Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Domain-agnostic multiplicative stress framework for systemic collapse

Junwon Lee

herrang0311@dankook.ac.kr

Dankook University

Article

Keywords:

Posted Date: May 7th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8968998/v2

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

ARTICLE

Domain-agnostic multiplicative stress framework for systemic collapse

Junwon Lee[1]*

[1] *Department of Public Administration, Dankook University, Republic of Korea*

* *Corresponding author: Junwon Lee (herrang0311@dankook.ac.kr)*

Abstract

Complex systems rarely collapse from a single perturbation; however, no framework has formalized the coincidence mechanism across domains. Here, we introduce a multiplicative stress integral, $S(t) = \rho(t) \times \Psi(t) \times \Omega(t)$, which decomposes systemic stress into external pressure (ρ), internal amplification (Ψ), and structural degradation (Ω). The multiplicative structure drives the stress toward zero when any channel is quiescent, thereby creating a coincidence filter. The cumulative index $\Pi(t)$ integrates the joint stress with a fixed equal weighting (1:1:1) and no weight optimization. Applied to five major crises, the 2008 financial crisis, Terra–Luna collapse, Fukushima disaster, COVID-19 pandemic, and global supply-chain disruption, the crisis periods separate from controls (1.9× to 3,627×; Fisher $p = 1.66 \times 10^{-11}$), with the multiplicative formulation matching or outperforming alternative models in all cases. When the channels are redundant, the filter dilutes rather than amplifies the signal, which is a falsifiable prediction that is empirically confirmed. Analysis of Π accumulation dynamics reveals three failure modes, ductile, brittle, and preloaded, suggesting that co-occurrence of independent stresses is a shared structural feature of systemic collapse.

Introduction

The collapse of complex systems is one of the most common phenomena in modern society.

The 2008 global financial crisis erased approximately $10 trillion in market value within weeks[1], while the 2022 Terra–Luna collapse destroyed $40 billion in digital assets within 72 h[2]. The 2011 Fukushima Daiichi disaster and the COVID-19 pandemic each triggered cascading cross-sector failures across the energy, finance, public health, social, and global supply chain systems[3–4]. Despite the diversity of these events in terms of domain, mechanism, and scale, a recurring question remains: Do systemic collapses share an underlying structural pattern?

Each of these domains has developed specialized analytical frameworks for crisis detection, yet none provides a common language for cross-domain comparisons. In finance, value-at-risk and its systemic extensions, such as CoVaR[5], quantify tail-risk exposure; DebtRank[6] measures cascading default risk through interbank exposures; and network-based models[7] characterize the propagation of local shocks to system-wide instability. Composite stress indices such as CISS[8] aggregate multiple market indicators into a single measure of financial stress. In epidemiology, compartmental models[9] describe disease transmission but do not capture the economic consequences. In geophysics, the Gutenberg–Richter relation[10] and Omori's aftershock law[11] characterize seismic hazards without addressing coupled financial disruptions. In supply chain analysis, bullwhip-effect models[12] address demand shock propagation, and recent supply-network resilience frameworks[13] emphasize viability under disruption; however, these approaches lack connections to financial subsystems.

Efforts to identify the common structural patterns across domains have led to significant advances. For instance, Sornette[14] proposed that financial crashes exhibit log-periodic power-law signatures analogous to critical phenomena in physics. Scheffer et al.[15–17] demonstrated that generic early warning signals—rising autocorrelation, increasing variance, and critical slowing down—may precede the tipping points across diverse systems. Battiston et al.[6] introduced DebtRank to quantify how network centrality determines systemic importance.

Despite these advances[18], to the best of our knowledge, no existing framework provides a single mathematical structure applicable across fundamentally different domains using common equations with equal (1:1:1) channel weighting and no fitted weights.

Our approach is inspired by the engineering science of cumulative damage. The Palmgren–Miner fatigue damage rule[19–20], introduced in 1945, quantifies how cyclic mechanical loads accumulate damage in materials until failure. This rule has remained the industry standard for eight decades[21] because it captures the essential physics of cumulative degradation without requiring detailed knowledge of the microstructural failure mechanisms. We propose an analogous approach: a multiplicative stress integral that accumulates the damage from three co-occurring stress channels. In materials science, a steel beam does not fail from load, defects, or degradation in isolation; it fails when these factors coincide. This multiplicative logic offers a natural analogy for systemic collapse.

Based on this analogy, we propose the structural stability index (Π), defined as $S(t) = \rho(t) \times \Psi(t) \times \Omega(t)$, where ρ represents external pressure, Ψ internal amplification, and Ω structural degradation. The cumulative damage integral $\Pi(t) = \int S(\tau)\, d\tau / \tau_0$ accumulates stress over time. This study makes three contributions: a domain-agnostic framework with fixed equal weighting (1:1:1) and no weight optimization, validation across five domains using publicly available data, and the identification of three failure modes—ductile, brittle, and preloaded—with distinct temporal signatures. We do not claim to have discovered a universal law; instead, we present evidence that a single framework captures collapse dynamics across five diverse domains and investigate whether this structural regularity reflects deeper commonalities across complex systems approaching failure.

Results

Framework definition

We define the instantaneous systemic stress as the product of three normalized components:

$$\mathbf{S(t) = \rho(t) \times \Psi(t) \times \Omega(t),}$$

where ρ quantifies external pressure, Ψ captures internal amplification, and Ω measures structural degradation. Each component is normalized by its 99th-percentile value (P-limit) from a stable reference period, yielding dimensionless quantities, where values near 1 represent the boundary of the stable-period variation, and values exceeding 1 indicate anomalous stress. The separation ratios were invariant to the choice of the P-limit percentile between 95th and 99.5th (Supplementary Table 1). The cumulative damage integral is

$$\mathbf{\Pi(t) = \int_0^t S(\tau)\, d\tau / \tau_0,}$$

where τ_0 denotes the observation interval expressed in years (τ_0 = 1/365 y for daily data and 1/12 y for monthly data; see Methods). All three components contribute with equal weight (1:1:1); no domain-specific optimization was applied. Table 1 summarizes the assignment of the variables.

Theoretical motivation

The multiplicative form S = ρ × Ψ × Ω is motivated by an information-theoretic argument. When the three stress channels are stochastically independent, the joint probability of simultaneous exceedance factorizes, and the associated Shannon surprise[22–23] (negative log-probability) decomposes additively over channels. As P-limit normalization maps each channel magnitude to a monotone proxy for its marginal exceedance probability, the channel product serves as a parameter-free approximation of the joint surprise, providing a principled

rationale for the multiplicative aggregation of co-occurring independent stresses (Supplementary Note). This reasoning also predicts the boundary condition observed empirically: when channels are correlated, the factorization overstates joint surprise, and the multiplicative product can dilute instead of amplify the stress signal—exactly the pattern observed in the Supply Chain case ($r(\rho,\Omega) = 0.952$). Thus, the nonredundancy criterion ($\max|r| < 0.7$) has a theoretical basis for the independence assumption underlying the multiplicative formulation.

Cross-domain validation

We applied the framework to five systemic crises: traditional finance (2008 financial crisis), digital assets (2022 Terra–Luna collapse), physical infrastructure (2011 Fukushima disaster), public health (2020 COVID-19 pandemic), and global logistics (2021 global supply chain disruption). For each case, we computed Π over a crisis period and a matched control period under noncrisis conditions.

All five cases exhibited crisis Π values consistently exceeding controls, with separation ratios from 1.9× (Terra–Luna) to 3,627× (COVID-19; Fig. 1). Notably, the COVID-19 separation reflects both the unprecedented deviation of pandemic stress indicators from pre-2020 baselines and a floor effect inherent in the multiplicative structure: because confirmed cases (ρ) were zero for 98% of the control period, $S(t) = 0$ by construction during the prepandemic baseline, producing a near-zero denominator. Therefore, the extreme magnitude is not directly comparable to other cases but correctly captures the binary transition from nonpandemic to pandemic conditions. The 2008 financial crisis exhibited the second-highest separation (18.6×). We computed Π across 14 sliding 180-d control windows for the 2008 case to address dependence on a single control period (Supplementary Table 7). The crisis Π exceeded all control windows, with a mean separation of 57.6× (95% CI: 41.0×–107.5×;

Cohen's d = 202.8), confirming that the single control is representative of the stable-period baseline (Fig. 2).

Time-normalized stress intensity

Crisis and control windows differ in duration (Table 2); hence, we define the time-normalized mean stress intensity $\overline{S} = \Pi/T$, where T denotes the window duration in years, to control for window-length bias (Supplementary Table 5). Three cases retain strong separation under time normalization: 2008 Financial ($\overline{S}$ separation 10.8×), COVID-19 (2,573.2×), and Supply Chain (3.3×). However, Terra–Luna (1.1×) and Fukushima (1.2×) show that their Π-based separation is substantially attributable to longer crisis windows. This indicates that these two cases exhibited sustained but low-intensity stress elevation instead of the concentrated multiplicative co-occurrence observed in the 2008 Financial and COVID-19 cases. Thus, the time-normalized metric distinguishes cases in which the stress is both elevated and temporally concentrated from those in which it is diffusely distributed.

Multiplicative structure outperforms alternative formulations

We compared multiplicative ($S = \rho \times \Psi \times \Omega$), additive ($S = \rho + \Psi + \Omega$), and maximum ($S = \max(\rho,\Psi,\Omega)$) formulations across all five cases. The multiplicative structure achieved the highest or equal-highest separation in all cases (Table 3). The advantage was most pronounced in the 2008 Financial case, where multiplication yields 18.6× versus 2.7× for addition, a seven-fold improvement. In the Terra-Luna case, the multiplicative and maximum formulations yielded identical separation (1.9×), consistent with the low overall stress differentiation in this case. This reflects the coincidence filter property: S(t) approaches zero unless all three components are simultaneously elevated.

The Fukushima case empirically demonstrates this. On March 11, 2011, seismic energy (ρ)

spiked but financial markets were closed ($\Psi \approx \Omega \approx 0$); multiplicative stress remained low. On March 14 (the first trading day), all three channels were elevated, producing 35× higher stress. The additive formulation incorrectly identified March 11 as the peak stress, whereas multiplication correctly identified March 14 (Fig. 3).

Channel ablation analysis

We performed a systematic channel ablation analysis comparing nine stress definitions to test whether the three-channel multiplicative formulation is necessary: three single-channel baselines (ρ only, Ψ only, Ω only), three dual-channel products ($\rho \times \Psi$, $\rho \times \Omega$, $\Psi \times \Omega$), and three triple-channel aggregations ($\rho + \Psi + \Omega$, $\max(\rho,\Psi,\Omega)$, $\rho \times \Psi \times \Omega$) across all five domains (Supplementary Table 4).

In the 2008 Financial case, where all three channels are statistically independent ($\max|r| = 0.206$), a clear stepwise pattern emerges: single-channel separation ranges from 2.1× to 5.4×, dual-channel from 3.2× to 14.4×, and the full three-channel product achieves 18.6×. This confirms that each additional independent channel amplifies discriminative power. COVID-19 exhibits a similar hierarchy, though dominated by the ρ channel (647× alone), reflecting the extreme deviation of pandemic case counts from pre-2020 baselines.

However, in the Supply Chain case, where ρ and Ω are highly correlated ($r = 0.952$), the three-channel product (9.2×) is outperformed by the dual-channel $\Psi \times \Omega$ (34.5×). Adding a redundant channel dilutes rather than amplifies the signal, which is consistent with the theoretical expectation that multiplicative amplification requires nonredundant inputs. The Terra–Luna and Fukushima cases showed marginal differences across all methods (1.6×–2.3×), indicating moderate but uniform stress elevation across channels, instead of sharp multiplicative co-occurrence. These results identify the conditions under which the multiplicative formulation provides the greatest benefit when the three channels capture the

independent stress dimensions.

Statistical validation through permutation testing

We tested whether the temporal coincidence of elevated ρ, Ψ, and Ω during crisis periods is attributable to chance by independently shuffling time indices 10,000 times per case[24]. Four of the five cases rejected the null hypothesis at $p < 0.001$, with $z = 12.77$ (2008 Financial) and $z = 17.98$ (COVID-19). The results remained significant after Bonferroni correction[25]. The Supply Chain case did not reach significance ($p = 0.26$) owing to the monthly resolution ($n = 47$); the 9.2× separation remained substantial. Fisher's combined $p = 1.66 \times 10^{-11}$ (ref. [26]) (Supplementary Table 19). Robustness analyses confirmed insensitivity to variable perturbations; 50% Gaussian noise changed the 2008 Financial separation by less than 4% (Supplementary Table 2; Fig. 6). To verify that independent shuffling does not inflate significance by destroying temporal autocorrelation, we repeated the permutation test using block permutation (block sizes B = 5, 10, and 20 d; B = 2, 3, and 4 m for the Supply Chain), which preserved within-block serial dependence while destroying cross-channel alignment (Supplementary Table 8). All four previously significant cases remain significant under block permutation at every block size tested: at B = 20, z-scores are 6.33 (2008 Financial), 3.98 (Terra–Luna), 3.29 (Fukushima, $p = 0.0025$), and 5.16 (COVID-19). The Supply Chain case remains insignificant. The z-scores decreased with increasing block size, as expected from the reduced number of effective permutation units; however, significance was preserved throughout (Fig. 4).

Emergence of three failure modes from collapse dynamics

Three failure modes emerge on examination of Π@collapse/Π_max and temporal lead time. Ductile failures (2008 Financial, Supply Chain) gradually accumulate (50–66% at collapse, 12–18 m lead), analogous to metal fatigue. Preloaded failures (Terra–Luna) exhibit 93% pre-

accumulated failures before the trigger, which is analogous to creep rupture. Brittle failures (Fukushima, COVID-19) exhibit a rapid onset with minimal lead time, analogous to glass fractures. These classifications carry retrospective implications: ductile failures exhibit extended rising Π periods whereas brittle failures offer minimal advance signals. However, prospective validation is required before operational application. Sensitivity analysis across alternative classification thresholds (Supplementary Table 6) confirmed that the Terra–Luna (preloaded) and COVID-19 (brittle) classifications were invariant to all tested cutoff variations. The remaining three cases exhibited boundary sensitivity: the 2008 Financial case (lead = 396 d) transitions from ductile to brittle only when the lead threshold is raised to 400 d, Fukushima (lead = 226 d) transitions from brittle to ductile only when the threshold is lowered to 200 d, and the Supply Chain case transitions when the onset threshold is raised above 10%. These transitions reflect the proximity to the classification boundaries instead of the instability of the taxonomy (Fig. 5; Supplementary Table 20).

Comparison with existing crisis detection methods

We computed crisis–control separation ratios for six methods across all five cases using identical data to benchmark the multiplicative framework against established alternatives: the multiplicative Π, critical slowing down indicators (rolling variance and lag-1 autocorrelation), first principal component (PCA-PC1)[8], additive aggregation, and channel maximum (Supplementary Table 16). The multiplicative Π achieved the highest separation in three of five cases: 2008 Financial (18.6× vs 12.1× for CSD variance), Terra–Luna (1.9×), and Fukushima (2.3×). CSD variance performed best in the COVID-19 case (9,675× vs. 3,627×), where the floor effect in the multiplicative control baseline amplifies the ratios of both methods. PCA-PC1 narrowly outperformed the other methods in the Supply Chain case (10.2× vs 9.2×), consistent with the known ρ–Ω redundancy. CSD autocorrelation—the

canonical early-warning signal16,27—produced near-unity separation (1.1×–1.2×) in four cases; the Supply Chain case yielded a negative ratio (−1.1×), reflecting sign reversal in control-period autocorrelation that renders ratio-based comparison inappropriate for this indicator. Overall, increasing autocorrelation does not reliably precede these diverse collapse types. These results suggest that the multiplicative coincidence filter provides competitive or superior crisis discrimination, particularly when the channel independence holds.

Pseudo-prospective signal analysis

We computed rolling mean stress $\overline{S}(t)$ using an expanding window and defined alert thresholds at 2σ above the control-period mean to assess whether the framework provides an advance signal before collapse (Supplementary Table 17; Fig. 7). In the 2008 Financial case, the 90-d rolling $\overline{S}$ first exceeded the 2σ threshold on August 14, 2007—5 d after the BNP Paribas liquidity freeze and 398 d (13 months) before the Lehman Brothers collapse. By the Lehman bankruptcy date, cumulative Π had already reached 49.9% of its final value, consistent with the ductile classification. Three of the five cases produced precollapse signals: 2008 Financial (398-d lead), Supply Chain (214-d lead), and COVID-19 (47-d lead). The Fukushima case produced no precollapse signal, consistent with its exogenous seismic trigger. The Terra-Luna case exhibited no signal, reflecting an abrupt stabilization mechanism. These results align with the failure mode taxonomy: ductile failures (2008, Supply Chain) exhibit extended pre-collapse warnings, whereas brittle (Fukushima) and preloaded (Terra–Luna) failures do not. We emphasize that this analysis is pseudo-prospective (the thresholds are calibrated from the full control period) and does not constitute a true out-of-sample forecast.

Discussion

The primary finding is that collapse across five fundamentally different domains—finance,

digital assets, geophysics, epidemiology, and logistics—is captured by a single mathematical structure: a multiplicative stress integral with fixed equal weighting (1:1:1) and no weight optimization. This cross-domain regularity is not self-evident; no a priori reason exists why banking-credit stress, stablecoin depegging, seismic–financial coupling, pandemic spillovers, and supply-chain disruptions share a common damage-accumulation pattern. This suggests that the framework detects a structural feature of how complex systems approach failure: collapse requires sustained co-occurrence of stresses across non-redundant channels[28–30].

An equal-weight (1:1:1) design was a deliberate analytical choice. Conventional risk models optimize parameters to fit historical data from a single domain, maximizing in-sample performance, but limiting transferability. We sacrifice potential in-sample gains to preserve domain portability by avoiding weight tuning. Unweighted multiplication produces a consistent separation across five unrelated domains without case-specific adjustment, suggesting that the framework captures the structural features of collapse rather than statistical artifacts.

The failure mode taxonomy has implications for the retrospective understanding of crises. Ductile failures exhibited 12–18 m retrospective lead windows during which Π was rising above baseline; real-time monitoring could in principle provide situational awareness—not deterministic prediction, but an escalating stress signal analogous to seismic hazard assessments[15,31]. For brittle failures, preparedness should focus on resilience instead of on anticipation. Preloaded failures are the most insidious; the system appears stable, although damage accumulates to near-critical levels. Pseudo-prospective analysis provides initial support for this distinction: rolling stress monitoring produced a 13-m advance signal for the 2008 Financial ductile failure but no precollapse signal for brittle (Fukushima) or preloaded (Terra–Luna) failures (Supplementary Table 17; Figure 7). We emphasize that these observations are retrospective, and prospective validation of future crises is required before

operational deployment.

The joint use of cumulative Π and time-normalized $\overline{S} = \Pi/T$ reveals a second dimension of crisis characterization beyond failure modes. Concentrated crises (2008 Financial and COVID-19) exhibited high separation under both metrics, indicating that stress channels are not only elevated but also temporally colocated. Diffuse crises (Terra–Luna, Fukushima) exhibited elevated cumulative Π but near-unity $\overline{S}$ separation, indicating that stress accrues over extended windows without sharp multiplicative peaks. This distinction has practical implications; concentrated crises may be more amenable to threshold-based early warnings, whereas diffuse crises require cumulative monitoring over longer horizons. Thus, the framework provides not only a binary crisis/non-crisis classification but also a two-dimensional characterization of crisis dynamics.

The superiority of the multiplicative structure reflects physical intuition: systemic collapse rarely results from a single perturbation, but requires the co-occurrence of multiple stress channels[29]. Multiplication enforces this mathematically; S(t) approaches zero when any component is near zero. This coincidence-filter property—paralleling the compound-event framework in climate science[32–33], where simultaneous extremes produce impacts disproportionate to individual hazards—admits an information-theoretic interpretation. If each channel has an exceedance probability of approximately 0.01 at the P-limit boundary, then, under independence, the joint probability is 10^{-6}, representing an rare event that the framework is designed to detect. We do not claim that multiplication is a unique combination rule[34]; instead, it functions as an effective coincidence detector that matches or outperforms additive and maximum alternatives in all tested cases, and achieves the highest separation among all nine ablation methods in four of the five domains. A broader benchmarking against critical slowing down indicators and PCA-based stress indices confirms this advantage: the multiplicative Π achieves the highest separation in three of five cases, whereas CSD lag-1

autocorrelation—the most widely cited early-warning signal[16,27]—produces near-unity or anomalous separation across all domains (Supplementary Table 16).

Despite its significant findings, this study had several limitations. N = 5 cases provide initial validation, but do not establish generalizability beyond the tested domains. Variable selection required domain expertise and iterative refinement (Methods); to constrain this degree of freedom, we reported variable substitution robustness (all configurations maintain Sep > 1; Supplementary Table 9), transform-window sensitivity (Sep ranges 16.6×–21.1× across k = 1–20; Supplementary Table 14), and matched-pipeline specificity tests (Supplementary Table 15). Time-normalized analysis revealed that Terra–Luna and Fukushima separation largely reflected longer crisis windows instead of concentrated stress ($\overline{S}$ separation 1.1× and 1.2×; Supplementary Table 5). The Supply Chain case illustrates the principal boundary condition of the framework: high ρ–Ω correlation (r = 0.952) violates channel independence, and the two-channel Ψ × Ω product (34.5×) outperforms the three-channel Π (9.2×). This is not a failure but a falsifiable prediction—multiplicative amplification that requires non-redundant inputs—as confirmed by both the supply chain result and the additional-case analysis (Supplementary Tables 10–12). Out-of-sample validation of the 2023 SVB crisis was not feasible because of LIBOR discontinuation.

The discriminative value of the framework was further demonstrated by a specificity analysis. We applied the 2008 banking-credit variables (DFF × TEDRATE × TOTBKCR) unchanged to the 2000 Dot-com crash and the 2019 repo near-miss, and separately applied supply chain variables to the 2011 floods in Thailand (Supplementary Table 10). Critically, the Dot-com and Repo cases were tested under both a minimal specification (raw weekly levels) and the full 2008 pipeline (|Δ5d| transforms, daily interpolation; Supplementary Table 15), yielding near-identical conclusions: Dot-com Sep = 1.2–1.3×, Repo Sep = 0.7–0.8×, compared with 18.6× for the 2008 crisis—a 14-fold difference that persists regardless of preprocessing. The

banking credit variables remain silent during an equity bubble (Dot-com) and a contained liquidity event (Repo), whereas the supply chain variables replicate elevated separation during a subsequent logistics disruption (Thailand, Sep = 3.5×). Both financial specificity cases violate the non-redundancy criterion (max|r| = 0.765 and 0.866), and both exhibit weak separation—precisely the outcome predicted by the independence-based theory (Supplementary Note 1). The Dot-com case additionally illustrates the distinction between temporal alignment and stress magnitude: permutation testing detects nonrandom co-movement ($z = 8.44$, $p < 0.001$), yet separation remains near unity (1.2×), confirming that channel synchrony alone does not constitute systemic stress without sustained co-elevation above stable-period baselines. Thus, the framework exhibits both sensitivity (detecting the stress type measured by its variables) and specificity (remaining inert to unrelated turbulence).

Several directions for future research have also emerged. Expanding the case set to N = 10–20 domains with preregistered variable protocols would strengthen the generalizability claims and further constrain the degrees of freedom of researchers. Automating the variable selection through mutual information screening can formalize the current expert-driven process. A deeper theoretical derivation from the maximum entropy principles or copula theory[34] would further ground the multiplicative formulation. Deploying Π as a real-time monitoring index would test whether retrospective patterns translate into operational signals, and analysis of near-miss events—systems approaching high Π without collapsing—could reveal conditions under which stress dissipates, providing insight into resilience[28,35].

Methods

Data acquisition

Data were obtained from public sources. Federal Reserve Economic Data (FRED,

https://fred.stlouisfed.org) provided macroeconomic and financial series: effective federal funds rate (DFF), TED spread (TEDRATE), total bank credit (TOTBKCR), personal consumption expenditures on durable goods (PCEDG), delivery times of manufacturers (DTCDISA066MSFRBNY), and freight PPI (WPU3012). Seismic data were obtained from the USGS Earthquake Hazards Program (https://earthquake.usgs.gov), with magnitude converted to energy via the Gutenberg–Richter relation[10] $\log_{10}E = 1.5M + 4.8$. COVID-19 data were obtained from the Johns Hopkins CSSE[36] (https://github.com/CSSEGISandData/COVID-19). Market data (BTC, ETH, LUNA/UST, Nikkei 225, USD/JPY, and the VIX) were obtained from Yahoo Finance and CoinGecko. High-yield credit spreads (BAMLH0A0HYM2) were obtained from FRED. All datasets were accessed between January and February 2025. No proprietary data were used in this study.

Variable design and physical interpretation

Variable selection followed a three-component causal protocol. Each system was decomposed into cause (ρ), response (Ψ), and state (Ω) channels satisfying four criteria. Criterion 1 (Causal ordering): ρ must be temporally upstream of Ψ and Ω. Criterion 2 (domain independence): Each component must originate from a distinct measurement domain. Criterion 3 (observability): All variables must be derived from authoritative public sources. Criterion 4 (non-redundancy): pairwise |r| during stable periods should not exceed 0.7. We acknowledge that the variable selection involved iterative refinement, and the protocol was formalized post hoc. Future applications will benefit from automated identification methods.

Window definitions and normalization

Each variable was normalized by its P-limit (99th percentile during a stable reference period): x_norm(t) = x(t) / P99_stable. For all daily cases, τ_0 = 1/365 y; for monthly cases (Supply Chain, Thailand), τ_0 = 1/12 y. As τ_0 cancels algebraically in the separation ratio

Π_crisis/Π_control, this choice does not affect the primary test statistic. When component series have different native frequencies, alignment is case-specific. For the 2008 Financial case, TOTBKCR (weekly) is linearly interpolated to the daily frequency of DFF and TEDRATE, and boundary values are forward-filled; ρ and Ψ are computed as 5-d absolute changes |Δ5d| of the interpolated daily series. For Terra–Luna, calendar-day series (BTC, LUNA, cross-asset correlation) are aligned via intersection with forward-fill for occasional missing dates. For Fukushima, seismic energy (calendar-day) is forward-filled to trading-day frequency to align with Nikkei and USD/JPY data. For COVID-19, daily case counts (7-d rolling average, calendar-day) and high-yield spreads (business-day) are forward-filled to the VIX trading-day index. For supply chain and the additional validation cases, all series share a common native frequency (monthly or weekly) and are aligned by date intersection without interpolation. The observation-weighted duration $T = N \times \tau_0$ used for time normalization (Supplementary Table 5) reflects the effective observation count after alignment. Three temporally ordered windows are defined for each case, serving distinct analytical roles. (i) The stable reference window is used exclusively for P-limit calibration. It precedes the crisis and contains no recognized stress events. No Π is computed over this window. (ii) The control evaluation window is the noncrisis comparison period over which Π_control is integrated. (iii) The crisis evaluation window is the period encompassing the stress event over which Π_crisis is integrated. The separation ratio Π_crisis / Π_control is the primary test statistic. In practice, windows may overlap depending on data availability and case structure (Supplementary Table 18 lists all dates and observation counts). Two overlap patterns arise. First, the stable and control windows can share a temporal range (e.g., in the 2008 Financial case, both span 2004–2006), because the control period is itself stable and suitable for P-limit calibration; this is conservative, as the same quiet data inform both the normalization baseline and the comparison benchmark. Second, in four of five cases

(Terra–Luna, Fukushima, COVID-19, and supply chain), the control window is a temporal subset of the crisis window: a single continuous observation period is collected, with the early precrisis part serving as control and the full window (including the stress episode) serving as crisis. This nested structure means Π_crisis includes both quiet and stressed days, which conservatively dilutes crisis-period stress intensity; any separation observed is therefore a lower bound on the true crisis–control difference. Critically, in all cases, the stable reference window used for P-limit calibration precedes the crisis episode entirely, preventing look-ahead bias: no information from the crisis period enters the normalization denominator. The identical S(t) calculation and Π integration procedure is applied to both control and crisis windows without modification. Preprocessing transforms (absolute changes, rolling averages, and percentage changes) are listed in Table 1 and were selected based on domain conventions: 5-d absolute changes capture business-week dynamics in interest rate markets; 7-d rolling averages remove weekly reporting cycles in epidemiological data; percentage changes render supply-chain indicators scale-invariant. Supplementary Table 14 lists sensitivity of the 2008 Financial case to alternative transform window widths (k = 1, 3, 5, 10, and 20 d), confirming that the separation ratio is robust to this preprocessing choice (Sep(Π) ranges from 16.6× to 21.1× across all tested k).

Permutation testing

Time indices of ρ, Ψ, and Ω were independently shuffled 10,000 times per case, preserving marginal distributions but destroying temporal alignment[24]. One-sided p-values were computed as the fraction of permuted Π values ≥ observed Π. Bonferroni correction[25] ($\alpha = 0.01$) was applied for multiple comparisons, and Fisher's method[26] combined independent p-values.

Failure mode classification

Failure modes were classified by Π@collapse/Π_max and temporal lead time (days from 10% threshold crossing to collapse). The 10% threshold is the first date Π(t)/Π_max ≥ 0.10. Cases with Π@collapse/Π_max > 0.80: preloaded; lead > 300 d: ductile; remaining: brittle.

Data availability

All data are publicly available from FRED (https://fred.stlouisfed.org), USGS (https://earthquake.usgs.gov), Johns Hopkins CSSE (https://github.com/CSSEGISandData/COVID-19), and Yahoo Finance (https://finance.yahoo.com). and CoinGecko (https://www.coingecko.com).

The source data for all figures and tables are provided with this paper.

Code availability

All analysis code is available at https://github.com/Junwon03/pi-framework-nature and archived at https://doi.org/10.5281/zenodo.18768101

All analyses were performed in Python 3.11 with NumPy 1.26 and Pandas 2.1.

The repository includes a unified pipeline (run_all.py) that reproduces all results reported in this paper, a channel ablation benchmark (run_benchmark.py) that reproduces Supplementary Table 4, a time-normalization and sensitivity analysis script (run_v12_enhancements.py) that reproduces Supplementary Tables 5 and 6, a sliding-window control and block permutation script (run_v13_enhancements.py) that reproduces Supplementary Tables 7 and 8, a variable substitution robustness test (run_variable_substitution.py) that reproduces Supplementary Table 9, a specificity validation script (run_additional_cases.py) that reproduces Supplementary Tables 10–13, sensitivity analysis scripts (sensitivity/) that reproduce

Supplementary Tables 14–15, and a method comparison and pseudo-prospective analysis script (run_method_comparison.py) that reproduces Supplementary Tables 16–17 and Figure 7.

References

1. Brunnermeier, M.K. Deciphering the Liquidity and Credit Crunch 2007-2008. *J. Econ. Perspect.* **23**, 77–100 (2009). https://doi.org/10.1257/jep.23.1.77

2. Briola, A. et al. Anatomy of a Stablecoin's failure: The Terra-Luna case. *Finance Res. Lett.* **51**, 103358 (2023). https://doi.org/10.1016/j.frl.2022.103358

3. Funabashi, Y. & Kitazawa, K. Fukushima in review: A complex disaster, a disastrous response. *Bull. At. Sci.* **68**, 9–21 (2015). https://doi.org/10.1177/0096340212440359

4. Baker, S.R. et al. The Unprecedented Stock Market Reaction to COVID-19. *Rev. Asset Pricing Stud.* **10**, 742–758 (2020). https://doi.org/10.1093/rapstu/raaa008

5. Adrian, T. & Brunnermeier, M.K. CoVaR. *Am. Econ. Rev.* **106**, 1705–1741 (2016). https://doi.org/10.1257/aer.20120555

6. Battiston, S. et al. DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk. *Sci. Rep.* **2**, 541 (2012). https://doi.org/10.1038/srep00541

7. Acemoglu, D., Ozdaglar, A. & Tahbaz-Salehi, A. Systemic Risk and Stability in Financial Networks. *Am. Econ. Rev.* **105**, 564–608 (2015). https://doi.org/10.1257/aer.20130456

8. Holló, D., Kremer, M. & Lo Duca, M. CISS—A Composite Indicator of Systemic Stress in the Financial System. ECB Working Paper Series No. 1426 (European Central Bank, 2012). https://doi.org/10.2139/ssrn.2018792

9. Kermack, W.O. & McKendrick, A.G. A Contribution to the Mathematical Theory of Epidemics. *Proc. R. Soc. A* **115**, 700–721 (1927). https://doi.org/10.1098/rspa.1927.0118

10. Gutenberg, B. & Richter, C.F. Frequency of earthquakes in California. *Bull. Seismol. Soc. Am.* **34**, 185–188 (1944). https://doi.org/10.1785/BSSA0340040185

11. Omori, F. On the after-shocks of earthquakes. *J. Coll. Sci. Imp. Univ. Tokyo* **7**, 111–200 (1894).

12. Lee, H.L., Padmanabhan, V. & Whang, S. The Bullwhip Effect in Supply Chains. *Sloan Manage. Rev.* **38**, 93–102 (1997).

13. Ivanov, D. & Dolgui, A. Viability of intertwined supply networks: extending the supply chain resilience angles towards survivability. *Int. J. Prod. Res.* **58**, 2904–2915 (2020). https://doi.org/10.1080/00207543.2020.1750727

14. Sornette, D. *Why Stock Markets Crash: Critical Events in Complex Financial Systems* (Princeton Univ. Press, 2003).

15. Scheffer, M. et al. Early-warning signals for critical transitions. *Nature* **461**, 53–59 (2009). https://doi.org/10.1038/nature08227

16. Dakos, V. et al. Slowing down as an early warning signal for abrupt climate change. *Proc. Natl Acad. Sci. USA* **105**, 14308–14312 (2008). https://doi.org/10.1073/pnas.0802430105

17. Lenton, T.M. et al. Tipping elements in the Earth's climate system. *Proc. Natl Acad. Sci. USA* **105**, 1786–1793 (2008). https://doi.org/10.1073/pnas.0705414105

18. Bisias, D., Flood, M., Lo, A.W. & Valavanis, S. A Survey of Systemic Risk Analytics. *Annu. Rev. Financ. Econ.* **4**, 255–296 (2012). https://doi.org/10.1146/annurev-financial-110311-101754

19. Palmgren, A. Die Lebensdauer von Kugellagern. *Z. Ver. Dtsch. Ing.* **68**, 339–341 (1924).

20. Miner, M.A. Cumulative Damage in Fatigue. *J. Appl. Mech.* **12**, A159–A164 (1945). https://doi.org/10.1115/1.4009458

21. Fatemi, A. & Yang, L. Cumulative fatigue damage and life prediction theories: a survey. *Int. J. Fatigue* **20**, 9–34 (1998). https://doi.org/10.1016/S0142-1123(97)00081-9

22. Shannon, C.E. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948). https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

23. Cover, T.M. & Thomas, J.A. *Elements of Information Theory* 2nd edn (Wiley, 2006). https://doi.org/10.1002/047174882X

24. Good, P.I. *Permutation, Parametric, and Bootstrap Tests of Hypotheses* 3rd edn (Springer, 2005). https://doi.org/10.1007/b138696

25. Bland, J.M. & Altman, D.G. Multiple significance tests: the Bonferroni method. *BMJ* **310**, 170 (1995). https://doi.org/10.1136/bmj.310.6973.170

26. Fisher, R.A. *Statistical Methods for Research Workers* 14th edn (Oliver & Boyd, 1970). https://doi.org/10.1007/978-1-4612-4380-9_6

27. Boers, N. Observation-based early-warning signals for a collapse of the Atlantic Meridional Overturning Circulation. *Nat. Clim. Change* **11**, 680–688 (2021). https://doi.org/10.1038/s41558-021-01097-4

28. Bak, P. *How Nature Works: The Science of Self-Organized Criticality* (Copernicus, 1996).

29. May, R.M., Levin, S.A. & Sugihara, G. Ecology for bankers. *Nature* **451**, 893–894 (2008). https://doi.org/10.1038/451893a

30. Bardoscia, M. et al. The physics of financial networks. *Nat. Rev. Phys.* **3**, 490–507 (2021).

https://doi.org/10.1038/s42254-021-00322-5

31. Scheffer, M. et al. Anticipating Critical Transitions. *Science* **338**, 344–348 (2012). https://doi.org/10.1126/science.1225244

32. Zscheischler, J. et al. Future climate risk from compound events. *Nat. Clim. Change* **8**, 469–477 (2018). https://doi.org/10.1038/s41558-018-0156-3

33. Leonard, M. et al. A compound event framework for understanding extreme impacts. *WIREs Clim. Change* **5**, 113–128 (2014). https://doi.org/10.1002/wcc.252

34. Nelsen, R.B. *An Introduction to Copulas* 2nd edn (Springer, 2006). https://doi.org/10.1007/0-387-28678-0

35. Haldane, A.G. & May, R.M. Systemic risk in banking ecosystems. *Nature* **469**, 351–355 (2011). https://doi.org/10.1038/nature09659

36. Dong, E., Du, H. & Gardner, L. An interactive web-based dashboard to track COVID-19 in real time. *Lancet Infect. Dis.* **20**, 533–534 (2020). https://doi.org/10.1016/S1473-3099(20)30120-1

Acknowledgements

This research received no external funding.

Author contributions

J.L. conceived the study, designed the framework, collected and curated all data, wrote analysis code, performed statistical validation, created all figures and tables, and wrote the manuscript.

I would like to thank Editage (www.editage.co.kr) for English language editing.

Competing interests

The author declares no competing interests.

AI disclosure

Large language models (Claude - opus 4.5, Anthropic) (GPT - 5.3 codex, OpenAI) were used as assistive tools during code implementation and manuscript editing.

All research design, theoretical formulation, variable selection, statistical interpretation, and scientific conclusions are solely the author's own.

Figure legends

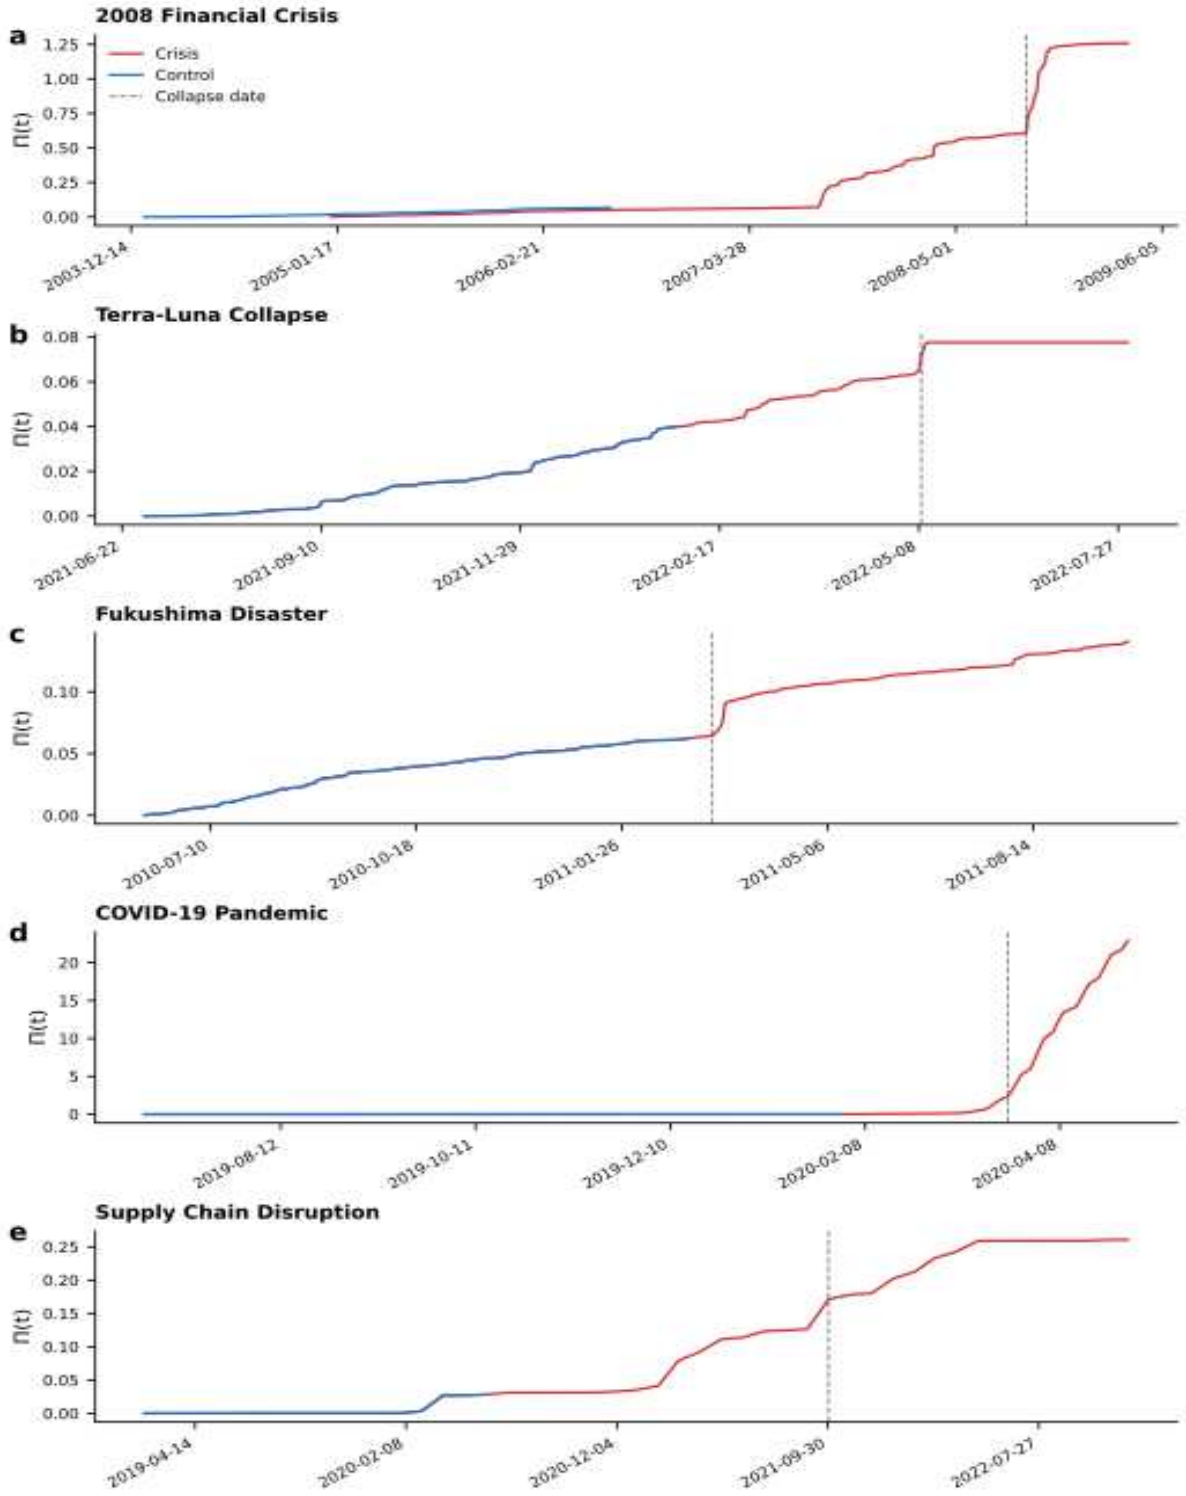


Figure 1 | Cross-domain Π trajectories. Cumulative damage integral Π(t) over the crisis period for each of the five cases: (a) 2008 Financial Crisis, (b) Terra-Luna collapse, (c) Fukushima disaster, (d) COVID-19 pandemic, (e) Supply chain disruption. Dashed vertical lines indicate the collapse date.

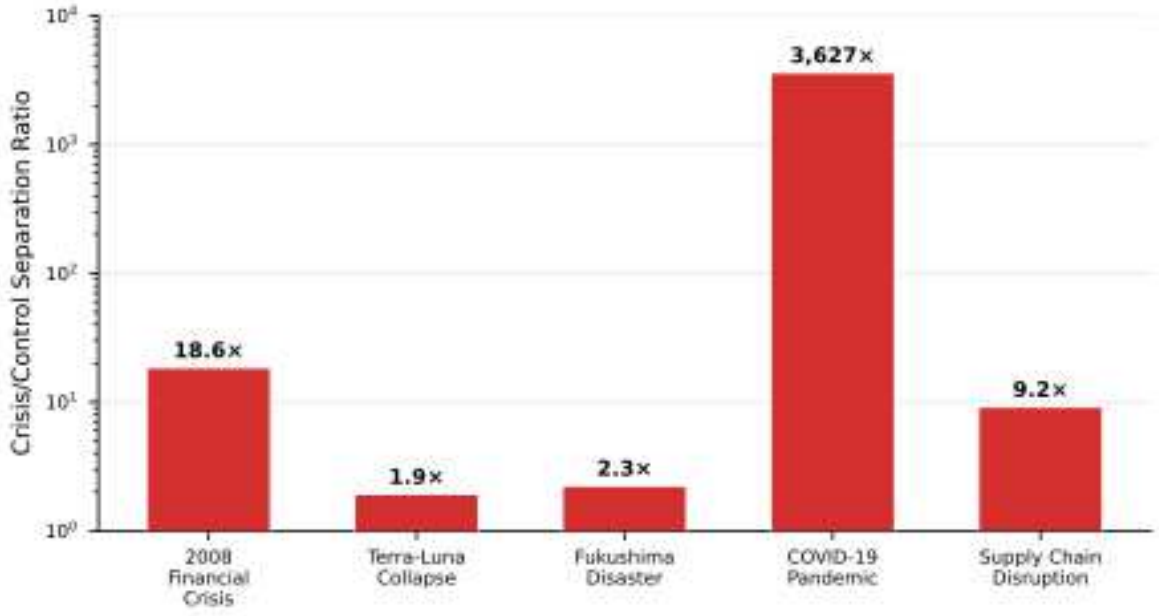


Figure 2 | Crisis versus control separation. Bar chart comparing final Π values between crisis and matched control periods for all five cases. Numbers above bars indicate separation ratios.

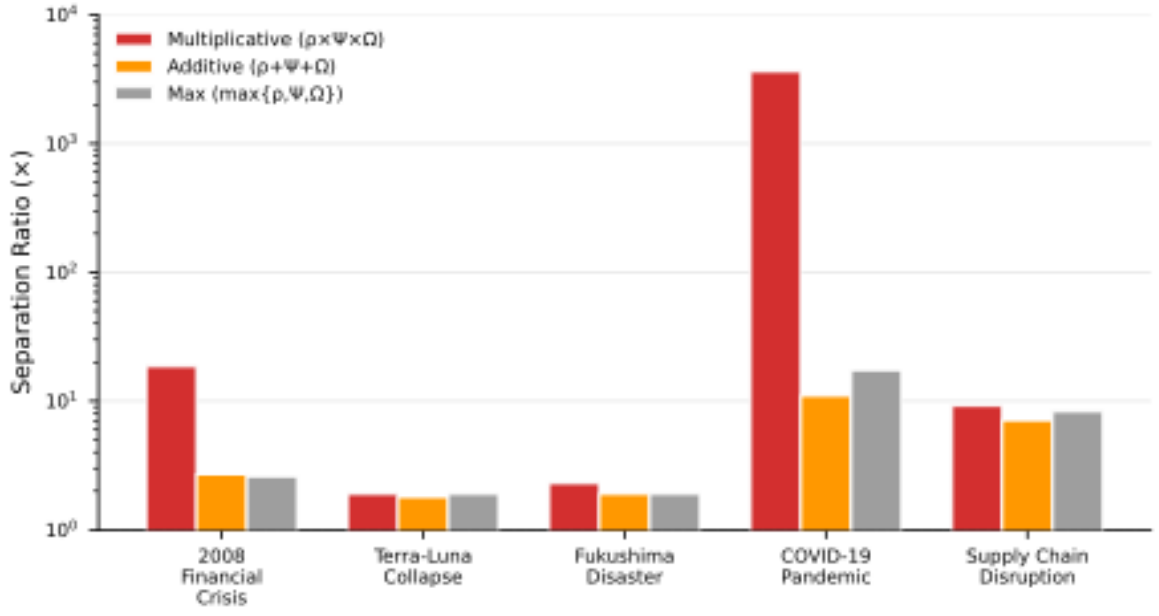


Figure 3 | Multiplicative structure as coincidence filter. Separation ratios for multiplicative, additive, and maximum formulations across all five cases.

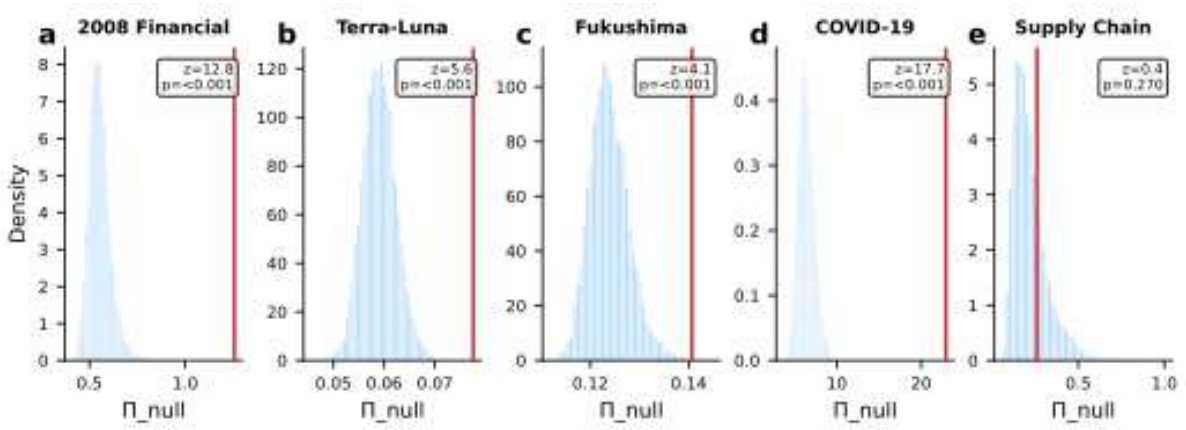


Figure 4 | Permutation test distributions. Histograms of Π from 10,000 permutations (grey) with observed Π (red arrow) for each case. Supply Chain (n = 47) does not reach significance.

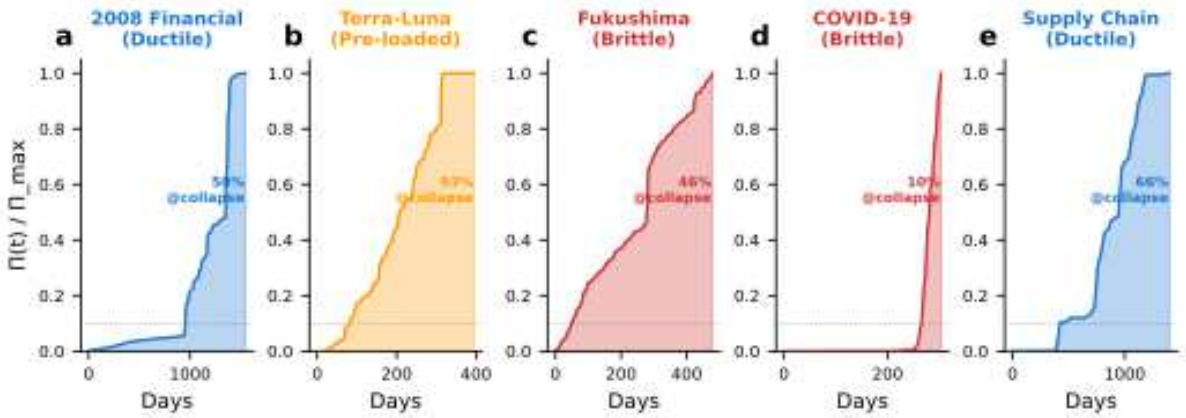


Figure 5 | Failure mode classification. Normalized Π(t)/Π_max trajectories color-coded by failure mode: Ductile (blue), Brittle (red), Preloaded (orange).

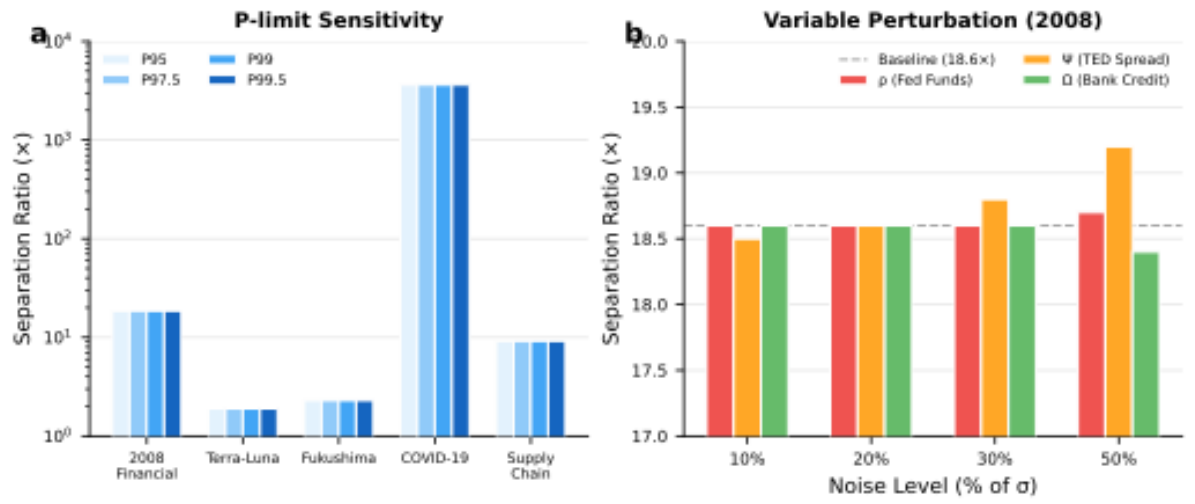


Figure 6 | Robustness analysis. (a) P-limit sensitivity: separation invariant across 95th–99.5th percentiles. (b) Variable perturbation: separation varies by less than ±3.2% under 50% noise.

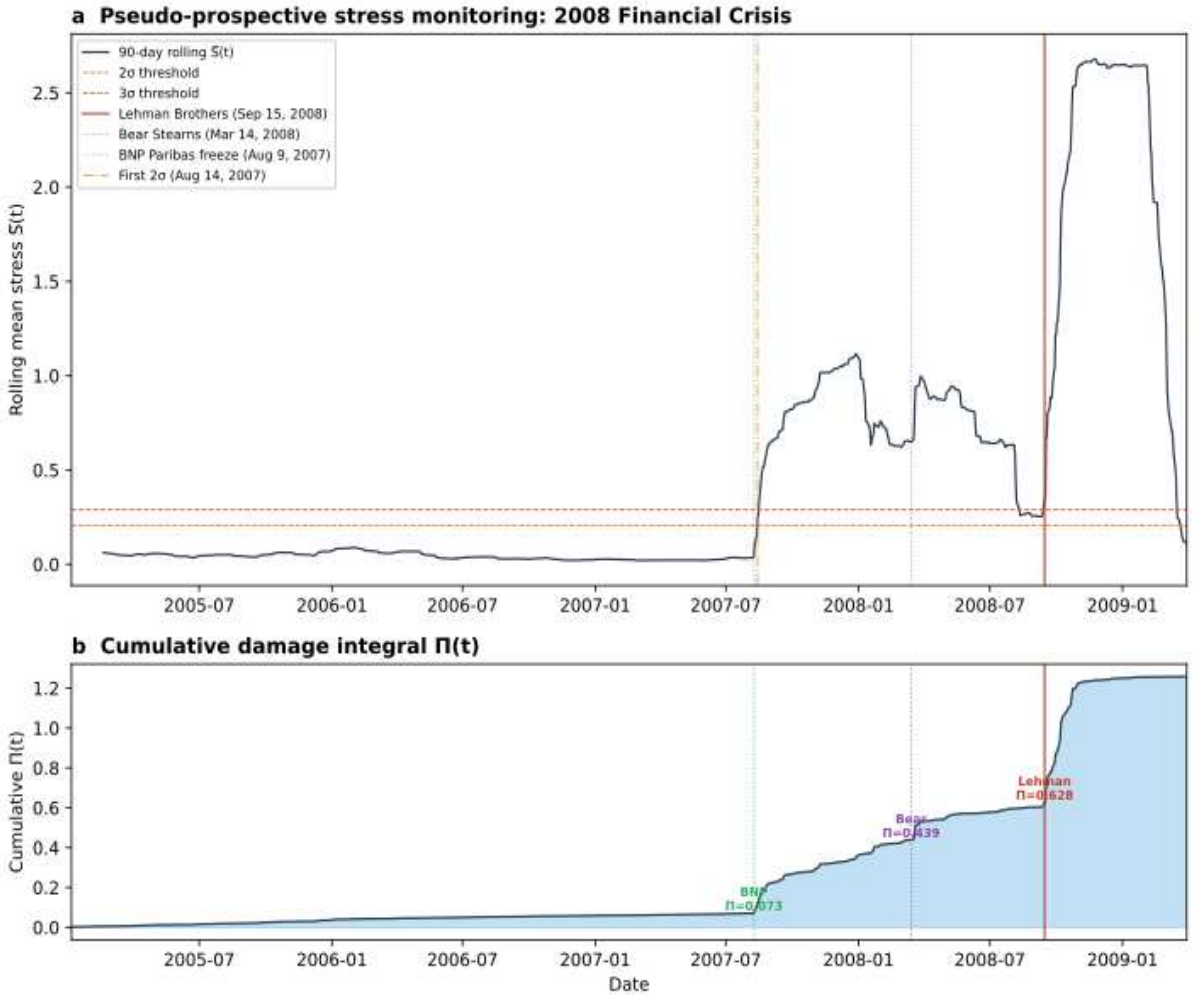


Figure 7 | Pseudo-prospective stress monitoring for the 2008 Financial Crisis. (a) 90-day rolling mean stress

$\bar{S}(t)$ with 2σ and 3σ alert thresholds derived from the control period. The first 2σ exceedance occurs on August 14, 2007, 398 days before the Lehman Brothers collapse. Key events annotated: BNP Paribas liquidity freeze (August 9, 2007) and Bear Stearns rescue (March 14, 2008). (b) Expanding-window cumulative Π(t) with values at key dates, showing progressive damage accumulation consistent with ductile failure dynamics.

Table 1 | Variable definitions and data sources for the five cases studied.

Case	ρ (External pressure)	Ψ (Internal amplification)	Ω (Structural degradation)	Freq.
2008 Financial	Fed Funds Rate |Δ5d| (FRED DFF)	TED Spread |Δ5d| (FRED TEDRATE)	Bank Credit (FRED TOTBKCR)	Daily
Terra-Luna	BTC Price |Δ5d| (Yahoo Finance)	UST/LUNA Price |Δ1d| (CoinGecko)	Cross-asset Correlation (Yahoo Finance)	Daily
Fukushima	Seismic Energy $\log_{10}E$ (USGS)	Nikkei 225 Vol. (Yahoo Finance)	USD/JPY |Δ%| (Yahoo Finance)	Daily
COVID-19	Daily New Cases, 7d avg (JHU CSSE)	VIX Index (CBOE)	HY Credit Spread (FRED BAMLH0A0HYM2)	Daily
Supply Chain	Durable Goods PCE |Δ%| (FRED PCEDG)	Supplier Delivery, clip ≥ 0 (FRED)	Freight PPI |Δ%| (FRED WPU3012)	Monthly

Table 2 | Cross-domain validation results. Separation ratio = Π_crisis / Π_control.

Case	Π_crisis	Π_control	Separation	N (crisis / control)
2008 Financial	1.2579	0.0677	18.6×	990 / 574
Terra-Luna	0.0776	0.0400	1.9×	396 / 215
Fukushima	0.1406	0.0625	2.3×	321 / 177
COVID-19	22.9154	0.00632	3,626.7×	210 / 149
Supply Chain	0.2606	0.0283	9.2×	47 / 17

Table 3 | Aggregation comparison. Separation ratios for three-channel formulations. The multiplicative structure achieves the highest or equal-highest separation in all five cases.

Case	ρ×Ψ×Ω (Multiplicative)	ρ+Ψ+Ω (Additive)	max(ρ,Ψ,Ω) (Maximum)
2008 Financial	18.6×	2.7×	2.6×
Terra-Luna	1.9×	1.8×	1.9×
Fukushima	2.3×	1.9×	1.9×
COVID-19	3,626.7×	11.0×	17.4×
Supply Chain	9.2×	7.1×	8.3×

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SUPPLEMENTARYformanuscript.docx
- Figurebenchmark.pdf