



When Harmless Nodes Become Systemic Chokepoints: Joint Chokepoint Amplification in Complex Systems

Raymond R. Tjandrawinata

School of Bioscience, Technology and Innovation Atma Jaya Catholic University of

Indonesia Jakarta, Indonesia

Email: raytjan@yahoo.com

Keywords:

complex systems, systemic risk, chokepoints, cascading failure, network resilience.

Abstract

Complex systems increasingly face systemic risks arising not only from individual component failures but also from interactions among multiple components that can generate nonlinear cascading effects. Conventional risk assessment approaches often rely on node-level indicators, such as centrality measures and individual impact assessments, which may overlook hidden vulnerabilities emerging from combinations of seemingly insignificant components. This study aimed to develop and characterize the concept of Joint Chokepoint Amplification (JCA) as a framework for identifying latent systemic vulnerabilities generated through interacting failures within complex networks. The research employed a qualitative conceptual design combined with structured literature analysis and controlled computational experiments. Secondary data from studies on network science, systemic risk, supply chain disruptions, and higher-order interactions were synthesized to formulate the JCA framework, while simulations using Erdős–Rényi, Barabási–Albert, and modular stochastic block model networks were conducted to examine amplification patterns. The results indicated that joint failures among individually low-impact nodes could produce disproportionately large systemic consequences, with amplification effects strongly influenced by network topology and system conditions. Modular networks demonstrated a higher concentration of latent joint chokepoints compared with other network structures. The findings revealed that conventional node-based assessments may fail to capture interaction-driven risks. This study concludes that systemic resilience requires evaluating not only critical nodes but also vulnerable configurations of interacting components, providing a foundation for developing more adaptive risk management strategies in complex systems.



This is an open access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.

Copyright © 2026 by Author. Published by Publikasiku Publisher

INTRODUCTION

Many systems can withstand the failure of component A or component B individually; however, when both components fail simultaneously, the system may collapse (Makoond et al., 2024; Maldonado et al., 2024; Xu et al., 2025). This phenomenon becomes understandable once nonlinear dependencies are considered. Nevertheless, a significant proportion of operational risk analysis still evaluates systems one component at a time by ranking critical nodes, identifying hubs, protecting major suppliers, securing high-throughput infrastructure, isolating large financial institutions, or prioritizing actors with the highest centrality. This approach is understandable because node-level importance is measurable. However, the ease of measurement may inadvertently lead to an assumption that measurable characteristics fully

represent systemic importance, overlooking relational vulnerabilities that emerge from interactions among components.

Network science has repeatedly highlighted the limitations of such simplification. Scale-free networks may demonstrate tolerance to random failures while remaining highly vulnerable to targeted attacks (Bi et al., 2025; Fu et al., 2024; Liu et al., 2026; Park, 2026; Yin et al., 2023). Load redistribution can transform the removal of a single node into an overload cascade. Interdependent networks can experience cascading failures across multiple layers even when individual layers appear robust in isolation. Financial networks may transition from stabilizing diversification to destabilizing propagation when shocks exceed critical thresholds. Production networks can transmit localized shocks through input specificity, downstream dependencies, and indirect exposures (Barrot & Sauvagnat, 2016; Carvalho et al., 2021). These findings indicate that systemic vulnerability is not solely a property of individual nodes but also a property of nodes operating within specific configurations and dynamic conditions.

Recent evidence has further emphasized this limitation. Fialkowski, Havlin, and Thurner (2026) reconstructed a national firm-level supply network and demonstrated that simultaneous firm failures could generate damage up to 257 times greater than the combined systemic-risk contributions of the individual firms. Only a small proportion of random pairs exhibited strong amplification effects, highlighting why conventional screening methods may overlook critical risks: dangerous combinations are relatively rare but can generate extreme consequences. Their amplification factor represents the ratio between joint impact and the sum of individual impacts. This study builds upon that empirical insight by examining whether the same mathematical principle can be generalized beyond supply-chain systems as a broader characteristic of complex systems.

This study introduces the concept of Joint Chokepoint Amplification (JCA). A chokepoint is commonly understood as a location, actor, or connection through which critical flows must pass. This study extends the concept by defining a joint chokepoint as a set of components whose systemic importance emerges through interaction rather than individual dominance. None of the components within the set needs to be critical independently; however, the combination becomes critical when simultaneous perturbations reduce redundancy, cross system thresholds, remove multiple network bridges, weaken shared buffers, synchronize losses, or activate feedback mechanisms. Therefore, the primary analytical focus shifts from individual nodes to interaction effects within the system impact function.

Despite substantial advances in network science, existing approaches remain limited in identifying latent systemic vulnerabilities. Conventional centrality measures, including degree centrality and betweenness centrality, are useful for identifying influential nodes but are primarily designed to evaluate individual importance rather than interaction-driven criticality. Consequently, systems may contain hidden failure configurations in which individual components have limited standalone impact, but their combined failure produces severe systemic consequences. Existing literature has not yet established a comprehensive framework for systematically defining, measuring, and identifying interaction-driven chokepoints across different domains. This conceptual limitation creates a gap between traditional node-based risk assessment and the growing recognition that systemic failures frequently emerge from nonlinear combinations of events.

The absence of such a framework creates practical challenges for organizations and

policymakers responsible for strengthening systemic resilience. In supply chains, organizations may overestimate resilience because multiple suppliers or alternative pathways appear available, even though these alternatives may depend on shared resources, geographic regions, infrastructure, or operational mechanisms. In financial systems, institutions may appear individually stable while collectively generating vulnerabilities through shared liquidity channels, collateral dependencies, or correlated exposures. Similar patterns can occur in transportation networks, healthcare systems, digital platforms, and organizational decision structures, where failures may emerge from interactions among multiple seemingly ordinary components rather than from a single dominant actor. Therefore, developing approaches capable of detecting latent joint vulnerabilities is increasingly important for improving risk preparedness and resilience strategies.

To address this research gap, this study proposes Joint Chokepoint Amplification (JCA) as a theoretical framework for identifying systemic vulnerabilities generated through combinations of interacting components. Unlike conventional chokepoint analysis, which focuses on highly influential nodes or connections, JCA conceptualizes criticality as a relational property that emerges when multiple components interact under specific system conditions. The framework distinguishes between visible criticality, which can be identified through conventional network metrics, and latent criticality, which emerges only when multiple individually low-impact components fail simultaneously. Through this perspective, the study argues that some of the most significant vulnerabilities in complex systems may exist among ordinary-looking components rather than clearly dominant nodes.

The novelty of this research lies in reframing systemic chokepoints from a node-based phenomenon into an interaction-based phenomenon. This study contributes by defining JCA as a set-function diagnostic framework, introducing the concept of latent joint chokepoints, and demonstrating that conventional centrality measures may fail to identify high-risk configurations in nonlinear systems. Furthermore, the study develops conceptual representations, including pairwise JCA matrices, higher-order JCA spectra, and chokepoint hypergraphs, to support the identification of hidden systemic risks. By integrating perspectives from network science, systemic risk theory, and higher-order interaction research, this study extends resilience analysis beyond individual components toward structural relationships among interconnected components.

The objective of this research was to develop and characterize a theoretical framework for understanding how combinations of individually low-impact components can generate disproportionate systemic consequences. Specifically, this study aimed to examine the mechanisms through which joint failures produce amplification effects, analyze the influence of network topology on the emergence of latent joint chokepoints, and demonstrate the limitations of conventional node-level risk assessment approaches. Through computational experiments involving different network structures, including Erdős–Rényi, Barabási–Albert, and modular stochastic block model networks, the study investigated how interaction-based vulnerabilities emerged under different system conditions.

The findings of this study provide theoretical and practical contributions to complex systems research and risk management. Theoretically, this research expands existing knowledge by demonstrating that systemic criticality can emerge from interactions among components rather than from individual dominance alone, contributing to the development of

higher-order network analysis. Practically, the JCA framework provides organizations, policymakers, and risk managers with a complementary approach for identifying hidden vulnerabilities that may remain undetected through conventional assessments. By recognizing that resilience depends not only on eliminating single points of failure but also on understanding dangerous combinations of dependencies, this research provides a foundation for developing more adaptive, robust, and strategically informed risk mitigation strategies across diverse complex systems.

METHOD

Research Design and Approach

This study employed a qualitative-conceptual research design that combined theoretical development through structured literature analysis (library research) with controlled computational experiments to examine and illustrate the proposed framework. The primary approach was analytical and interpretive rather than statistical-inferential, focusing on developing, defining, and characterizing the theoretical construct of Joint Chokepoint Amplification (JCA) through synthesis and critical analysis of existing literature on network science, systemic risk, and higher-order interactions.

Data Sources

The primary data sources are secondary and documentary: peer-reviewed journal articles, working papers, and preprints in complex-systems science, network theory, financial-network economics, supply-chain risk, and higher-order interaction research, as listed in the References. These sources were purposively selected for their theoretical and empirical relevance to node centrality, cascading failure, systemic-risk amplification, and higher-order network interactions. Synthetic network data generated through controlled simulation (Erdos-Renyi, Barabasi-Albert, and modular stochastic-block-model realizations, detailed in Section 4) serve as a secondary, illustrative data source used to probe the theoretical claims developed in Section 3.

Data Collection Technique

Data collection followed a structured literature review technique: relevant theoretical and empirical works were identified, read, and thematically organized around the core constructs of node centrality, cascading failure, higher-order interactions, and systemic-risk amplification. Key definitions, findings, and arguments were extracted and compared across sources to identify a conceptual gap, namely the absence of a general, cross-domain framework for interaction-generated systemic risk. This documentary analysis was complemented by a controlled computational experiment, in which synthetic network scenarios were generated and perturbed to observe the emergence of the theoretical construct under study.

Data Analysis Technique

Data were analyzed using conceptual-analytical techniques characteristic of theory-building qualitative research: comparison, synthesis, and critical evaluation of existing constructs, such as degree and betweenness centrality, collective influence, and higher-order interaction models, against the phenomenon of joint, interaction-generated systemic risk. This analysis was used to formally define Joint Chokepoint Amplification and the latent joint chokepoint, and to derive their theoretical implications, including the centrality-blind-spot result presented in Section 3.4. The resulting framework was then examined descriptively

through the computational experiment, whose design, parameters, and interpretive criteria are detailed in Section 4.

RESULTS AND DISCUSSION

Amplification is not exceptional; extreme amplification is

Across all three network families, average pairwise JCA was only modestly above one. That fact is important. If every pair were dramatically superadditive, the framework would be uninformative because 'everything interacts with everything.' Instead, the simulations reproduce the qualitative signature found in empirical supply-chain work: a large near-additive bulk and a thin, dangerous upper tail.

The tail differed sharply by topology. Across the 20 baseline realizations, the mean of each network's maximum JCA was 12.29 in ER, 10.57 in BA, and 14.45 in modular networks. The 99th percentile of pairwise JCA averaged 3.55, 2.62, and 7.02 respectively. The modular networks were especially distinctive not because every pair was dangerous, but because their upper tail was persistently heavier. All 20 modular realizations contained at least one latent joint chokepoint.

Table 3. Baseline JCA results across 20 realizations per topology

Topology	Mean JCA	P99 JCA	Max mean	JCA LJC (%)	share Runs with LJC	rho degree	rho overlap
Erdos-Renyi	1.144	3.55	12.29	0.514	16/20	0.21	0.44
Barabasi-Albert	1.085	2.62	10.57	0.164	11/20	0.20	0.20
Modular SBM	1.285	7.02	14.45	3.739	20/20	0.16	0.37

LJC share is the fraction of pairs satisfying $\epsilon \leq 0.05$ for each member, joint impact ≥ 0.20 , and $JCA \geq 2$. Correlations are mean within-network Spearman coefficients.




Figure 2. Mean per-network maximum JCA across 20 baseline realizations. Error bars show one standard deviation. Extreme amplification appears in every topology, but modular networks sustain the strongest upper-tail pattern.

The most dangerous pairs can look ordinary one at a time

The strongest single baseline example is intentionally difficult to forget. In one ER realization, two nodes each produced only one failed node when perturbed separately: $I(a) = I(b) = 1/60 = 0.0167$. Perturbed together, they triggered complete collapse: $I(\{a,b\}) = 1$. The resulting JCA was 30. Neither singleton gave any warning in its impact score. The interaction term contained almost the entire story.

This is the operational meaning of latent criticality. If a risk committee had ranked the 60 nodes solely by singleton damage and stopped after the obvious top tier, these two nodes could have been dismissed as noise. Yet they formed a perfect two-key switch for system-wide failure. The example is synthetic, but the mechanism is not. Fialkowski et al. (2026) observe the empirical analogue in supply chains: combinations of individually low-ESRI firms can eliminate substitution and trigger near-total losses.

The title's 'harmless nodes' is therefore provocative but precise under the declared ε rule. The nodes are not metaphysically harmless. They are harmless under the diagnostic that most organizations actually use: test one thing at a time, measure standalone consequence, and rank. JCA reveals where that diagnostic changes the answer.

Modularity concentrates latent joint chokepoints

The largest topology contrast appeared in the density of latent pairs. At the baseline threshold, latent pairs represented about 3.74% of individually harmless pairs in modular networks, compared with about 0.51% in ER and 0.16% in BA networks. All 20 modular realizations contained latent joint chokepoints; 16 of 20 ER and 11 of 20 BA realizations did so.

Why should modularity matter? A modular network creates local neighborhoods in which several nodes share downstream neighborhoods, bridges, or threshold exposure. One failure may be absorbed inside the module. Two failures can cross a local threshold simultaneously and then export the cascade through the sparse inter-module structure. In other words, modularity creates pockets of redundancy and pockets of conjunctive fragility at the same time. The same architecture that contains one shock can synchronize two.

This result should not be generalized mechanically to all modular systems. Modularity can also protect systems by compartmentalizing contagion. The point is conditional: when local threshold rules create conjunctive activation, modules can harbor many more dangerous combinations than singleton statistics imply.

Centrality sees part of the risk, not the interaction term

Pair-averaged degree centrality had only modest association with JCA: mean Spearman ρ was 0.21 in ER, 0.20 in BA, and 0.16 in modular networks. Betweenness was similarly modest. Neighborhood overlap, by contrast, reached $\rho = 0.44$ in ER and 0.37 in modular networks, although it was weaker in BA networks ($\rho = 0.20$). In the modular family, only about one third of the top 1% JCA pairs were captured by screening the top 10% of pairs ranked by mean degree centrality.

This is not a contest in which overlap 'beats' centrality universally. BA networks remind us why: hubs genuinely matter in preferential-attachment structures, and degree screening captured a larger share of extreme pairs there. The result is subtler. When amplification depends on shared neighborhoods or joint threshold exposure, pair structure becomes information that no average of two node scores can fully reconstruct.

Argument about flow assumptions has an analogue here. Centrality measures encode a process model, usually at the node level. JCA exposes when the relevant process is conjunctive. A centrality map may still identify where risk travels, while a JCA map identifies which perturbations must occur together to open the route.




Figure 3. Mean Spearman association between pairwise JCA and pair descriptors. Neighborhood overlap carries more information than degree or betweenness in ER and modular networks, but not uniformly across topologies.

Distance from threshold is a system condition, not a footnote

The sensitivity analysis produced the most important system-level result. Lowering the mean node threshold from 0.38 toward 0.30 sharply increased latent joint chokepoints, especially in ER and modular networks. At mean threshold 0.30, modular networks averaged more than 12% latent pairs among the individually harmless pair set in the six sensitivity realizations; at 0.38 the corresponding share fell below 1%. The upper tail of maximum JCA moved in the same direction.

This means that joint chokepointness is not a permanent label attached to a node pair. It is a conditional property of a pair under a system state. A pair can be innocuous when buffers are thick and dangerous when thresholds tighten. In supply chains, inventory and spare supplier capacity can move the threshold. In finance, capitalization and liquidity do. In organizations, staffing and decision slack do. The same structural pair can migrate in and out of latent criticality without any edge in the network changing.

That observation changes how JCA should be used. It is not enough to calculate a static 'dangerous pairs' list once a year. The relevant object is $JCA(S | C)$, where C denotes system condition: demand pressure, liquidity, spare capacity, substitutability, geopolitical constraint, information load, or another state variable that changes propagation.




Figure 4. Sensitivity of latent joint chokepoints to the mean threshold. Lower thresholds place the system closer to cascade susceptibility, making interaction-generated vulnerability much denser.

The Redundancy Paradox: Why Backup Can Hide Joint Failure

Redundancy is usually treated as the opposite of chokepoint risk. Often it is. A second supplier, second route, second approval path, or second data source reduces dependence on the first. But redundancy can also make a dangerous configuration invisible because each component looks safe precisely while its substitute remains available. The risk is not removed; it is transferred from the node to the set.

Consider two suppliers, each individually replaceable. A conventional audit records low single-source risk. If both draw on the same upstream input, port, certification body, cloud region, financing channel, or labor pool, the backups are not independent. The apparent redundancy is conditional redundancy. When the common condition fails, the system discovers that it had two names for one dependency. Fialkowski et al. (2026) give an empirical form of this mechanism: replaceability breaks down when simultaneous failures remove enough of the supplying industry.

The paradox becomes sharper in organizational systems. Two executives may each have a deputy, but both deputies may depend on the same legal interpretation. Two laboratories may be geographically separate but share the same analytical standard. Two digital platforms may run on separate applications but the same identity provider. Two committees may have separate mandates but both require the same final signature. A resilience map that counts backups without mapping shared fallback mechanisms can overstate resilience.

JCA reframes redundancy as a question about independence of compensation. A backup is not robust because it exists. It is robust to the extent that the mechanisms sustaining it fail independently of the mechanisms sustaining the primary path. This yields a practical search heuristic: high-JCA candidates are often pairs or sets that share a substitution pool, downstream gate, buffer, or recovery mechanism.

Cross-Domain Implications

1. Supply chains and production networks

Production networks are the most immediate application because the empirical phenomenon has already been demonstrated. The network-origins-of-aggregate-fluctuations literature established how supplier position and input-output asymmetry transmit local shocks to aggregate outcomes, a line of work recently extended to adaptive, reshaping supply networks (König et al., 2022) and synthesized into a general micro-to-macro propagation framework (Baqae & Rubbo, 2023). Natural-disaster evidence shows that input specificity magnifies propagation (Barrot & Sauvagnat, 2016) and that disruptions can move upstream and downstream through firm networks (Carvalho et al., 2021). JCA adds a compound-shock lens: which supplier combinations jointly exhaust substitution?

This suggests that supplier-risk programs should stop at neither spend concentration nor single-source flags. They should test pair and set failures conditioned on shared product classes, geographic hazards, logistics routes, certification dependencies, and substitute capacity. A pair of Tier-2 suppliers with modest revenue may be more systemically important together than the largest Tier-1 supplier alone. The obvious supplier can be visible risk. The dangerous combination can be hidden risk.

2. Financial systems

Financial-network research already emphasizes that connectivity can diversify and propagate risk at the same time (Amelkin et al., 2024; Glasserman & Young, 2016). Systemwide stress testing increasingly integrates banks, nonbank financial intermediaries, margin calls, fire sales, market impact, and cross-sector liquidity channels (Oura et al., 2026). The natural JCA question is not merely which institution is systemically important, but which combination breaches a market-wide compensatory mechanism.

Candidate joint chokepoints include institutions sharing collateral, crowded portfolios, liquidity providers, clearing arrangements, or margin dynamics. A pair can be individually absorbable but jointly force fire sales into an illiquid market. This is exactly the kind of configuration for which local resilience can coexist with system fragility. JCA would not replace institution-level systemic-risk scores; it would expose interaction risk left in their residual.

3. Infrastructure and maritime chokepoints

The word chokepoint is literal in transport infrastructure. Recent work quantifies the systemic impact of maritime chokepoint disruptions and explicitly calls attention to joint disruptions and co-occurring hazards (Verschuur et al., 2025). Agent-based work on Suez and Panama disruptions likewise shows that timing and rerouting behavior shape recovery (Wen et al., 2026). JCA gives this literature an additional question: does losing two routes together impose merely the sum of their individual delay costs, or does the joint loss eliminate rerouting and create nonlinear congestion?

The answer may depend on time. Two disruptions separated by six months may be additive; the same disruptions overlapping for ten days may be superadditive. A mature JCA model should therefore eventually become temporal: $JCA(S, t, \Delta t)$, with jointness defined not only by membership but by overlap in shock duration and recovery state.

4. Organizations, governance, and AI systems

Organizations are networks of decisions, expertise, permissions, routines, and legitimacy. Their formal organization charts reveal only part of the flow. Polishchuk and Polishchuk (2026) make a parallel point in network vulnerability: flow cores can be more informative

than topology alone because what actually moves through the network matters. In organizations, the relevant flow may be approval, information, budget authority, technical sign-off, or political legitimacy.

A joint organizational chokepoint can therefore involve two units that are individually replaceable but jointly necessary for a critical workflow. The same logic applies to AI-agent societies. Lu et al. (2026) show that collective behavior can be bottlenecked by interaction topology even when individual models are capable. In multi-agent architectures, a JCA-style stress test could perturb combinations of agents, tools, memory channels, or routing nodes to reveal configurations whose joint loss destroys collective performance. The system can be rich in agents and poor in independent pathways.

Adaptive response complicates the picture. Liu et al. (2026) show that interventions can change network structure through rewiring and even backfire above certain regimes. A static JCA estimate should therefore be treated as conditional on the response rule. After a joint failure, actors may reroute, substitute, hoard, coordinate, or defect. The system that is being measured is partly created by the shock itself.

5. An Operational Search Strategy

The theoretical problem becomes combinatorial quickly. A system with n nodes has $n(n-1)/2$ pairs and $C(n,k)$ sets of size k . Exhaustive enumeration is feasible for dozens of nodes and impossible for national-scale firm networks. Fialkowski et al. (2026) address the same explosion through a large-set perturbation and minimal-subset extraction strategy. The broader JCA framework suggests a layered search process rather than brute force.

- 1) Stage 1 - Define the impact function. Choose a consequence that actually matters operationally: lost output, unserved demand, failed service, equity loss, time-to-recovery, or another system outcome.
- 2) Stage 2 - Compute singleton impacts. This establishes the visible-risk baseline and prevents JCA from confusing already-dangerous nodes with interaction-generated risk.
- 3) Stage 3 - Build candidate sets using mechanism clues. Prioritize shared substitute pools, common downstream neighbors, common buffers, cross-layer dependencies, correlated hazards, or common recovery channels. Do not restrict candidates to adjacent nodes.
- 4) Stage 4 - Enumerate pairs where feasible. Pairwise JCA is the lowest-cost interaction screen and often reveals the dominant mechanisms that can guide higher-order search.
- 5) Stage 5 - Escalate to triples and larger sets selectively. Use random chemistry, greedy extraction, rare-event sampling, or other combinatorial search methods rather than naive enumeration.
- 6) Stage 6 - Condition on system state. Recalculate under stress regimes - low inventories, tight liquidity, high demand, staff shortages, degraded information, or restricted rerouting - because joint chokepointness can appear only near thresholds.
- 7) Stage 7 - Intervene on mechanisms, not only nodes. If JCA is created by shared substitution, increase substitute independence. If it is created by a conjunctive approval gate, redesign decision rights. If it is created by a common buffer, diversify the buffer.
- 8) This last stage is where JCA becomes more than a risk score. Protecting both members of a dangerous pair may be expensive and brittle. Removing the interaction mechanism can be cheaper. If two suppliers form a latent joint chokepoint because they share one

port, securing a third supplier that uses the same port does not solve the problem. A structurally independent route does.

- 9) The design objective can therefore be written as reducing high-impact JCA edges or hyperedges subject to cost. This invites future optimization work: minimize the upper-tail JCA spectrum rather than simply maximize average robustness. A system that looks excellent on average but contains a handful of catastrophic joint chokepoints may still be strategically fragile.

What is genuinely new here?

The most important claim of this paper is not that simultaneous failures can be worse than separate failures. That idea is old, and the 2026 supply-chain evidence quantifies it with unusual force. Nor is the JCA ratio arithmetically new; it is the generalized form of the amplification factor already used by Fialkowski et al. (2026). The novelty lies in recognizing what follows once that ratio is treated as a general complex-systems object.

First, it creates a formal category - the latent joint chokepoint - for configurations whose members are individually low-impact. Second, it establishes that no node-only centrality screen can be complete in nonlinear threshold systems, because joint amplification can become arbitrarily large while individual centrality vanishes. Third, it separates structural higher-order interaction from impact higher-order interaction. Fourth, it creates a risk representation - the chokepoint hypergraph - that can exist even when the underlying physical network is pairwise. Fifth, it places system condition inside the definition of practical criticality: the same set can move from safe to dangerous as thresholds change.

These points turn a domain result into a theory of diagnostic failure. The diagnostic failure is subtle because singleton analysis can be locally correct. Node A really is harmless alone. Node B really is harmless alone. The error occurs only when those correct local statements are silently composed into a false systemic conclusion: therefore A and B are harmless together. Additivity is the hidden assumption. JCA makes that assumption testable.

The paper's most counterintuitive implication

The deepest implication is that resilience can hide its own failure mode. If a system absorbs every single-node failure, management may become confident that redundancy is working. But perfect singleton robustness can coexist with catastrophic pair fragility. In fact, the better the backup arrangement works one failure at a time, the more ordinary its components may look in a node ranking. The vulnerability migrates into the dependency among backups.

This is why a 'zero single point of failure' declaration is not equivalent to systemic resilience. A system can have no single point of failure and still have a two-point failure set, a three-point failure set, or a shared-condition failure that behaves like a higher-order chokepoint. The relevant maturity question is therefore not whether single points of failure have been eliminated. It is how the minimum cardinality of catastrophic failure sets changes under stress, and whether those sets share a common mechanism that can be redesigned.

Relationship to resilience and critical-transition theory

JCA complements rather than replaces resilience metrics. Gao, Barzel, and Barabasi (2016) seek universal resilience patterns by reducing complex network dynamics; Scheffer et al. (2009) identify generic early-warning signals near critical transitions; Fang et al. (2026) classify percolation dynamics through temporal self-similarity and show that real networks occupy universality classes not captured by idealized models. These approaches ask how close a system

is to losing resilience. JCA asks which combination of perturbations can exploit that proximity.

The synthesis suggests a two-axis risk map. One axis is susceptibility: how close is the current system condition to a transition? The second is combinatorial trigger structure: which sets can cross it? Near a critical transition, a small increase in pairwise or higher-order JCA may be more informative than a large change in average centrality. Conversely, high JCA in a deeply stable regime may remain a theoretical curiosity if absolute joint impact stays small.

Limitations and Research Agenda

Several limitations are deliberate and point directly to the next studies. First, the computational model is a generic threshold cascade. It does not claim to reproduce power grids, banks, hospitals, supply chains, or organizations. Its role is to show that the formal phenomenon is not tied to the production function used in the Ecuadorian supply-chain study. Domain applications require domain-specific impact functions and propagation rules.

Second, the baseline uses 60-node synthetic networks, although every pair is enumerated across 60 independent realizations and the sensitivity study adds 54 further networks. Larger systems will require sampling or extraction methods. The combinatorial challenge is part of the phenomenon, not merely a computational inconvenience: if high-JCA sets are rare, random pair screening can miss the tail. Rare-event methods and guided candidate generation are therefore central to future work.

Third, the present model is static during each cascade. Real systems adapt. Suppliers rewire, banks hoard liquidity, ships reroute, teams bypass formal authority, and AI agents change communication patterns. Adaptive response can suppress or magnify JCA. The natural extension is a dynamic JCA in which the network and the response rule co-evolve after perturbation.

Fourth, JCA is a ratio and ratios become unstable when the denominator is very small. This is why the framework also uses absolute joint impact and $\Delta(S)$, and why latent-joint classification requires a κ threshold. Future work should compare JCA with log-ratio transformations, Shapley-style interaction indices, supermodularity measures, and causal interaction metrics.

Fifth, hazard dependence has not been modeled. A dangerous pair that almost never fails together may contribute less expected systemic risk than a modestly amplified pair exposed to the same cyclone, cyber event, sanction, or policy shock. A full risk formulation should combine JCA with joint perturbation probability: expected interaction loss depends on both the consequence of the set and the probability that the set is jointly activated.

Sixth, the topology result should be interpreted mechanistically rather than universally. Modular networks produced the densest latent joint chokepoints under the chosen threshold process. Other dynamics can reverse the ranking. The higher-order literature already shows that representation and dynamics can change qualitative conclusions (Zhang et al., 2023). JCA must therefore travel with its impact model, not above it.

CONCLUSION

This study concludes that systemic vulnerabilities in complex systems cannot be fully understood through conventional node-based assessments that focus solely on individual importance, centrality, or standalone failure impacts. The introduction of Joint Chokepoint Amplification (JCA) demonstrates that combinations of individually low-impact components can generate disproportionate consequences when they interact under specific system

conditions. Computational experiments across different network topologies revealed that latent joint chokepoints exist within complex structures, with modular networks showing stronger concentrations of interaction-driven vulnerabilities. These findings highlight that systemic risk is not merely a characteristic of critical nodes but also emerges from relationships, dependencies, and collective failure mechanisms among multiple components. Therefore, resilience strategies should incorporate interaction-based assessments to identify hidden configurations that may threaten system stability.

Future research should extend the JCA framework by applying it to empirical datasets from real-world domains, including global supply chains, financial networks, transportation infrastructures, healthcare systems, and digital ecosystems. Further studies should develop dynamic JCA models that incorporate adaptive responses, temporal variations, uncertainty, and the probability of simultaneous disruptions. Comparative research between JCA and other systemic risk measurement approaches, such as Shapley interaction indices, supermodularity analysis, and causal interaction models, is also needed to evaluate its effectiveness and practical applicability. Additionally, future investigations should explore optimization strategies for reducing high-impact joint chokepoints while maintaining operational efficiency, enabling organizations and policymakers to design more robust and adaptive complex systems.

REFERENCES

- Amelkin, V., Venkatesh, S., & Vohra, R. (2024). Contagion and equilibria in diversified financial networks. *Journal of Economic Theory*, 217, 105819. <https://doi.org/10.1016/j.jet.2024.105819>
- Baqee, D., & Rubbo, E. (2023). Micro propagation and macro aggregation. *Annual Review of Economics*, 15(1), 91–123. <https://doi.org/10.1146/annurev-economics-082322-012833>
- Barrot, J.-N., & Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics*, 131(3), 1543–1592. <https://doi.org/10.1093/qje/qjw018>
- Bi, S., Yuan, X., Hu, S., Li, K., Ni, W., Hossain, E., & Wang, X. (2025). Resilience and failure analysis in next-generation communication networks: A contemporary survey. *IEEE Transactions on Network Science and Engineering*.
- Carvalho, V. M., Nirei, M., Saito, Y. U., & Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the Great East Japan Earthquake. *Quarterly Journal of Economics*, 136(2), 1255–1321. <https://doi.org/10.1093/qje/qjaa044>
- Fang, S., Meng, J., Lin, Q., Chen, B., Dong, G., Bai, C., Nagler, J., Deng, Y., Lambiotte, R., & Fan, J. (2026). Temporal self-similarity reveals percolation universality classes in complex networks. *Nature Communications*, 17, 8605. <https://doi.org/10.1038/s41467-026-75436-4>
- Fialkowski, J., Havlin, S., & Thurner, S. (2026). *Catastrophic disruption cascades driven by the nonlinearity of systemic risk* [Preprint]. arXiv. <https://arxiv.org/abs/2607.20068>
- Fu, X., Wang, Y., Huang, X., & Li, W. (2024). Enhancing intrusion tolerance of clustered wireless sensor networks using a double-stage scale-free topology evolution model. *IEEE Sensors Journal*, 24(16), 26716–26730.
- Gao, J., Barzel, B., & Barabási, A.-L. (2016). Universal resilience patterns in complex networks. *Nature*, 530, 307–312. <https://doi.org/10.1038/nature16948>
- Glasserman, P., & Young, H. P. (2016). Contagion in financial networks. *Journal of Economic Literature*, 54(3), 779–831. <https://doi.org/10.1257/jel.20151228>
- König, M. D., Levchenko, A., Rogers, T., & Zilibotti, F. (2022). Aggregate fluctuations in

- adaptive production networks. *Proceedings of the National Academy of Sciences*, 119(38), e2203730119. <https://doi.org/10.1073/pnas.2203730119>
- Liu, L., Han, Z., Chen, X., Luo, C., Zheng, H., Wang, X., & Tang, S. (2026). *Control of harmful information spreading on adaptive higher-order networks via group dissolution* [Preprint]. arXiv. <https://arxiv.org/abs/2608.07874>
- Liu, Z., Xie, X., Shi, M., Tan, H., Luo, H., Ye, L., & Kao, Q. (2026). Complex network approaches for vulnerability assessment and resilience enhancement in power systems. *Sustainable Energy Technologies and Assessments*, 86, 104859.
- Lu, S., Meng, W., Yu, A., Shao, K., Luan, J., He, R., & Liang, J. (2026). *Topological collapse of higher-order interactions bottlenecks collective intelligence in AI agent societies* [Preprint]. arXiv. <https://arxiv.org/abs/2608.15519>
- Makoond, N., Setiawan, A., Buitrago, M., & Adam, J. M. (2024). Arresting failure propagation in buildings through collapse isolation. *Nature*, 629(8012), 592–596.
- Maldonado, D., Cruz, E., Torres, J. A., Cruz, P. J., & Benitez, S. del P. G. (2024). Multi-agent systems: A survey about its components, framework and workflow. *IEEE Access*, 12, 80950–80975.
- Oura, H., Arnould, G., Ding, X., Grippa, P., Gross, M., Laliotis, D., Leika, M., Lepore, C., Letizia, E., Valderrama, L., & Zhang, Y. (2026). *Systemwide stress test at the IMF: Integrating nonbank financial intermediary risks* (IMF Working Paper No. 2026/170). International Monetary Fund. <https://doi.org/10.5089/9798229054980.001>
- Park, C. (2026). Network topology and recovery delay thresholds determine cascading failure vulnerability in sports systems. *Scientific Reports*.
- Polishchuk, D., & Polishchuk, O. (2026). Flow cores and the problem of vulnerability of complex network systems. *Scientific Reports*. <https://doi.org/10.1038/s41598-026-60866-3>
- Verschuur, J., Lumma, J., & Hall, J. W. (2025). Systemic impacts of disruptions at maritime chokepoints. *Nature Communications*, 16, 10421. <https://doi.org/10.1038/s41467-025-65403-w>
- Wen, T., Chen, Y.-W., Ma, J., Syed, T. A., & De, A. (2026). Evaluating disruption risks at the chokepoints of the maritime transport network to support re-routing decision-making. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2026.07.006>
- Xu, J., Zhang, Q., Zhong, Z., He, S., Zhang, C., Lin, Q., Pei, D., He, P., Zhang, D., & Zhang, Q. (2025). OpenRCA: Can large language models locate the root cause of software failures? *International Conference on Learning Representations*, 2025, 84562–84590.
- Yin, R., Ma, X., Yuan, H., Zhai, M., & Guo, C. (2023). A distributed adaptive routing against selective forwarding attack in scale-free network considering cascading failure. *Journal of Complex Networks*, 11(3), cnad021.
- Zhang, Y., Lucas, M., & Battiston, F. (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>