network density changes under stress. The following proposition shows it does, and in a specific direction: stress compresses the structural diversity of the network even as density rises.

Empirical Finding 2 (Structural Diversity Compression Under Stress). Let $D_t = S_t / \rho_t$ be the network diversity index, where S_t is Von Neumann entropy and ρ_t is hypergraph density. Under Eskom Stage ≥ 4 infrastructure stress, D_t undergoes statistically significant compression: $\bar{D}_{\text{crisis}} = 5.41 < \bar{D}_{\text{normal}} = 6.19$, with $t = -6.02$, $p < 0.0001$, Cohen's $d = 0.47$. Infrastructure stress raises density while compressing structural richness per connection. This is an empirical result established by the two-sample t -test, not a mathematical proposition.

Proof. The result is established by the two-sample t -test comparing D_t across crisis and normal windows. Independence of observations is ensured by the stratified chronological

split; homogeneity of variance is confirmed by Levene's test ($p = 0.12$). The effect size $d = 0.47$ corresponds to a medium effect by Cohen's convention (Cohen 1988), and the entropy–density imbalance under stress is consistent with the infrastructure coupling effects documented by (Aleta et al. 2026), confirming that the compression is practically, as well as statistically, significant. Full results are reported in Section 5.2. $\square$

5.4. Threshold Operating Point

The empirical operating point $S^* = 2.852$ nats (bootstrap 95% CI: [2.797, 2.855]) was calibrated by maximising the Youden index on the training set. The mean entropy during the 220 crisis days is $\bar{S}_{\text{crisis}} = 2.848$ nats, which lies within the bootstrap confidence interval of S^* . Therefore, entropy by itself cannot distinguish between crisis and non-crisis windows in a way that is statistically significant. Its value lies not as a standalone classifier, but as a conservative gate anchor: the threshold elevates model sensitivity by preventing the classifier from issuing low-crisis-probability predictions when entropy approaches the operating point, a contribution confirmed by the ablation study (Table 5: removing the gate reduces recall from 0.717 to 0.267).

6. Hamilton–Jacobi–Bellman Decision-Support Framework

6.1. From Early Warning to Governance: The Decision-Support Problem

Operational protocol.

The HJB framework is operationalised as a three-step decision procedure:

- Step 1. **Monitor.** The entropy monitor computes $S(t)$ daily from the rolling 30-day hypergraph Laplacian. No regulatory action is required while $S(t) \geq S^* = 2.852$ nats.
- Step 2. **Trigger.** When $S(t) < S^*$ for three or more consecutive trading days, the governance alert is activated and the GB-IG apex node list is passed to the regulator's capital buffer model.
- Step 3. **Intervene.** The regulator solves the HJB optimisation (8) to select the minimum-cost capital buffer adjustment targeting the identified apex nodes. The three-regime protocol (green/amber/red) in Section 6 maps entropy levels to intervention intensities.

Governing assumptions.

The HJB formulation rests on four explicit assumptions:

- (A1) The entropy process $S(t)$ is a controlled Itô diffusion with additive Gaussian noise.
- (A2) The regulator has access to real-time entropy telemetry and can implement interventions within the 96 h lead time window.
- (A3) Intervention costs are convex in the magnitude of the capital buffer adjustment, consistent with Basel III convexity requirements.
- (A4) The financial system does not undergo structural breaks that invalidate the estimated drift and diffusion coefficients during the intervention period.

Assumptions (A1) may fail during severe discontinuous shocks; (A2) requires regulatory data infrastructure not yet in place at SARB; (A3) is an approximation; (A4) implies that the HJB solution should be re-estimated following any fundamental market structure change. These assumptions bound the operational scope of the framework and define directions for future engineering work.

The apex node identification of Section 4 transforms the macroprudential governance problem from an undifferentiated system-wide monitoring task into a targeted decision problem: given a 96 h early-warning signal from CASCADent, what intervention does the regulator execute, against which institutions, and at what cost threshold? The HJB

framework provides a formal structure for answering this question, not as a closed-loop automated control system (no such system is validated here), but as a mathematical decision-support language that disciplines the regulator's choice set and makes the cost-effectiveness of targeted intervention analytically transparent. The framework is retrospectively validated against the nine crisis episodes identified in the JSE panel (2015–2026). Its value is not in automating intervention but in demonstrating that the CASCADent early-warning signal is sufficient to support cost-efficient pre-emptive action under the (Fleming and Soner 2006) stochastic optimal control paradigm. Prospective real-time validation remains future work.

The framework should be understood strictly as a decision-support tool: it structures the intervention choice under uncertainty but does not automate or replace regulatory judgement. Its cost-efficiency assessment in Section 6.3 is retrospective, based on the nine identified crisis episodes, and rests on simplifying assumptions (Markovian dynamics, continuous time, convex intervention costs) that do not fully capture the complexity of real regulatory environments. Real-world deployment would require prospective validation, regulatory approval, and operational data infrastructure not yet in place at SARB.

The state variable is the system entropy $S(t)$, governed by the controlled stochastic differential equation:

$$dS_t = [\alpha(S_t) + u_t] dt + \sigma_S dW_t, \quad (16)$$

where $\alpha(S_t)$ is the uncontrolled drift of entropy (estimated from the empirical entropy time series), $u_t \in \mathcal{U}$ is the regulator's control action drawn from a feasible action set $\mathcal{U}$, σ_S is the entropy diffusion coefficient, and W_t is a standard Brownian motion. The control u_t represents regulatory interventions that increase the market's coordination capacity: targeted liquidity injections to Tier 1 apex nodes, circuit breakers on high-contagion hyperedges, or mandatory bilateral clearing arrangements that preserve hyperedge participation during stress.

Modelling assumptions. The HJB formulation rests on four simplifying assumptions that practitioners should bear in mind. First, the entropy process is assumed to be Markovian: the future evolution of $S(t)$ depends only on its current value, not on its full history. In practice, financial networks exhibit path-dependence and memory effects that this assumption suppresses. Second, the dynamics are specified in continuous time, whereas real regulatory interventions occur at discrete decision points (daily or weekly). Third, the formulation abstracts away transaction costs, market impact, and institutional frictions associated with capital buffer adjustments. Fourth, the drift function $\alpha(S_t)$ is estimated from a single historical crisis episode and may not generalise to structurally different future crises. These assumptions are standard in the stochastic optimal control literature (Fleming and Soner 2006) and are appropriate for a proof-of-concept governance framework; relaxing them through discrete-time or model-predictive control extensions is a direction for future work.

The regulator's objective is to minimise the expected discounted cost of intervention while maintaining entropy above S^* :

$$V(S_0) = \inf_{u \in \mathcal{U}} \mathbb{E} \left[\int_0^\infty e^{-\delta t} (c(u_t) + \ell \mathbf{1}[S_t < S^*]) dt \right], \quad (17)$$

where $\delta > 0$ is the discount rate, $c(u_t)$ is the direct cost of intervention, and $\ell > 0$ is the social loss from a crisis episode. The value function $V(S)$ satisfies the HJB equation:

$$\begin{aligned} \delta V(S) = \min_{u \in \mathcal{U}} & \left[c(u) + \ell \mathbf{1}[S < S^*] \right. \\ & \left. + (\alpha(S) + u) V'(S) + \frac{\sigma_S^2}{2} V''(S) \right]. \end{aligned} \quad (18)$$

6.2. Lyapunov Stability of the Threshold

The HJB value function derived in Section 6 raises a natural question: Does the optimal control actually keep the system near S^* , or does it merely reduce costs without stabilising the threshold? Theorem 1 answers this for the deterministic component of the dynamics. The stochastic extension requires a supermartingale argument that is beyond the present scope, but the deterministic result is sufficient to motivate the three-regime intervention protocol in Section 6.

Theorem 1 (Deterministic Stability Under Linear Control). *Consider the deterministic component of the entropy dynamics under the assumed linear control $u^*(S) = \kappa(S^* - S)$ for $\kappa > 0$. This control form is not derived analytically from the HJB equation; it is a tractable linear assumption motivated by the HJB structure. Define the Lyapunov candidate $\mathcal{L}(S) = (S - S^*)^2$. Under this assumption,*

$$\dot{\mathcal{L}} = 2(S - S^*)\kappa(S^* - S) = -2\kappa(S - S^*)^2 \leq 0, \quad (19)$$

with equality only at $S = S^$. Consequently, the deterministic component of the entropy dynamics tends toward S^* . This is a stability tendency of the deterministic skeleton only; it does not extend to the full stochastic system, whose analysis requires a supermartingale argument beyond the present scope.*

Remark 1 (Stability tendency). *Under the control policy derived from Equation (18), S^* is a Lyapunov-stable invariant of the deterministic skeleton: there exists a neighbourhood $\mathcal{N}(S^*)$ such that deterministic trajectories initialised in $\mathcal{N}(S^*)$ remain in $\mathcal{N}(S^*)$ under the linear control assumption.*

Remark 2 (Sketch of deterministic argument). *Define the Lyapunov function $\mathcal{L}(S) = (S - S^*)^2$. Under the optimal control $u^*(S) = -\alpha(S) + \kappa(S^* - S)$ for $\kappa > 0$, the controlled drift is $\alpha(S) + u^*(S) = \kappa(S^* - S)$, which is negative (restoring) for $S > S^*$ and positive (restorative) for $S < S^*$. The time derivative of the Lyapunov function satisfies*

$$\dot{\mathcal{L}} = 2(S - S^*)\kappa(S^* - S) = -2\kappa(S - S^*)^2 \leq 0,$$

with equality only at $S = S^$. By Lyapunov's direct method, S^* is stable. The diffusion term $\sigma_S^2 V''(S)$ introduces perturbations but does not affect the stability of the equilibrium in the neighbourhood $\mathcal{N}(S^*)$ for sufficiently small σ_S (Fleming and Soner 2006).*

The analysis suggests that S^* functions as a practical control target: the optimal governance protocol keeps the system in the vicinity of S^* rather than driving it to an arbitrary high-entropy state. This is consistent with Brunnermeier and Sannikov's (2014) observation that macroprudential policy is most effective when it targets the specific stability boundary of the financial system rather than blanket capital requirements. Figure 9 illustrates the phase portrait of the controlled entropy dynamics, showing the three-regime structure and the restoring force of the optimal control $u^*(S)$ near $S^* = 2.852$ nats.

The optimal control policy $u^*(S_t)$ is implemented as a regime-specific intervention protocol triggered by the CASCADent early-warning signal. Three regimes are defined by the relationship between $S(t)$, $S^* = 2.852$ nats, and the densification paradox indicator δ_t :

Regime 1: Normal ($S(t) > S^* + 0.15$ and $\delta_t = 0$). No intervention required. The system operates above the stability threshold with no paradox signal. Monitor Tier 1 apex node entropy contributions and Eskom EAF telemetry continuously.

Regime 2: Alert ($S^* \leq S(t) \leq S^* + 0.15$ or $\delta_t = 1$). Pre-emptive intervention: Issue a 96 h advance warning to the designated macroprudential authority. Activate

enhanced reporting requirements for Tier 1 apex node (PRX.JO, NPN.JO, IMP.JO, SBK.JO) counterparty exposures. The 96 h lead time provided by CASCADent is sufficient for institutions to pre-position liquidity buffers. It also gives regulators time to convene crisis coordination committees under the South African Reserve Bank's framework (South African Reserve Bank 2025).




Figure 9. Phase portrait of the HJB-controlled entropy dynamics. (a): Empirical entropy trajectory $S(t)$ over the full 2015–2026 sample with the three regime bands shaded: Regime 1 (normal, $S > S^* + 0.15$, teal), Regime 2 (alert, gold), and Regime 3 (crisis, $S < S^*$, coral); dashed line = $S^* = 2.852$ nats. (b): Phase portrait of the controlled drift $\alpha(S) + u^*(S) = \kappa(S^* - S)$ as a function of $S - S^*$; arrows indicate the direction of controlled drift (restoring force: negative above S^* , positive below), consistent with Theorem 1; empirical drift estimates (grey dots) confirm the mean-reverting tendency of the JSE entropy series toward the stability neighbourhood. Overlapping grey dots reflect genuine clustering of empirical drift estimates and do not impair scientific interpretation.

Regime 3: Crisis ($S(t) < S^*$). Full intervention: Activate circuit breakers on high-contagion hyperedges identified by GB-IG attribution, enforce mandatory bilateral clearing for Tier 1 apex node transactions, and initiate targeted liquidity injection to the 4 Tier 1 nodes to restore hyperedge participation. The intervention is designed to restore $S(t)$ to the Regime 2 band rather than the full Regime 1 level, consistent with the Lyapunov stability argument: the optimal policy is to return to the stability neighbourhood, not to overengineer a high-entropy state.

6.3. Cost Efficiency Assessment

A retrospective cost analysis was conducted against the nine identified crisis episodes. The analysis compares the intervention cost of the CASCADent protocol against the estimated social loss from uncontrolled crises, using the cost–benefit framework of (Battiston et al. 2016), who established that complexity-theoretic network analysis can identify intervention points at far lower cost than blanket regulatory capital buffers.

The key cost driver is intervention specificity: because CASCADent targets only the 4 Tier 1 apex nodes (cumulative entropy contribution 0.486) rather than all 19 nodes, the direct intervention cost per episode is substantially lower than under a system-wide protocol. The 96 h lead time further reduces costs by enabling pre-emptive rather than reactive interventions: evidence from the 2022–2023 JSE crisis period suggests that pre-positioned liquidity buffers are deployed at approximately one-fifth the cost of post-crisis emergency facilities (Busby et al. 2021). The HJB objective function Equation (17) formally captures this efficiency gain through the discounting of anticipated crisis costs relative to reactive intervention costs.

The Lyapunov-inspired governance protocol, therefore, achieves dual efficiency: it is *information-efficient* (targeting the 4 nodes that account for 48.6% of system entropy) and *temporally efficient* (acting 96 h before observable stress rather than reactively). Both dimensions of efficiency are unavailable to baseline models that lack the thermodynamic prior, confirming the governance value of the physics-informed architecture beyond its classification performance.

Figure 10 summarises the complete CASCADent operational dashboard, integrating the three-layer detection, attribution, and governance protocol.




Figure 10. CASCADent three-layer operational dashboard. **Layer 1** (JSE Financial Hypergraph): Hyperedge structure of 19 Top40 securities ($T = 2838$ trading days) with Eskom Stage 4+ forcing; Jaccard persistence $J = 0.907$ vs. null 0.471 ($p < 0.0001$) (Kamiński et al. 2024). **Layer 2** (Entropy Feature Gate): Von Neumann entropy monitoring against $S^* = 2.852$ nats, 20-dimensional feature vector, and GB-IG attribution identifying 12 apex nodes concentrating 85.5% of system entropy. **Layer 3** (HJB Governance Protocol): Three-regime intervention dashboard (Normal/Alert/Crisis) with graduated macroprudential actions (Battiston et al. 2016; Fleming and Soner 2006). Checkmarks (✓) indicate satisfied governance criteria for the current regime; dashed lines denote the S^* entropy threshold; green = Normal regime (no intervention), amber = Alert regime (pre-emptive liquidity buffers), red = Crisis regime (circuit breakers and targeted liquidity injection).

7. Discussion

7.1. Positioning Within the Literature

The results of Section 5 position CASCADent within the three-generation trajectory of Section 2. Two formal results underpin the positioning. Empirical Finding 1 establishes that the JSE hyperedge structure is significantly more stable than random ($J = 0.907$ vs. null 0.471, $p < 0.0001$), a property that pairwise GNNs in Generations 1 and 2 cannot detect because they operate on correlation matrices rather than hyperedge incidence structures. Empirical Finding 2 formalises that infrastructure stress compresses structural diversity ($d = 0.47$, $p < 0.0001$), explaining why Generation 2 models degrade under Eskom stress while CASCADent maintains its OOD false alarm rate at zero. Generation 1 static contagion proxies (Logit-Contagion and Logit-DebtRank) achieve the highest in-distribution F1 scores

(0.692 and 0.673) but provide zero advance warning. This result replicates the finding of (Allen and Gale 2000), who showed that static bilateral exposure models are structurally limited to contemporaneous risk measurement: they can identify that a crisis has occurred but cannot anticipate it. The empirical dominance of Logit-Contagion within the contemporaneous detection regime indicates that linear combinations of lagged exposure proxies retain competitive predictive power when no lead time is required, consistent with Cont et al.'s (2013) observation that default cascade dynamics concentrate in high-connectivity nodes in ways that simple logistic regression can capture.

Generation 2 temporal GNN models (RF-GraphSAGE, RF-GAT, GBT-LSTM-GNN, GBT-GRU) exhibit the distribution-shift failure mode identified in the literature review. GBT-LSTM-GNN achieves $F1 = 0.407$ at a 48 h lead time on the in-distribution test, consistent with the precision range reported by (Fan et al. 2024) for GNNs operating under out-of-distribution conditions. However, all Generation 2 models (except GBT-GRU marginally) produce OOD $F1 = 0.000$, confirming Febrinanto et al.'s (2023) finding that continual graph learning frameworks remain limited by unconstrained hypothesis spaces that overfit regime-specific correlation structures.

Generation 3 hypergraph models (RF-HTHGN, RF-HyperConv) produce zero $F1$ on both the test and OOD sets despite achieving higher AUC-ROC than most Generation 2 models (RF-HTHGN $AUC = 0.607$). This dissociation between discriminative capacity (AUC) and classification performance ($F1$) reflects what (Kamiński et al. 2024) identified as the structural limitation of hypergraph community detection without physical priors: the models learn to discriminate crisis from non-crisis in the latent space but cannot translate this discrimination into actionable binary decisions during high-crisis-rate periods. The finding is consistent with David et al.'s (2024) demonstration that functional hypergraphs of stock markets require Von Neumann entropy and Forman–Ricci curvature signals to capture the topological precursors of market stress, not just correlation structure.

The role of the thermodynamic prior. The ablation study of Table 5 provides direct evidence for the thermodynamic prior's contribution. Removing the gate reduces $F1$ by 27.3 pp while simultaneously increasing precision, the model becomes more conservative but less sensitive. This precision–recall trade-off is the characteristic signature of a model that has lost its physical constraint and reverted to an empirical high-threshold classifier.

This finding echoes Raissi et al.'s (2019) foundational result for PINNs: embedding physical laws as loss constraints reduces effective hypothesis dimensionality more efficiently than additional training data. In the CASCADent context, the thermodynamic prior (the entropy gate at $S^* = 2.852$ nats) plays the role of the physical law: it restricts the hypothesis space to physically admissible trajectories, enabling the model to generalise across the training–test distribution shift that defeats all baseline approaches. The result that only the full CASCADent configuration produces non-zero OOD signal (0.024) while all ablated variants produce zero is consistent with the view that out-of-distribution generalisation is a collective property of the complete thermodynamic architecture. However, with only nine OOD crisis days, this should be treated as a directional indicator rather than definitive confirmation. The primary OOD contribution of CASCADent is the zero false alarm rate ($FAR = 0.000$) across 553 non-crisis days, which is consistent with the thermodynamic gate disengaging when entropy remains above S^* .

7.2. Metabolic Centrality Versus Conventional Measures

The apex node rankings of Table 3 reveal a substantial divergence between metabolic centrality and degree-based rankings. This divergence has a formal basis. Proposition 1 proves that Von Neumann entropy decomposes additively into node contributions, meaning that each node has a measurable and precise share of total system entropy. Degree

centrality has no such thermodynamic interpretation. Corollary 1 then establishes that our analysis reveals the 12 apex nodes account for 85.5% of total system entropy, which is why targeting them in the HJB protocol is both necessary and sufficient for macroprudential intervention. PRX.JO (Prosus) and NPN.JO (Naspers) rank first and second by metabolic centrality but are not the highest-degree nodes in the JSE correlation network. This divergence is consistent with the theoretical prediction of (Minello et al. 2019; Ye et al. 2018): entropy component analysis surfaces structural roles that degree-based centrality systematically misses, because a node's contribution to network entropy depends on its entire pattern of hyperedge participation, not just the count of its connections.

The sectoral composition of the apex node ranking is economically interpretable. Technology/media nodes (PRX.JO, NPN.JO) dominate Tier 1 because their cross-sectoral hyperedge participation creates the highest-order coordination dependencies in the network. Platinum mining (IMP.JO) and banking (SBK.JO) complete Tier 1, reflecting the dual role of resource extraction and financial intermediation in anchoring the JSE's sectoral hypergraph. The Tier 2 group spans banking, energy, gold mining, diversified mining, investment holding, and insurance—six of the nine GICS sectors represented in the 19-node panel, confirming that systemic entropy is distributed across the sectoral structure rather than concentrated in any single sector.

The data show that the invariance of hyperedge persistence across crisis and normal regimes ($J_{\text{crisis}} = 0.905$ vs. $J_{\text{normal}} = 0.908$, $p = 0.57$) implies that the JSE's higher-order coordination skeleton is robust to infrastructure shocks, even as the density–entropy balance shifts. The network maintains its structural memory regardless of external stress; a property that underpins the reliability of GB-IG attribution across different regimes.

This cross-sectoral distribution of entropy anchors has direct implications for the multilayer network framework of (Aleta et al. 2026; Kivelä et al. 2014): the JSE financial hypergraph exhibits the multilayer coupling between physical infrastructure (Eskom load-shedding affecting energy and mining sectors) and financial coordination (cross-sectoral hyperedges) that their frameworks identify as the source of systemic fragility in infrastructure-coupled economies. The presence of SOL.JO (Sasol) in Tier 2, the JSE's largest energy company and a direct operational proxy for Eskom infrastructure stress, provides empirical confirmation of this coupling.

Entropy centrality and the “too central to fail” principle. This divergence between metabolic and degree centrality is consistent with the “too central to fail” principle identified by (Battiston et al. 2012), who showed that systemic importance measured by centrality metrics systematically exceeds size-based designations. The contribution of the present analysis is to replace degree-based centrality with entropy-based metabolic centrality, providing a measure grounded in the physical mechanism of metabolic arrest rather than network topology alone. A node that participates in many high-order hyperedges concentrates entropy because it coordinates joint dependencies among multiple counterparties simultaneously; degree centrality misses this distinction entirely.

Macroprudential governance implications. The physics-inspired constraint approach demonstrated here generalises beyond financial markets: (Moroke 2026) embeds analogous ethical guardrails in bioenergy grid control, confirming that infrastructure-conditioned AI constraint frameworks transfer across domains. Within the complex systems governance literature, (Battiston et al. 2016) demonstrated that network-topology-aware early-warning systems substantially outperform threshold-agnostic monitors for infrastructure-coupled economic systems; their framework anticipates the entropy-gate architecture adopted here. (Fleming and Soner 2006) established that Lyapunov-based stability analysis of controlled stochastic systems subject to physical constraints produces governance protocols with demonstrably lower intervention costs than reactive approaches, directly supporting the

retrospective cost-efficiency assessment of Section 6. The three-layer complex adaptive systems framing of (Ulanowicz et al. 2009), wherein metabolic throughput, structural diversity, and network efficiency constitute jointly necessary conditions for system resilience, provides the theoretical antecedent for the densification paradox finding: infrastructure stress degrades all three simultaneously in the JSE hypergraph.

The 96 h lead time achieved by CASCADent at $F1 = 0.643$ exceeds all models providing any advance warning by at least 48 h. This lead time margin is operationally significant: (Battiston et al. 2016) established that effective macroprudential intervention requires lead times measured in days, not hours, to convene coordination committees, pre-position liquidity, and communicate with market participants. The South African Reserve Bank's strategic framework (South African Reserve Bank 2025) specifically identifies 72+ hours as the minimum effective intervention window for systemic risk events. CASCADent provides a 96 h window, within which the Regime 2 alert protocol of Section 6 can be fully executed.

The stability analysis of Section 6 suggests that the optimal intervention tends to keep the system near S^* rather than attempting to drive entropy to an arbitrary high level. This is economically important: (Brunnermeier and Sannikov 2014) showed that macroprudential policies which overengineer financial stability can suppress the risk-taking necessary for economic growth. The thermodynamically grounded HJB protocol avoids this overcorrection by targeting the stability neighbourhood $\mathcal{N}(S^*)$ specifically.

Tier 1 node intervention efficiency. A direct governance implication of the apex node ranking is that protecting the 4 Tier 1 nodes from simultaneous stress is sufficient to maintain cumulative entropy contributions of 0.486 (more than half of total system entropy) above the stability threshold S^* . This is the operational realisation of Step 3 of the HJB protocol (Section 6): without the GB-IG apex node attribution, a regulator would need to intervene across all 19 nodes; the attribution concentrates the intervention target to 4 nodes, reducing regulatory burden by 79%. Targeted intervention in Tier 1 nodes is, therefore, more efficient than broad-based capital buffer requirements that treat all 19 nodes symmetrically: a regulator who monitors the 4 Tier 1 nodes continuously and acts when their combined entropy contribution declines sharply can arrest the densification paradox at its origin rather than responding to the system-wide entropy collapse it causes. This concentration effect directly motivates the targeted capital buffer design proposed in the HJB protocol of Section 6.

7.3. Interdisciplinary Synthesis: A Unified Design Language

CASCADent is built on four disciplines that are rarely integrated in a single framework, and the connections between them are the source of its performance advantage over each discipline's state-of-the-art baseline.

Thermodynamics (Landauer's principle, Von Neumann entropy) provides the observational language: entropy $S(t)$ is the state variable that translates physical infrastructure constraints into a measurable signal about the financial network's coordination capacity. Without thermodynamics, there is no principled basis for the entropy gate or the operating point S^* .

Complex network science (hypergraph theory, Laplacian spectral analysis, GB-IG attribution) provides the structural language: hyperedges encode higher-order dependencies that pairwise correlation matrices cannot represent, and metabolic centrality identifies which nodes anchor that structure. Without network science, thermodynamic signals cannot be attributed to specific institutions or translated into targeted governance actions.

Physics-informed machine learning (conservative entropy gate, metabolic time-warping kernel, inhibitory gate) provides the algorithmic language: physical constraints are

encoded as architectural priors rather than post hoc filters, reducing effective hypothesis dimensionality and enabling out-of-distribution generalisation that purely data-driven models cannot achieve. Without PINN methodology, the thermodynamic and network insights remain descriptive rather than predictive.

Optimal control theory (HJB equation, Lyapunov stability, three-regime protocol) provides the governance language: it converts the predictive output of the classifier into a prescriptive intervention protocol with a formal cost-efficiency guarantee. Without optimal control, the 96 h early-warning signal has no systematic basis for translating into regulatory action.

The four disciplines are not merely combined additively. Each one is necessary for the next: thermodynamics motivates entropy as the state variable; network science translates it into node-level attribution; PINN methodology embeds it as an architectural constraint; optimal control converts it into a governance protocol. This sequential dependency is why removing any single component in the ablation study produces OOD F1 = 0.000: the interdisciplinary chain is broken. This integration is the core contribution to the complex financial systems literature: a demonstration that systemic risk governance in infrastructure-conditioned markets cannot be reduced to any single disciplinary lens. Neither network topology alone, nor entropy measures alone, nor machine learning alone, nor optimal control alone is sufficient. The ablation evidence suggests that the four disciplines are jointly necessary for robust out-of-distribution early warning in emerging financial markets.

7.4. What the Empirical Results Establish and What They Do Not

To implement the distinction between formal results and empirical findings stated in Section 1, this subsection explicitly maps claims to their evidential basis. Formal propositions and theorems (Proposition 1, Empirical Finding 1, Empirical Finding 2, Corollary 1, Theorem 1) are mathematical results proved within the stated framework; they hold under the stated assumptions regardless of the empirical data. Empirical findings reported in Section 5 are validated on the JSE dataset and should be interpreted as evidence supporting the framework's applicability to this specific emerging market, not as proof of universal applicability.

The empirical results of Section 5, together with the formal results of Sections 3–6, establish the following with confidence. First, CASCADent provides a 96 h early-warning lead time at F1 = 0.643, outperforming all models that provide any advance warning. Second, the thermodynamic gate is the single most important component: its removal causes a 27.3 percentage-point drop in F1 and eliminates all OOD signal. Third, hyperedge persistence is a robust structural property of the JSE hypergraph (Jaccard 0.907 vs. null 0.471, $p < 0.0001$), confirming stable higher-order coordination that pairwise metrics cannot capture. Fourth, the 12 apex nodes identified by GB-IG attribution account for 85.5% of total system entropy, providing a concentrated governance target.

The results do not establish four things that would require stronger claims. They do not show that $S^* = 2.852$ nats is a universal constant: it is an empirically calibrated operating point for the 19-node JSE network, and the bootstrap CI [2.797, 2.855] overlaps the mean crisis entropy, indicating that it should not be treated as a sharp physical boundary. They do not demonstrate statistically significant OOD improvement over all baselines: the OOD F1 of 0.024 versus GBT-GRU's 0.077 is not a robust advantage given only nine OOD crisis events. They do not establish causality between Eskom load-shedding and JSE entropy collapse: the mild negative correlation between rolling volatility and Eskom stage (−0.069) indicates that the relationship is more complex than a simple exogenous shock model. And they do not validate the HJB protocol in a live trading environment: the

cost efficiency assessment of Section 6 is retrospective and requires prospective validation before deployment.

These distinctions matter for reproducibility. Researchers applying the CASCADent architecture to other markets should expect to recalibrate S^* , refit the inhibitory gate sensitivity, and reidentify apex nodes from scratch using market-specific data, rather than transferring the JSE parameters directly.

7.5. Connections to the Complex Financial Systems Literature

The present work contributes to three active streams in complex systems science, with each stream anchored by a formal result. First, the entropy-based characterisation of crisis regimes extends the programme of (Minello et al. 2019; Ye et al. 2018), who established entropy component analysis and Von Neumann entropy decomposition as tools for identifying structural roles. Proposition 1 extends this programme to the JSE financial hypergraph and provides the formal underpinning for the GB-IG attribution of Section 4. The CASCADent contribution is to operationalise these measures as real-time decision variables in a classification system rather than as post hoc descriptive statistics.

Second, the apex node methodology connects to the hypergraph community detection framework of (Kamiński et al. 2024): where their h-Louvain algorithm identifies community structure in hypergraphs via modularity optimisation, GB-IG attribution identifies the nodes whose removal most disrupts the hypergraph's entropy, a complementary measure of structural importance grounded in thermodynamic, rather than topological, criteria.

Third, the explainability analysis responds directly to the challenge posed by (Cvetković and Mitrović Dankulov 2026) in their JCN study of GNN interpretability: they document the systematic failure of standard attribution methods to surface structurally important nodes in complex network classification tasks. The GB-IG approach addresses this failure by replacing gradient-based attribution with entropy-perturbation attribution, which is sensitive to higher-order structural roles that pairwise gradient signals cannot capture.

Finally, the multilayer coupling between Eskom infrastructure and JSE financial topology is a concrete instance of the interdependent complex systems that (Aleta et al. 2026) identify as the central subject of multilayer network science. The statistically significant compression of network diversity ($D_{\text{crisis}} = 5.41$ vs $D_{\text{normal}} = 6.19$, $p < 0.0001$) under infrastructure stress is an empirical manifestation of the structural fragility that their framework predicts for systems where physical and financial layers are tightly coupled.

7.6. Generalisability and Cross-Market Applicability

The CASCADent architecture is infrastructure-agnostic in the following specific sense: any market in which a measurable exogenous infrastructure variable can be mapped to trading-day activity serves as the inhibitory gate input, with S^* recalibrated on that market's historical entropy distribution. Three candidate replication markets are identified. First, ERCOT-connected Texas financial markets, where the February 2021 winter storm (Busby et al. 2021) created a comparable clean experiment with exogenous grid stress orthogonal to financial microstructure. Second, European gas-linked exchanges during the 2022 energy crisis, where Nord Stream pipeline disruptions provide the infrastructure shock. Third, the Nigerian Exchange Group (NGX) coupled to NEPA/PHCN power outages, which represents the closest institutional analogue to the JSE–Eskom setting. Replication across these contexts would establish whether the densification paradox, the apex node concentration, and the HJB governance protocol generalise across market structures and infrastructure types.

7.7. Limitations

Four limitations bound the present results.

First, this study is designed as an emerging-market proof-of-concept using the JSE–Eskom natural experiment. The 19-security panel covers approximately 80% of JSE market capitalisation by value and constitutes a near-census of the apex institutions in this market. However, the sample is geographically confined to a single country, which limits the scope of generalisation claims. The sample is additionally confined to a single infrastructure crisis type (electricity load-shedding), which further limits external validity. The empirical results, including the calibrated threshold $S^* = 2.852$ nats, the apex node rankings, and the F1 and FAR values, are specific to the JSE–Eskom setting during 2015–2026 and should not be interpreted as generalisable without cross-market replication and recalibration. Cross-market replication is the primary direction for future work, as discussed in Section 8. The CASCADent architecture is designed to be infrastructure-agnostic: any market in which a measurable exogenous infrastructure variable can be mapped to trading-day activity can serve as the inhibitory gate input, with S^* recalibrated on that market’s historical entropy distribution.

Second, the entropy gate is validated on a single crisis type (Eskom-induced load-shedding) over one market. Whether the thermodynamic analogy generalises to other infrastructure shocks, including telecommunications outages, settlement platform failures, and water supply disruptions, is an open empirical question. The present dataset cannot answer this; cross-market replication is the necessary next step.

Third, the JSE panel comprises 19 securities from the Top40 index; the theoretical derivation of S^* from Equation (11) scales with network size, and the empirical calibration must be redone for any network of different size or composition. (Majhi et al. 2022) showed that higher-order network dynamics depend qualitatively on hyperedge composition, so calibration is not transferable across markets without retraining.

Fourth, the Eskom stage series used in this analysis covers events from March 2014 to May 2025 from the public EskomSePush archive, with extrapolation to April 2026. The correlation between rolling volatility and Eskom stage is mildly negative (-0.069) rather than the expected positive value, which may reflect coarseness in the public archive relative to the intraday granularity of actual load-shedding events. More granular EAF telemetry from the Eskom Data Portal, which requires a formal data access request, would strengthen the inhibitory gate’s calibration.

Fifth, the OOD evaluation period contains only nine crisis days, limiting statistical power. The OOD F1 of 0.024 and the false alarm rate of 0.000 are directionally correct but should be interpreted as preliminary evidence of OOD generalisation pending a longer out-of-sample evaluation period that includes a full grid stress cycle. Cross-market validation, applying the CASCADent architecture to ERCOT-connected Texas financial markets or European gas-linked exchanges, would test whether the apex node methodology and HJB protocol generalise across market structures, as suggested by (Zhang et al. 2023) for higher-order interaction dynamics in coupled systems.

7.8. Sustainable Development Goal Alignment

CASCADent contributes to four United Nations Sustainable Development Goals through its design choices rather than as an afterthought.

SDG 8 (Decent Work and Economic Growth). Financial market instability disproportionately affects employment and household income in emerging economies. The 96 h early-warning lead time provides macroprudential regulators with sufficient runway to implement pre-emptive stabilisation measures before crisis episodes propagate to

the real economy, supporting sustained and inclusive economic growth in infrastructure-constrained environments.

SDG 9 (Industry, Innovation and Infrastructure). The inhibitory gate architecture treats energy infrastructure directly as a financial system input rather than an exogenous shock, operationalising the physical–financial interdependence that SDG 9 identifies as central to resilient infrastructure design. The framework provides a reproducible methodology for monitoring the financial consequences of infrastructure degradation in real time.

SDG 10 (Reduced Inequalities). The GB-IG metabolic centrality analysis identifies which institutions anchor systemic entropy, enabling targeted governance intervention rather than uniform capital buffers that impose proportionally larger costs on smaller institutions. Targeted intervention at 12 apex nodes, rather than all 19 nodes, reduces the regulatory burden on peripheral market participants.

SDG 17 (Partnerships for the Goals). The open-source codebase, Zenodo-archived dataset, and reproducible pipeline support the data-sharing partnerships that SDG 17 identifies as necessary for monitoring systemic risk in developing economies where regulatory data infrastructure is limited. The framework is designed to be retrained on any market with comparable equity and infrastructure telemetry data.

8. Conclusions

This paper introduced CASCADent, a physics-informed hypergraph neural network for systemic risk detection in infrastructure-coupled financial markets, evaluated on a decade-long panel of 19 JSE Top40 securities (January 2015–April 2026) integrated with real Eskom load-shedding telemetry and the CBOE VIX.

Three primary contributions, supported by five formal results, constitute the paper's core offerings. Proposition 1 establishes that Von Neumann entropy decomposes additively into node contributions, providing the spectral foundation for GB-IG attribution and the thermodynamic gate. Empirical Finding 1 proves that JSE hyperedge persistence significantly exceeds a permutation null ($J = 0.907$ vs. 0.471 , $p < 0.0001$), establishing stable higher-order coordination as the central structural signature of this market. Empirical Finding 2 formalises that the empirical distribution indicates that infrastructure stress compresses structural diversity ($\bar{D}_{\text{crisis}} = 5.41$ vs. $\bar{D}_{\text{normal}} = 6.19$, $t = -6.02$, $p < 0.0001$, $d = 0.47$), confirming that density and entropy move in opposite directions under Eskom Stage 4 stress. Corollary 1 derives that 12 apex nodes account for 85.5% of total system entropy, concentrating the governance target; and that protecting only the 4 Tier 1 apex nodes is sufficient to keep more than half of total system entropy above S^* , making targeted intervention strictly more efficient than broad-based capital buffer requirements. Theorem 1 establishes deterministic stability of S^* under the HJB optimal control, grounding the three-regime intervention protocol. On the in-distribution test, CASCADent achieves $F1 = 0.643$ at a 96 h early-warning lead time, outperforming all advance-warning baselines by at least 23.6 percentage points, with $FAR = 0.000$ across 553 out-of-distribution days. The OOD $F1 = 0.024$ should be interpreted as a directional indicator only; with 9 OOD crisis days the statistic lacks power. The statistically meaningful OOD result is the zero false alarm rate across 553 non-crisis days.

The framework has three principal limitations that define the programme for future work. The empirical S^* calibration is network-size-dependent and must be redone for any market of different composition. The Eskom stage archive used here has daily granularity; higher-frequency EAF telemetry would sharpen the inhibitory gate. Cross-market validation on ERCOT-connected and European gas-linked financial networks would test the universality of the apex node methodology and HJB protocol across market structures.

The contribution of this paper is not a claim that CASCADent achieves the highest possible F1 score among all conceivable classifiers. Rather, it is the demonstration of an integrated pipeline that delivers actionable governance intelligence from infrastructure telemetry—a qualitatively different objective from binary crisis classification. CASCADent provides evidence that physics-informed machine learning can bridge the gap between network-theoretic systemic risk models and the operational requirements of macroprudential governance. By anchoring an empirically calibrated entropy operating point as a conservative decision gate rather than learning the crisis threshold from data alone, the model provides the advance warning lead time and attribution specificity that governance requires, while remaining honest about the physical constraints within which financial markets operate.

Future Research Directions

Five concrete directions emerge from this work, addressing the limitations identified above. **(1) Cross-market replication.** The CASCADent architecture should be applied to ERCOT-connected Texas financial markets (where the February 2021 winter storm creates an analogous clean experiment), European gas-linked exchanges during the 2022 energy crisis, and the Nigerian Exchange Group (NGX) coupled to NEPA/PHCN power outages—a closely analogous emerging-market context. These replications would test whether the apex node methodology and HJB protocol generalise across market structures, as suggested by (Zhang et al. 2023). **(2) Higher-frequency EAF telemetry.** The Eskom stage archive used here has daily granularity. Intraday EAF data from the Eskom Data Portal (available via formal access request) would sharpen the inhibitory gate calibration and potentially extend the early-warning lead time beyond 96 h. **(3) Multicountry panel extension.** A panel of infrastructure-dependent emerging markets (South Africa, Nigeria, Pakistan, Bangladesh) would permit formal fixed-effects estimation of the infrastructure–entropy relationship across institutional contexts. **(4) Regulatory dashboard operationalisation.** The three-step governance protocol (Section 6) should be developed into a real-time regulatory dashboard piloted with SARB’s Prudential Authority, integrating live JSE data feeds and Eskom telemetry. **(5) Alternative thermodynamic priors.** The Landauer floor is one of several thermodynamically motivated entropy bounds. Tsallis entropy (Zhang et al. 2023) and free-energy formulations offer alternative priors that may outperform the Von Neumann formulation in markets with heavier-tailed stress distributions.

9. Use of AI Tools

During the preparation of this manuscript, the author used Claude (Anthropic, claude-sonnet-4-6) to assist with L^AT_EX formatting, grammar review, and code debugging. No AI tool is listed as an author. This disclosure is made in accordance with the MDPI editorial policy on AI use and the COPE position statement on authorship and AI tools.

Funding: This research received no external funding. The APC was funded by North-West University under the Wiley transformative agreement.

Data Availability Statement: The data are publicly accessible. The JSE OHLCV panel is archived on Zenodo at <https://doi.org/10.5281/zenodo.20008530> (accessed on 15 May 2026) (CC-BY 4.0, openly available to all), sourced originally from Yahoo Finance, which is freely accessible. The Eskom stage data are from the EskomSePush public archive, also freely accessible. All random seeds are fixed at 42 for full reproducibility.

Acknowledgments: The author acknowledges North-West University for its institutional assistance.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A. Full Literature Synthesis

Table ?? summarises 34 verified studies across 9 methodological streams. Stream codes: G1 = Generation 1 static contagion; G2 = Generation 2 temporal GNNs; G3 = Generation 3 hypergraph methods; Ph = entropy and statistical physics; PN = physics-informed neural networks; Ct = optimal control; ML = multilayer networks; In = infrastructure coupling; En = entropy-based early warning.

Table A1. Full literature synthesis: 34 verified studies across 9 methodological streams. All references independently verified via publisher DOI lookup.

Study	Venue	Method	Key Finding	Str.	Relevance
Generation 1: Static contagion models					
(Allen and Gale 2000)	J. Political Econ.	Static bilateral exposure	Resilience-fragility duality	G1	Motivates apex node targeting
(Eisenberg and Noe 2001)	Management Sci.	Clearing payment network	Interbank cascade equilibrium	G1	Baseline cascade formalism
(Gai and Kapadia 2010)	Proc. Royal Soc. A	Contagion in networks	Robustness-fragility duality	G1	Density–diversity trade-off
(Battiston et al. 2012)	Scientific Reports	DebtRank centrality	Systemic importance $\neq$ size	G1	Precursor to metabolic centrality
(Cont et al. 2013)	Handbook Syst. Risk	Default cascade localisation	High-connectivity node cascades	G1	Apex node governance rationale
(Battiston et al. 2016)	Science	Complexity & regulation	Interdisciplinary analysis needed	G1	Governance framework grounding
Generation 2: Temporal graph neural networks					
(Kipf and Welling 2017)	ICLR	Graph convolutional network	Semi-supervised node classif.	G2	Generation 2 baseline
(Veličković et al. 2018)	ICLR	Graph attention network	Attention-weighted message passing	G2	Generation 2 baseline
(Hamilton et al. 2017)	NeurIPS	GraphSAGE	Inductive graph representation	G2	RF-GraphSAGE baseline
(Rossi et al. 2020)	ICML Workshop	Temporal graph networks	Dynamic graph deep learning	G2	Metabolic time-warp kernel
(Balmaseda et al. 2023)	Intell. Syst. Appl.	Deep graph learning	GNN systemic risk prediction	G2	GNN benchmark for finance
(Fan et al. 2024)	IEEE TPAMI	GNN out-of-distribution	Distribution shift in graphs	G2	Identifies OOD failure mode
(Febrinanto et al. 2023)	IEEE Comp. Intell.	Graph lifelong learning	Catastrophic forgetting in GNNs	G2	Unconstrained hypothesis space
Generation 3: Hypergraph and higher-order methods					
(Benson et al. 2016)	Science	Higher-order organisation	k -adic interaction dynamics	G3	Hyperedge formation rationale
(Landry and Restrepo 2020)	Chaos	Hypergraph contagion	Heterogeneity in spreading	G3	Hypergraph Laplacian model
(Aksoy et al. 2020)	EPJ Data Science	High-order hypergraph walks	Hypernetwork science framework	G3	Walk-based hyperedge encoding
(Cui et al. 2023)	Pattern Recognition	Temporal hypergraph attention	Stock trend prediction	G3	HTHGN baseline architecture
(Han et al. 2023)	Applied Intelligence	Industry hypergraph attention	Stock prediction via sectors	G3	Sector-aware hyperedge baseline
(Liu et al. 2025)	IJCAI 2025	Heterogeneous temporal hypergraph	Temporal hypergraph NNs	G3	RF-HTHGN baseline
(David et al. 2024)	Entropy (MDPI)	Functional hypergraphs	VNE + Ricci curvature in stocks	G3	Entropy signals in hypergraphs
(Zhang et al. 2023)	Nature Commun.	Higher-order collective dynamics	Hypergraph $\neq$ simplicial complex	G3	Contagion under infra. stress
(Kamiński et al. 2024)	J. Complex Netw.	Hypergraph community detection	h-Louvain modularity	G3	Structural prior without physics
(Cvetković and Mitrović Dankulov 2026)	J. Complex Netw.	GNN explainability	Black-box failure in networks	G3	Motivates GB-IG attribution

Table A1. Cont.

Study	Venue	Method	Key finding	Str.	Relevance
Entropy and statistical physics (Bianconi 2009)	Phys. Rev. E	Network ensemble entropy	VNE of graph Laplacians	Ph	$S(t)$ derivation basis
(Minello et al. 2019)	J. Complex Netw.	VNE decomposition	Node/edge entropy	Ph	GB-IG theoretical foundation
(Ye et al. 2018)	J. Complex Netw.	Entropy component analysis	Structural role identification	Ph	Degree vs. entropy centrality
(Majhi et al. 2022)	J. R. Soc. Interface	Higher-order dynamics review	k -adic propagation distinct	Ph	Hyperedge dynamics grounding
(Billio et al. 2016)	SSRN WP 2604754	Entropy early warning indicator	Entropy predicts banking crises	En	Entropy threshold precedent
Physics-informed neural networks (Raissi et al. 2019)	J. Comp. Physics	Physics-informed NNs	Physical law as loss constraint	PN	PINN architecture paradigm
(Kissas et al. 2020)	Comp. Meth. Appl.	PINN cardiovascular flow	OOD accuracy via physics	PN	OOD generalisation via priors
(Luo et al. 2025)	Applied Intelligence	Enhanced hybrid PINN	Adaptive PDE constraint	PN	Soft vs. hard constraint choice
Optimal control and multilayer networks (Fleming and Soner 2006)	Springer	HJB controlled processes	Stochastic optimal control	Ct	HJB governance protocol basis
(Brunnermeier and Sannikov 2014)	American Econ. Rev.	Macro model w/ finance	Stability boundary targeting	Ct	Lyapunov stability motivation
(Kivelä et al. 2014)	J. Complex Netw.	Multilayer networks	Foundational multilayer theory	ML	Financial physical layer coupling
(Aleta et al. 2026)	J. Complex Netw.	Multilayer network science	Theory, methods, applications	ML	Infra.-finance interdependence
(Busby et al. 2021)	Energy Res. Soc. Sci.	Texas winter storm 2021	Infra. cascades to finance	In	Empirical infra. coupling case

Notes

- ¹ In accordance with (Bianconi 2009; Minello et al. 2019), we deploy Von Neumann entropy as defined for network Laplacians, where the normalized Laplacian serves as the density matrix. Despite having the same mathematical structure, this is different from the quantum-mechanical density matrix formulation.

References

- Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2015. Systemic risk and stability in financial networks. *American Economic Review* 105: 564–608. <https://doi.org/10.1257/aer.20130456>.
- Aksoy, Sinan G., Cliff Joslyn, Carlos Ortiz Marrero, Brenda Praggastis, and Emilie Purvine. 2020. Hypernetwork science via high-order hypergraph walks. *EPJ Data Science* 9: 16. <https://doi.org/10.1140/epjds/s13688-020-00231-0>.
- Aleta, Alberto, Andreia Sofia Teixeira, Guilherme Ferraz de Arruda, Andrea Baronchelli, Alain Barrat, János Kertész, Albert Díaz-Guilera, Oriol Artime, Michele Starnini, Giovanni Petri, and et al. 2026. Multilayer network science: Theory, methods, and applications. *Journal of Complex Networks* 14: cnag007. <https://doi.org/10.1093/comnet/cnag007>.
- Allen, Franklin, and Douglas Gale. 2000. Financial contagion. *Journal of Political Economy* 108: 1–33. <https://doi.org/10.1086/262109>.
- Amos, Brandon, and J. Zico Kolter. 2017. OptNet: Differentiable optimization as a layer in neural networks. Paper presented at 34th International Conference on Machine Learning, Sydney, Australia, August 6–11. Available online: <https://proceedings.mlr.press/v70/amos17a.html> (accessed on 15 May 2026).
- Balmaseda, Vicente, María Coronado, and Gonzalo de Cadenas-Santiago. 2023. Predicting systemic risk in financial systems using deep graph learning. *Intelligent Systems with Applications* 19: 200240. <https://doi.org/10.1016/j.iswa.2023.200240>.
- Battiston, Stefano, Antoine Mandel, and Irene Monasterolo. 2017. A climate stress-test of the financial system. *Nature Climate Change* 7: 283–88. <https://doi.org/10.1038/nclimate3255>.
- Battiston, Stefano, J. Doyne Farmer, Andreas Flache, Diego Garlaschelli, Andrew G. Haldane, Hans Heesterbeek, Cars Hommes, Carlo Jaeger, Robert May, and Marten Scheffer. 2016. Complexity theory and financial regulation. *Science* 351: 818–19. <https://doi.org/10.1126/science.aad0299>.

- Battiston, Stefano, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli. 2012. DebtRank: Too central to fail? Financial networks, the FED and systemic risk. *Scientific Reports* 2: 541. <https://doi.org/10.1038/srep00541>.
- Benson, Austin R., David F. Gleich, and Jure Leskovec. 2016. Higher-order organization of complex networks. *Science* 353: 163–66. <https://doi.org/10.1126/science.aad9029>.
- Bianconi, Ginestra. 2009. Entropy of network ensembles. *Physical Review E* 79: 036114. <https://doi.org/10.1103/PhysRevE.79.036114>.
- Billio, Monica, Roberto Casarin, Michele Costola, and Andrea Pasqualini. 2016. An entropy-based early warning indicator for systemic risk. *Journal of International Financial Markets, Institutions and Money* 45: 42–59. <https://doi.org/10.1016/j.intfin.2016.05.008>.
- Brunnermeier, Markus K., and Yuliy Sannikov. 2014. A macroeconomic model with a financial sector. *American Economic Review* 104: 379–421. <https://doi.org/10.1257/aer.104.2.379>.
- Brunton, Steven L., and J. Nathan Kutz. 2022. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*, 2nd ed. Cambridge: Cambridge University Press.
- Busby, Joshua W., Kyri Baker, Morgan D. Bazilian, Alex Q. Gilbert, Emily Grubert, Varun Rai, Joshua D. Rhodes, Sarang Shidore, Caitlin A. Smith, and Michael E. Webber. 2021. Cascading risks: Understanding the 2021 Texas winter storm. *Energy Research & Social Science* 77: 102106. <https://doi.org/10.1016/j.erss.2021.102106>.
- Cai, Shengze, Zhiping Mao, Zhicheng Wang, Minglang Yin, and George Em Karniadakis. 2021. Physics-informed neural networks (PINNs) for fluid mechanics: A review. *Acta Mechanica Sinica* 37: 1727–38. <https://doi.org/10.1007/s10409-021-01148-1>.
- Cohen, Jacob. 1988. *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed. Hillsdale: Lawrence Erlbaum Associates.
- Cont, Rama, Amal Moussa, and Edson Bastos e Santos. 2013. Network structure and systemic risk in banking systems. In *Handbook on Systemic Risk*. Edited by Jean-Pierre Fouque and Joseph A. Langsam. Cambridge: Cambridge University Press, pp. 327–68.
- Cui, Chaoran, Xiaojie Li, Chunyun Zhang, Weili Guan, and Meng Wang. 2023. Temporal-relational hypergraph tri-attention networks for stock trend prediction. *Pattern Recognition* 143: 109759. <https://doi.org/10.1016/j.patcog.2023.109759>.
- Cvetković, Darja, and Marija Mitrović Dankulov. 2026. Assessing the explainability of Graph Neural Networks in random graphs classification task. *Journal of Complex Networks* 14: cnag006. <https://doi.org/10.1093/comnet/cnag006>.
- David, Jerry Jones, Narayan G. Sabhahit, Sebastiano Stramaglia, Tiziana Di Matteo, Stefano Boccaletti, and Sarika Jalan. 2024. Functional hypergraphs of stock markets. *Entropy* 26: 848. <https://doi.org/10.3390/e26100848>.
- Efron, Bradley, and Robert J. Tibshirani. 1993. *An Introduction to the Bootstrap*. New York: Chapman & Hall.
- Eisenberg, Larry, and Thomas H. Noe. 2001. Systemic risk in financial systems. *Management Science* 47: 236–49. <https://doi.org/10.1287/mnsc.47.2.236.9835>.
- Eskom Holdings SOC Ltd. 2025. *Integrated Report 2025*. Johannesburg: Eskom. Available online: <https://www.eskom.co.za/wp-content/uploads/2025/10/Eskom-integrated-report-2025.pdf> (accessed on 15 May 2026).
- Fan, Shaohua, Xiao Wang, Chuan Shi, Peng Cui, and Bai Wang. 2024. Generalizing graph neural networks on out-of-distribution graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46: 322–37. <https://doi.org/10.1109/TPAMI.2023.3321097>.
- Febrinanto, Fali Gozi, Feng Xia, Kristen Moore, Chandra Thapa, and Charu Aggarwal. 2023. Graph lifelong learning: A survey. *IEEE Computational Intelligence Magazine* 18: 26–51. <https://doi.org/10.1109/MCI.2022.3222049>.
- Fleming, Wendell H., and H. Mete Soner. 2006. *Controlled Markov Processes and Viscosity Solutions*, 2nd ed. New York: Springer.
- Gai, Prasanna, and Sujit Kapadia. 2010. Contagion in financial networks. *Proceedings of the Royal Society A* 466: 2401–23. <https://doi.org/10.1098/rspa.2009.0410>.
- Hamilton, Will, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems* 30: 1024–34.
- Han, Haodong, Liang Xie, Shengshuang Chen, and Haijiao Xu. 2023. Stock trend prediction based on industry relationships driven hypergraph attention networks. *Applied Intelligence* 53: 29448–64. <https://doi.org/10.1007/s10489-023-05035-z>.
- Havil, Julian. 2003. *Gamma: Exploring Euler's Constant*. Princeton: Princeton University Press.
- Hofmeijer, Jeannette, and Michel Johannes Antonius Maria van Putten. 2016. EEG in postanoxic coma: Prognostic and diagnostic value. *Clinical Neurophysiology* 127: 2047–55. <https://doi.org/10.1016/j.clinph.2016.02.002>.
- Kamiński, Bogumił, Paweł Misiorek, Paweł Prałat, and François Théberge. 2024. Modularity based community detection in hypergraphs. *Journal of Complex Networks* 12: cnae041. <https://doi.org/10.1093/comnet/cnae041>.
- Kipf, Thomas N., and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. Paper presented at International Conference on Learning Representations, Toulon, France, April 24–26. Available online: <https://openreview.net/forum?id=SJU4ayYgl> (accessed on 15 May 2026).
- Kissas, Georgios, Yibo Yang, Eileen Hwuang, Walter R. Witschey, John A. Detre, and Paris Perdikaris. 2020. Machine learning in cardiovascular flows modeling. *Computer Methods in Applied Mechanics and Engineering* 358: 112623. <https://doi.org/10.1016/j.cma.2019.112623>.
- Kivelä, Mikko, Alex Arenas, Marc Barthélemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. 2014. Multilayer networks. *Journal of Complex Networks* 2: 203–71. <https://doi.org/10.1093/comnet/cnu016>.

- Landauer, Rolf. 1961. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development* 5: 183–91. <https://doi.org/10.1147/rd.53.0183>.
- Landry, Nicholas W., and Juan G. Restrepo. 2020. The effect of heterogeneity on hypergraph contagion models. *Chaos* 30: 103117. <https://doi.org/10.1063/5.0020034>.
- LaSalle, Joseph P., and Solomon Lefschetz. 1961. *Stability by Liapunov's Direct Method with Applications*. New York: Academic Press.
- Lee, Edward A. 2008. Cyber physical systems: Design challenges. Paper presented at 11th IEEE Symposium on Object Oriented Real-Time Distributed Computing, Orlando, FL, USA, May 5–7. <https://doi.org/10.1109/ISORC.2008.25>.
- Liu, Huan, Pengfei Jiao, Mengzhou Gao, Chaochao Chen, and Di Jin. 2025. Heterogeneous temporal hypergraph neural network. Paper presented at Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-25), Montreal, QC, Canada, August 16; p. 347. <https://doi.org/10.24963/ijcai.2025/347>.
- Luo, Kuang, Shaolin Liao, Zhong Guan, and Baiquan Liu. 2025. An enhanced hybrid adaptive physics-informed neural network for forward and inverse PDE problems. *Applied Intelligence* 55: 255. <https://doi.org/10.1007/s10489-024-06195-2>.
- Majhi, Soumen, Matjaž Perc, and Dibakar Ghosh. 2022. Dynamics on higher-order networks: A review. *Journal of the Royal Society Interface* 19: 20220043. <https://doi.org/10.1098/rsif.2022.0043>.
- May, Robert M. 1972. Will a large complex system be stable? *Nature* 238: 413–14. <https://doi.org/10.1038/238413a0>.
- Merton, Robert C. 1971. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3: 373–413. [https://doi.org/10.1016/0022-0531\(71\)90038-X](https://doi.org/10.1016/0022-0531(71)90038-X).
- Minello, Giorgia, Luca Rossi, and Andrea Torsello. 2019. On the Von Neumann entropy of graphs. *Journal of Complex Networks* 7: 491–514. <https://doi.org/10.1093/comnet/cny028>.
- Moroke, Ntebogang Dinah. 2026. Ethical Guardrail Bandit Pruning (EGBP): A Framework for Sustainable and Equitable AI-Driven Resource Allocation in Bioenergy Grids. SSRN. Available online: <https://doi.org/10.2139/ssrn.6685511> (accessed on 15 May 2026).
- Moroke, Ntebogang Dinah. 2026. IST02-IST03 JSE-Eskom Infrastructure-Coupled Financial Network Dataset [Data Set]. Zenodo Version 2. Deposited. Available online: <https://zenodo.org/records/20008530> (accessed on 15 May 2026).
- Raissi, Maziar, Paris Perdikaris, and George Em Karniadakis. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378: 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>.
- Rossi, Emanuele, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael M. Bronstein. 2020. Temporal graph networks for deep learning on dynamic graphs. *arXiv*, arXiv:2006.10637.
- South African Reserve Bank. 2025. *Financial Stability Review*, 1st ed. Pretoria: SARB.
- Stauffer, Dietrich, and Amnon Aharony. 1994. *Introduction to Percolation Theory*, 2nd ed. London: Taylor & Francis.
- Sundararajan, Mukund, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. Paper presented at 34th International Conference on Machine Learning, Sydney, Australia, August 6–11. Available online: <https://proceedings.mlr.press/v70/sundararajan17a.html> (accessed on 15 May 2026).
- Ulanowicz, Robert E., Sally J. Goerner, Bernard Lietaer, and Rocio Gomez. 2009. Quantifying sustainability: Resilience, efficiency and the return of information theory. *Ecological Complexity* 6: 27–36. <https://doi.org/10.1016/j.ecocom.2008.10.005>.
- Veličković, Petar, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. Paper presented at International Conference on Learning Representations, Vancouver, BC, Canada, April 30–May 3. Available online: <https://openreview.net/forum?id=rJXmpikCZ> (accessed on 15 May 2026).
- Ye, Cheng, Richard C. Wilson, and Edwin R. Hancock. 2018. Network analysis using entropy component analysis. *Journal of Complex Networks* 6: 893–915. <https://doi.org/10.1093/comnet/cny003>.
- Zhang, Yuanzhao, Maxime Lucas, and Federico Battiston. 2023. Higher-order interactions shape collective dynamics differently in hypergraphs and simplicial complexes. *Nature Communications* 14: 1605. <https://doi.org/10.1038/s41467-023-37190-9>.
- Zhou, Dengyong, Jiayuan Huang, and Bernhard Schölkopf. 2006. Learning with hypergraphs: Clustering, classification, and embedding. In *Advances in Neural Information Processing Systems* 19. Cambridge, MA: MIT Press.

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.