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.

Systemic Trade Collapse Probability under Compounded Geopolitical Shocks: A Max-Stable Process Approach with Bivariate H\"{u}sler-Reiss Spectral Density for the Trump Tariff Regime and Iran Conflict

JANARDAN BEHERA

janardanbeheragreetsyou@gmail.com

Ravenshaw University

Research Article

Keywords: Systemic trade collapse, extreme value theory, max-stable processes, H¨usler-Reiss distribution, bivariate Hawkes process, geopolitical tail risk, trade network fragility, risk amplification

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9190559/v1

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

Additional Declarations: No competing interests reported.

Systemic Trade Collapse Probability under Compounded Geopolitical Shocks: A Max-Stable Process Approach with Bivariate Hüsler-Reiss Spectral Density for the Trump Tariff Regime and Iran Conflict

Dr. Janardan Behera

* Email:janardan@ravenshawuniversity.ac.in

Abstract

Classical approaches to modeling geopolitical trade shocks treat them as additive perturbations to a fundamentally stable system. The simultaneous imposition of sweeping US tariff regimes under the Trump administration and escalating military pressure on Iran expose the inadequacy of this assumption at a historically consequential scale. These are not perturbations. They are correlated tail events whose joint dependence structure in the extremes differs qualitatively from their center-of-distribution behavior, and modeling them additively systematically underestimates the probability of catastrophic multi-corridor trade collapse. We develop a formal framework for the systemic trade collapse probability $\mathrm{TCEP}(\lambda_T, \lambda_I, \Gamma, u)$, defined as the probability that $d \geq 2$ critical bilateral corridors simultaneously fall below threshold fraction u of pre-shock volumes. The framework is grounded in max-stable process theory with a bivariate Hüsler-Reiss spectral density characterizing tail dependence between tariff-induced and conflict-induced disruptions. Joint shock arrivals are modeled as a bivariate marked Hawkes process with a shared CIR intensity component. The TCEP admits a closed-form expression, and the risk amplification factor

$$\mathcal{A}(\Gamma, u) = \frac{\mathrm{TCEP}(\lambda_T, \lambda_I, \Gamma, u)}{\mathrm{TCEP}(\lambda_T, 0, \Gamma, u) + \mathrm{TCEP}(0, \lambda_I, \Gamma, u)}$$

exceeds unity for all $\Gamma < \infty$, with amplification strictly increasing as tail dependence strengthens. Empirical estimation from bilateral trade flow data during the 2018 US-China tariff war and the 2019 Strait of Hormuz crisis yields $\hat{\Gamma} = 1.84$ with 95% confidence interval $[1.31, 2.47]$, corresponding to $\hat{\mathcal{A}} = 2.13$ at the 10% collapse threshold: the compound shock is more than twice as dangerous as either shock alone. Sensitivity analysis across fifteen systemic corridors identifies Hormuz-dependent routes as the dominant systemic risk nodes. Policy implications for reserve management, trade diversification, and multilateral coordination are developed.

Keywords: Systemic trade collapse, extreme value theory, max-stable processes, Hüsler-Reiss distribution, bivariate Hawkes process, geopolitical tail risk, trade network fragility, risk amplification.

MSC 2020: 60G70, 62G32, 62P20, 91B84.

JEL Classification: C58, F13, F51, F52, G15.

1 Introduction

Something broke in early 2025. Not a single trade route, not a bilateral relationship, not a supply contract. The architecture broke. The United States, under the Trump administration, imposed sweeping reciprocal tariffs affecting the majority of its trading partners, pushing the average effective tariff rate to 7.7 percent, the highest since 1947, and generating bilateral tariff walls against China exceeding 47.5 percent [33]. The value of goods flowing from China to the United States fell by nearly 25 percent in the first three quarters of 2025. Simultaneously, the geopolitical pressure on Iran escalated through a series of military confrontations in June 2025 and the catastrophic escalation of February 28, 2026, in which coordinated US-Israeli strikes on Iranian military and nuclear infrastructure triggered retaliatory attacks across the Gulf, the de facto closure of the Strait of Hormuz, and a 70 percent decline in tanker transit through a waterway carrying approximately 20 million barrels of oil per day, representing 27 percent of global maritime oil trade [9]. These are not two separate stories. The tariff regime and the Iran conflict are compounding shocks acting on the same global trading system through different but correlated channels, and their simultaneous occurrence creates a risk environment that no model designed to handle either shock in isolation can adequately represent.

Economists have responded to the tariff shock with gravity model regressions, general equilibrium simulations, and vector autoregressive analyses. These are the right tools for measuring the average effect of a policy change on bilateral trade flows. Jean et al. [22] estimated that global trade flows would contract by between 5.5 and 8.5 percent under various tariff scenarios, with sectors deeply integrated into global value chains absorbing the largest output losses. Abukari and Saluja [1] documented the long-run scenario consequences across three policy paths: technological protectionism, strategic recalibration, and bloc polarization. Each of these contributions is valuable. None of them addresses the tail of the distribution. None of them asks: what is the probability that the trading system does not merely contract by 5.5 percent but experiences a simultaneous collapse of multiple critical trade corridors, each falling below the threshold at which normal commercial operations become unviable? That is a different question, and it requires a different statistical framework.

The literature on extreme value theory in economics and finance has developed powerful tools for exactly this kind of tail-risk quantification. Poon et al. [25] established that the tail dependence structure of financial markets during extreme events is qualitatively different from their dependence in ordinary market conditions, and that ignoring this difference leads to systematic underestimation of joint failure probabilities. Ling and Liu [23] showed through a general contagion risk framework grounded in EVT that the limit behavior of systemic risk measures depends critically on both the tail index of the marginal distributions and the strength of dependence in the extremes, and that the two components cannot be estimated separately without bias. Their result is directly relevant here: when tariff shocks and conflict shocks both push trade corridor volumes toward their lower tails, the joint probability of

simultaneous collapse depends on how the two shock types are correlated in the tail, not in the center of their joint distribution. Standard correlation measures, Pearson or Spearman, tell us nothing about tail dependence. They can be arbitrarily close to zero while the coefficient of tail dependence, defined as $\chi_{ij} = \lim_{u \to 1^-} \mathbb{P}(F_j(X_j) > u \mid F_i(X_i) > u)$, is strictly positive, exactly the regime relevant to systemic collapse risk.

The mathematical object suited to this problem is the max-stable process. A stochastic process $\{Z(s)\}_{s \in \mathcal{S}}$ is max-stable if for any $n \in \mathbb{N}$, the pointwise maximum $\bigvee_{k=1}^{n} Z_k(s)$ of n independent copies, after suitable normalization, has the same distribution as $Z(s)$ itself. Max-stable processes arise as the only possible limits of normalized pointwise maxima of independent and identically distributed stochastic processes, making them the natural class for modeling the extremes of spatial or multivariate fields. In the bivariate case, the Hüsler-Reiss distribution [21] provides the canonical parametric family, characterized by a single variogram parameter $\Gamma > 0$ that governs the strength of tail dependence: as $\Gamma \to 0$ the joint distribution approaches complete dependence, and as $\Gamma \to \infty$ the marginals become asymptotically independent. Engelke et al. [15] developed the peaks-over-threshold estimation framework for Hüsler-Reiss distributions and their spatial extensions, the Brown-Resnick processes, providing computationally efficient estimators based on the result that conditional increments of any process in the max-domain of attraction of the Hüsler-Reiss model asymptotically follow a multivariate Gaussian distribution. Their methodology is the statistical engine of the empirical analysis in this article.

What the existing EVT literature has not done is apply this framework to the problem of compounded geopolitical shocks on trade networks. Zhang [36] provided a comprehensive review of extreme value methods for systematic risk and extreme co-movements but confined the analysis to financial markets where the relevant processes are return series rather than physical trade flows. The trade network literature, reviewed thoroughly by Dadkhah [10] in the context of the current tariff regime, treats bilateral trade flows as fundamentally Gaussian processes perturbed by policy shocks, ignoring the heavy tails and asymptotic dependence that characterize extreme disruption events. The gap is structural. Closing it requires specifying a joint probability model for bilateral trade flow collapses under compounded shocks, estimating its parameters from historical disruption data, and deriving the systemic collapse probability as an analytically tractable function of the shock intensities and the tail dependence parameter.

This article provides that specification, estimation, and derivation. The central quantity of interest is the systemic trade collapse probability $\mathrm{TCEP}(\lambda_T, \lambda_I, \Gamma, u)$, defined as the probability that at least d^* of the d trade corridors in a designated systemic set each experience flow reductions exceeding $u \in (0, 1)$ of their pre-shock volumes simultaneously, where $\lambda_T > 0$ and $\lambda_I > 0$ are the intensities of the tariff-shock and Iran-conflict-shock point processes respectively, $\Gamma > 0$ is the Hüsler-Reiss variogram parameter governing their tail dependence, and u is the collapse threshold. The key theoretical contribution is the derivation of a closed-form expression for TCEP in terms of the bivariate Hüsler-Reiss spectral density,

proving that the compounding effect strictly amplifies the systemic collapse probability relative to the sum of the individual shock probabilities for all $\Gamma < \infty$, and quantifying the amplification through the risk amplification factor $\mathcal{A}(\Gamma, u) > 1$.

The empirical contribution is the estimation of $\hat{\Gamma} = 1.84$ from historical bilateral trade flow data during two reference crisis episodes: the 2018-2019 US-China tariff escalation, which generated documented bilateral flow reductions of 22-38 percent in affected product categories [2], and the 2019 Strait of Hormuz tanker seizure incidents, which temporarily reduced transit volumes by 15-18 percent. The resulting amplification factor $\mathcal{A} = 2.13$ at the 10 percent collapse threshold means that the compound shock of the current tariff regime plus the active Strait of Hormuz disruption is more than twice as dangerous as the sum of the individual shocks, a result that no linear additive model can generate and that standard gravity model simulations systematically miss.

The article is organized as follows. Section 2 reviews the three research streams that the framework bridges. Section 3 develops the bivariate point process model for compounded geopolitical shocks. Section 4 introduces the max-stable process framework and the Hüsler-Reiss spectral density. Section 5 derives the systemic trade collapse probability and the amplification factor. Section 6 develops the statistical estimation procedure. Section 7 reports the empirical results. Section 8 conducts sensitivity analysis across fifteen trade corridors. Section 9 develops policy implications, and Section 10 concludes.

2 Literature Review

Three bodies of literature converge in the framework developed in this article. None of them, individually, is sufficient. The extreme value theory literature has built a mathematically complete machinery for modeling tail dependence, max-stable processes, and joint exceedance probabilities, but it has been applied almost exclusively to financial returns and environmental extremes, never to bilateral trade flow collapses under compounded geopolitical shocks. The trade network and geopolitical risk literature has documented, with increasing empirical precision, that geopolitical distance depresses trade flows and that compounding political tensions amplify this effect, but its statistical tools, gravity model regressions and VAR systems, are designed for average effects, not tail probabilities. The systemic risk and financial contagion literature has developed point process models for compounded shock arrival and shown that independence assumptions systematically underestimate joint failure probabilities, but it has not migrated these tools to the trade context where the relevant extremes are flow volumes rather than asset returns. The gap between these three streams is the intellectual space this article occupies.

2.1 Extreme Value Theory: Foundations and Statistical Methodology

The mathematical foundations of extreme value theory were laid by Fisher and Tippett in 1928 and formalized by Gnedenko in 1943, who established that the only possible non-degenerate limit distributions for normalized maxima are members of the generalized extreme value family. The modern treatment, covering both the block maxima and peaks-over-threshold approaches with full asymptotic theory, is presented in the authoritative monograph by de Haan and Ferreira [12], which remains the standard reference for the theoretical foundations exploited in this article. For the multivariate extension, the central object is the stable tail dependence function $\ell : [0,\infty)^d \to [0,\infty)$, which satisfies $\ell(\boldsymbol{y}) = \lim_{t\to\infty} t\,\mathbb{P}(1 - F_1(X_1) \le y_1/t \text{ or } \ldots \text{ or } 1 - F_d(X_d) \le y_d/t)$ and characterizes the multivariate extreme value distribution through $G(\boldsymbol{x}) = \exp\{-\ell(-\log G_1(x_1), \ldots, -\log G_d(x_d))\}$.

The peaks-over-threshold approach, which forms the statistical basis for the estimation methodology of Section 6, was developed rigorously in the multivariate setting by Rootzén and Tajvidi [28], who established that the natural limit distribution for threshold exceedances above high multivariate thresholds is the multivariate generalized Pareto distribution (mGPD), characterized by the exponent measure ν satisfying $\nu(\{\boldsymbol{z} \in [0,\infty)^d : z_j > y_j \text{ for some } j\}) = \ell(\boldsymbol{y})$. Their result is the multivariate analog of the Pickands-Balkema-de Haan theorem and provides the theoretical justification for using threshold exceedances rather than block maxima for estimating the tail dependence structure Γ in the Hüsler-Reiss model. The statistical inference for this class of models has been substantially advanced by Wadsworth and Campbell [34], who developed a geometric approach to estimation of the limit set shape that accommodates both asymptotic dependence and asymptotic independence simultaneously, extending the applicability of multivariate EVT beyond the classical asymptotic dependence assumption that would otherwise restrict the analysis to situations where the tail dependence coefficient $\chi_{ij} > 0$.

Extremes in financial time series received comprehensive treatment in the foundational monograph by Embrechts, Klüppelberg and Mikosch [13], which established the connections between heavy tails, subexponential distributions, and extreme value theory in an actuarial and financial context. Their treatment of the compound Poisson approximation for rare event probabilities, specifically the result that for a Poisson process with intensity λ and i.i.d. marks $\{X_k\}$ with distribution F, the compound Poisson measure satisfies $\mathbb{P}(N(t) = n) = e^{-\lambda t}(\lambda t)^n/n!$, provides the building block for the bivariate point process structure developed in Section 3. Coles [8] gave the accessible statistical treatment of extreme value methodology that remains the standard entry point for applied researchers, introducing the profile likelihood approach to uncertainty quantification in extreme value parameter estimation and the diagnostic toolkit for assessing goodness of fit at high thresholds. Both references inform the estimation and uncertainty quantification procedures of Section 6.

The extension of univariate EVT to portfolio risk contexts under multivariate tail dependence was developed by Božović [4], who combined orthogonalized return residuals with univariate extreme value

estimation to produce a fully multivariate tail risk model that captures both marginal tail behavior and extremal co-movement in a single tractable framework. Their approach is directly analogous to the strategy adopted here for trade corridors: estimate marginal tail behavior for each corridor separately, then estimate the joint tail dependence structure through the Hüsler-Reiss variogram. The tail dependence and smoothness analysis of Ferreira and de Haan [17], which connected the tail dependence coefficient χ to the smoothness properties of stationary processes in the ℓ_1-embeddable framework, provides the theoretical background for the exceedance set analysis used to characterize the systemic collapse threshold $u^*(\Gamma)$ in Section 5.

2.2 Max-Stable Processes and the Hüsler-Reiss Family

Max-stable processes are the natural continuous- space generalization of multivariate extreme value distributions, representing the only possible limits of normalized pointwise maxima of i.i.d. stochastic processes. The representation theorem of de Haan [11], which states that every simple max-stable process $Z(s)$ has the spectral representation $Z(s) = \max_{i \geq 1} \xi_i \, f(s; \boldsymbol{\omega}_i)$ where $\{\xi_i, \boldsymbol{\omega}_i\}$ is a Poisson process on $(0, \infty) \times \Omega$ with intensity $\xi^{-2} d\xi \times d\nu(\boldsymbol{\omega})$ and $\{f(\cdot; \boldsymbol{\omega})\}$ is a family of non-negative functions satisfying $\int_\Omega f(s; \boldsymbol{\omega}) \, d\nu(\boldsymbol{\omega}) = 1$ for all $s \in \mathcal{S}$, underpins the entire parametric modeling strategy of the present article. The spectral measure ν encodes the tail dependence structure, and the Hüsler-Reiss model corresponds to the specific choice of a Gaussian spectral density. Smith [30] independently characterized the class of max-stable processes through a spatial storm process representation and showed that the resulting class of distributions, parameterized by a variogram function $\gamma(h) = \mathbb{E}[(W(s+h) - W(s))^2]/2$ where W is a Gaussian process, coincides with the Brown-Resnick processes introduced by Brown and Resnick [5].

The Schlather process [29], another parametric family of max-stable processes, uses a correlation function $\rho(h)$ rather than a variogram to parameterize tail dependence, and is asymptotically equivalent to the Hüsler-Reiss model for slowly varying ρ. Its bivariate distribution is $\mathbb{P}(Z(s) \leq z_1, Z(t) \leq z_2) = \exp\{-\frac{1}{z_1} - \frac{1}{z_2} + \frac{1}{z_1 z_2} \, T(\rho(h), z_1, z_2)\}$, where T is a bivariate correction factor depending on the correlation and the two threshold levels. The choice between the Schlather and Brown-Resnick families for applied modeling depends on the decay properties of the empirical variogram: the Brown-Resnick model is more flexible for processes whose tail dependence weakens monotonically with increasing spatial distance, which is the relevant behavior for bilateral trade corridor distances in geopolitical space. The composite likelihood approach to max-stable process estimation was developed by Huser and Davison [20], who showed that pairwise likelihood maximization produces consistent and asymptotically normal estimators of the Hüsler-Reiss variogram parameters even when the full likelihood of a max-stable process is computationally intractable. Their result makes the estimation of Γ in Section 6 feasible from moderate-sized historical crisis datasets.

The sparse structure of extremal dependence in high-dimensional settings, where most pairs of variables are asymptotically independent but a small number of pairs exhibit strong tail dependence, was characterized by Engelke and Ivanovs [14] through the concept of the extremal graph, a graph whose edges represent pairs with strictly positive tail dependence. Their result that the joint tail probability factorizes over the connected components of the extremal graph provides the theoretical foundation for treating the d trade corridors as a partially connected extremal network rather than requiring full dependence among all corridor pairs.

The theoretical connection between tail dependence, exceedance sets, and metric embeddings was rigorously established by Strokorb and Schlather [32], who showed that the extremal coefficients of a random vector can be expressed through random exceedance sets in a framework directly extending the theory of ℓ_1-embeddable metrics. Their connection between the multivariate regular variation framework and Bernoulli compatibility provides the algebraic structure underlying the systemic collapse probability derivation in Section 5, specifically the result that the probability of simultaneous exceedance of all d^* corridors decomposes into a sum over subsets of the extremal coefficient matrix.

2.3 Trade Network Fragility and Geopolitical Risk

The gravity model has been the workhorse of empirical trade analysis since Tinbergen's original formulation, and its extensions to incorporate political and institutional determinants of bilateral flows represent the primary quantitative tool available for measuring the impact of geopolitical shocks on trade. Gupta et al. [19] were among the first to incorporate the Caldara-Iacoviello geopolitical risk index directly into a gravity model specification, finding significant negative effects of geopolitical risk on bilateral trade flows across 164 countries over 1985-2013. A one-unit increase in the GPR index reduces bilateral trade by approximately 0.5%, a result whose modest magnitude conceals a distributional heterogeneity that gravity model regressions cannot detect: the average effect is small because geopolitical risk episodes are typically short-lived, but the tail effect during severe episodes is orders of magnitude larger.

This distributional heterogeneity is exactly the empirical phenomenon that motivates the EVT approach. Caldara and Iacoviello [6] constructed the comprehensive geopolitical risk index that has become the standard input for empirical studies in this area, demonstrating through structural VAR analysis that geopolitical risk shocks reduce investment, output, and trade flows through a combination of heightened uncertainty and direct supply chain disruption. Their index is constructed from newspaper text counts and is therefore a continuous measure of geopolitical stress rather than a binary event indicator, which is important for the point process specification of Section 3 where escalation is modeled as a threshold crossing of the geopolitical stress intensity.

The geoeconomic fragmentation literature, represented by Gopinath et al. [18], has documented the structural shift in global trade patterns since 2016 through a comprehensive analysis of trade and financial

flows disaggregated by geopolitical bloc membership. Their finding that trade between hypothetical US-aligned and China-aligned blocs has grown significantly slower than intra-bloc trade since 2022 reflects a long-run structural change that is distinct from but correlated with the tariff-specific and conflict-specific shocks modeled here. The WTO [35] documented through high-dimensional fixed effects gravity regressions that bilateral trade flows became significantly more sensitive to geopolitical distance after 2022, establishing that the regime shift documented by gravity model estimates is statistically robust to standard endogeneity corrections. A one-standard-deviation increase in geopolitical distance reduces bilateral trade by approximately 5%, with the effect concentrated in sectors with high global value chain integration, precisely the corridors with the highest systemic collapse risk.

The energy-trade nexus under geopolitical stress was examined by Son and Ryu [31], who constructed the first comprehensive gravity model of military alliance effects on bilateral energy trade across seven energy categories. Their finding that military alliance membership increases bilateral energy trade by 15-22% while geopolitical distance reduces it proportionally provides the energy market calibration for the Strait of Hormuz corridor analysis of Section 7. The network structure of global trade under sequential geopolitical shocks was analyzed through a graph-theoretic lens by Climent-Serra et al. [7], who showed that the global trade network's giant component remains intact through most shocks but that centrality measures for systemically important nodes deteriorate sharply during compounding events. Their identification of bridge countries as the highest-risk nodes in the trade network is consistent with the systemic collapse probability analysis of Section 8.

2.4 Compounded Shock Models and Systemic Risk

The modeling of compounded shocks in systemic risk contexts has a substantial literature rooted in financial contagion analysis. The key insight, going back to Allen and Gale's network model of financial contagion, is that the probability of systemic failure under compounded shocks exceeds the sum of the individual failure probabilities whenever the shocks share a common causal driver. In the present context, the tariff regime and the Iran conflict share several common drivers: domestic political pressures in the United States, global energy market conditions, and the strategic competition between the United States and China that motivates both the tariff escalation and the Iran policy posture. Any model that treats the two shocks as independent misspecifies the joint arrival process and underestimates the compounding risk.

The bivariate Hawkes process, in which each event in one sequence increases the intensity of both sequences, provides the natural model for this kind of inter-process contagion. Fajgelbaum et al. [16] established empirically that the 2018 tariff escalations produced retaliatory dynamics consistent with a self-exciting point process, where each US tariff announcement increased the probability of a Chinese retaliatory action within the following ninety days. Their result motivates the bivariate marked point

process specification of Section 3, though the present article adds the Iran conflict sequence as a second interacting process rather than modeling only US-China retaliation dynamics. Beirlant et al. [3] provided the comprehensive statistical treatment of point process representations in extreme value theory, specifically the result that the point process $\mathcal{N}_n = \sum_{i=1}^{n} \delta_{(i/n,\, \boldsymbol{X}_i/n^{1/\alpha})}$ converges weakly to a Poisson process on $(0,1] \times ([0,\infty]^d \setminus \{\mathbf{0}\})$ with intensity measure $dt \times d\nu(\boldsymbol{x})$, where ν is the exponent measure of the limiting multivariate extreme value distribution. This convergence result is the theoretical bridge between the bivariate point process model for shock arrivals in Section 3 and the max-stable process model for trade flow extremes in Section 4.

The Pickands spectral representation [24], which characterizes multivariate extreme value distributions through a spectral measure H on the unit simplex Δ_{d-1} satisfying $\ell(\boldsymbol{y}) = \int_{\Delta_{d-1}} \max_{j=1,\ldots,d}(y_j w_j)\, dH(\boldsymbol{w})$, provides the canonical decomposition of the stable tail dependence function into a spectral component that captures the angular structure of extremal dependence and is used in Section 5 to derive the closed-form expression for the systemic trade collapse probability. For the Hüsler-Reiss model in the bivariate case, the spectral density with respect to the arc-length measure on $\Delta_1 = \{(w, 1-w) : w \in [0,1]\}$ takes the form derived by Engelke et al. [15], and is reproduced in Section 4.3 for completeness. The Resnick [27] formalization of regular variation and its connection to extreme value theory through the vague convergence framework provides the measure-theoretic foundation that underpins the TCEP derivation and ensures that the limit expressions derived in Section 5 are mathematically rigorous rather than merely heuristic. It is a foundational work that the present article draws on extensively in the proofs of Theorem 2 and Theorem 3.

3 Bivariate Point Process Model for Compounded Geopolitical Shocks

Before any probability can be attached to systemic trade collapse, the shock arrival process must be specified with mathematical precision. Two sequences of geopolitical escalation events are relevant here. The first is the sequence of tariff escalation events generated by the Trump administration, including announced tariff increases, expanded product scope, and retaliatory counter-tariffs that constitute the observable record of the trade war. The second is the sequence of Iran conflict escalation events, including military strikes, tanker seizures, Strait of Hormuz closure episodes, and sanctions intensifications that constitute the observable record of the Iran confrontation. Both sequences are random. Both have internal self-exciting dynamics. And, crucially, they are not independent: each Trump tariff escalation increases the strategic pressure on Iran by signaling US willingness to use economic coercion, and each Iran escalation increases the political cost of trade liberalization by reinforcing the domestic political coalition that supports tariff protection. The joint process must capture both the marginal dynamics of

each sequence and their mutual excitation.

3.1 Marked Bivariate Point Process Specification

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space. Define two sequences of random times

$$\mathcal{T}_T = \{T_1^T < T_2^T < \cdots\} \quad \text{and} \quad \mathcal{T}_I = \{T_1^I < T_2^I < \cdots\}, \tag{1}$$

representing the arrival times of tariff escalation events and Iran conflict escalation events respectively, defined on the same probability space. Each event carries a mark: $M_k^T \in \mathcal{M}_T$ is the severity of the k-th tariff escalation event, measured in percentage-point increase in the effective tariff rate on the affected product categories, and $M_j^I \in \mathcal{M}_I$ is the severity of the j-th Iran conflict escalation event, measured through a composite index incorporating the fraction of Strait of Hormuz transit disrupted, the number of military assets deployed, and the magnitude of new sanctions imposed.

The joint process $\mathbf{N}(t) = (N_T(t), N_I(t))$ where $N_T(t) = \sum_k \mathbf{1}\{T_k^T \le t\}$ and $N_I(t) = \sum_j \mathbf{1}\{T_j^I \le t\}$ counts the cumulative number of events in each sequence up to time t. The conditional intensity functions $\lambda_T(t)$ and $\lambda_I(t)$ of the two processes, representing the instantaneous arrival rates conditional on the full history $\mathcal{H}_t = \sigma(\mathbf{N}(s) : s \le t)$, are specified as a bivariate Hawkes process with mutual excitation:

$$\lambda_T(t) = \mu_T + \int_0^t \phi_{TT}(t-s)\, dN_T(s) + \int_0^t \phi_{IT}(t-s)\, dN_I(s), \tag{2}$$

$$\lambda_I(t) = \mu_I + \int_0^t \phi_{TI}(t-s)\, dN_T(s) + \int_0^t \phi_{II}(t-s)\, dN_I(s), \tag{3}$$

where $\mu_T > 0$ and $\mu_I > 0$ are the baseline intensities of the two processes, representing the unconditional arrival rates in the absence of any prior excitation. The kernel functions ϕ_{TT}, ϕ_{IT}, ϕ_{TI}, and ϕ_{II} govern the self-exciting and cross-exciting dynamics. Specifically, $\phi_{TT}(t-s) = \alpha_{TT} e^{-\beta_{TT}(t-s)}$ is the self-excitation kernel of the tariff process: a tariff escalation event at time s temporarily raises the intensity of subsequent tariff escalations by $\alpha_{TT} > 0$, decaying exponentially with rate $\beta_{TT} > 0$. The cross- excitation kernel $\phi_{IT}(t-s) = \alpha_{IT} e^{-\beta_{IT}(t-s)}$ models the effect of Iran conflict escalations on subsequent tariff escalation intensity. By symmetry, ϕ_{TI} and ϕ_{II} govern the cross- and self-excitation of the Iran process.

The full specification of the bivariate exponential Hawkes process is therefore parameterized by the vector $\boldsymbol{\theta} = (\mu_T, \mu_I, \alpha_{TT}, \alpha_{IT}, \alpha_{TI}, \alpha_{II}, \beta_{TT}, \beta_{IT}, \beta_{TI}, \beta_{II}) \in \mathbb{R}^{10}_{>0}$. Stationarity of the bivariate process requires the spectral radius condition

$$\rho\left(\begin{pmatrix} \frac{\alpha_{TT}}{\beta_{TT}} & \frac{\alpha_{IT}}{\beta_{IT}} \\ \frac{\alpha_{TI}}{\beta_{TI}} & \frac{\alpha_{II}}{\beta_{II}} \end{pmatrix}\right) < 1, \tag{4}$$

where $\rho(\cdot)$ denotes the spectral radius. When condition (4) holds, the stationary mean intensities satisfy $(\bar{\lambda}_T, \bar{\lambda}_I)^T = (I - A)^{-1}(\mu_T, \mu_I)^T$ where $A_{jk} = \alpha_{jk}/\beta_{jk}$ is the matrix of branching ratios.

3.2 Shared Intensity Component and Common Geopolitical Driver

Equations (2) and (3) model the interaction between the two shock sequences through their observable histories. But there is an additional source of dependence that is not captured by the observable event records: the latent common geopolitical driver that simultaneously raises the probability of both tariff escalation and Iran conflict escalation without necessarily manifesting as an observable event in either sequence. The domestic political cycle in the United States, the strategic posture of the executive branch, and global energy market conditions all represent latent drivers of this kind.

To capture this, we augment the bivariate Hawkes specification with a shared intensity component driven by a latent process $\{V(t)\}_{t\geq 0}$. The augmented intensity functions are

$$\lambda_T(t) = \mu_T + \kappa_T V(t) + \int_0^t \phi_{TT}(t-s)\, dN_T(s) + \int_0^t \phi_{IT}(t-s)\, dN_I(s), \tag{5}$$

$$\lambda_I(t) = \mu_I + \kappa_I V(t) + \int_0^t \phi_{TI}(t-s)\, dN_T(s) + \int_0^t \phi_{II}(t-s)\, dN_I(s), \tag{6}$$

where $\kappa_T \geq 0$ and $\kappa_I \geq 0$ are the loadings of each process on the common driver, and $V(t)$ is modeled as a Cox-Ingersoll-Ross process satisfying

$$dV(t) = a(b - V(t))\, dt + \sigma_V \sqrt{V(t)}\, dW(t), \tag{7}$$

with mean-reversion speed $a > 0$, long-run mean $b > 0$, volatility $\sigma_V > 0$, and $W(t)$ a standard Brownian motion independent of the event histories. The Feller condition $2ab > \sigma_V^2$ ensures that $V(t) > 0$ almost surely, keeping both intensities strictly positive. The long-run mean intensity of the tariff process under this specification is $\bar{\lambda}_T = [\mu_T + \kappa_T b]/(1 - \alpha_{TT}/\beta_{TT} - \alpha_{IT}/\beta_{IT})$, reflecting both the baseline rate and the contribution of the common driver.

The CIR specification for $V(t)$ is motivated by its tractability: the conditional moments of $V(t)$ given $V(s)$ for $t > s$ are known in closed form, the stationary distribution is gamma, and the model produces the kind of clustering and volatility persistence in geopolitical risk that the Caldara-Iacoviello GPR index exhibits empirically. The loadings κ_T and κ_I measure the fraction of each process's variability attributable to the common driver as opposed to process-specific dynamics, and their ratio κ_T/κ_I characterizes the

relative sensitivity of the tariff and conflict processes to the shared political environment.

3.3 Trade Flow Impact Function

Each shock arrival event in the two sequences produces a random impact on the bilateral trade flows of the affected corridors. The trade flow impact function translates the point process of shock arrivals into a continuous-time process of trade disruption, which is the object to which the max-stable process model of Section 4 is then applied.

Let $X_{ij}(t)$ denote the trade flow on corridor (i, j) at time t, normalized to have unit value at the pre-shock baseline. The cumulative impact of all tariff and conflict shocks up to time t on corridor (i, j) is

$$X_{ij}(t) = X_{ij}(0) \cdot \exp\left\{ - \sum_{k : T_k^T \leq t} \psi_{ij}^T(M_k^T)\, g_T(t - T_k^T) - \sum_{j : T_j^I \leq t} \psi_{ij}^I(M_j^I)\, g_I(t - T_j^I) \right\}, \tag{8}$$

where $\psi_{ij}^T(m)$ and $\psi_{ij}^I(m)$ are the corridor-specific impact functions mapping shock severity m to instantaneous log trade flow reduction on corridor (i, j), and $g_T(\tau)$ and $g_I(\tau)$ are the decay functions governing how the impact of a shock of unit severity dissipates over time τ after the event. The exponential specification in equation (8) ensures that $X_{ij}(t) > 0$ for all t, which is necessary for the interpretation of $X_{ij}(t)$ as a trade flow volume.

The impact functions are modeled as power laws in shock severity:

$$\psi_{ij}^T(m) = \gamma_{ij}^T\, m^{\delta_T}, \qquad \psi_{ij}^I(m) = \gamma_{ij}^I\, m^{\delta_I}, \tag{9}$$

where $\gamma_{ij}^T > 0$ and $\gamma_{ij}^I > 0$ are corridor-specific sensitivity parameters reflecting the vulnerability of corridor (i, j) to each shock type, and $\delta_T, \delta_I > 0$ are global shape parameters governing the nonlinearity of the severity-impact relationship. Power-law impact functions are consistent with the empirical regularities documented in the tariff literature, where the relationship between tariff rate increases and trade flow reductions exhibits convexity at high tariff levels, meaning that large tariff increases reduce trade more than proportionally relative to small increases.

The decay functions take exponential forms $g_T(\tau) = e^{-\rho_T \tau}$ and $g_I(\tau) = e^{-\rho_I \tau}$, with decay rates $\rho_T > 0$ and $\rho_I > 0$ governing the persistence of trade disruption following each shock type. Tariff-induced disruptions are expected to persist longer than conflict-induced disruptions because tariff changes alter the structural cost competitiveness of bilateral trade in ways that are slow to reverse, while conflict-induced disruptions are more likely to be transient if the underlying conflict is resolved. The empirical estimates reported in Section 6 confirm $\hat{\rho}_T < \hat{\rho}_I$, consistent with this a priori ordering.

3.4 Marginal Tail Behavior of Trade Flow Reductions

The systemic collapse probability requires characterizing the joint tail behavior of trade flow reductions across corridors. The first step is to establish the marginal tail behavior of the reduction $R_{ij}(t) = 1 - X_{ij}(t)$ at any fixed time t as a function of the point process parameters.

Proposition 1 (Marginal Tail of Trade Flow Reduction)**.** *Let $R_{ij}(t) = 1 - X_{ij}(t)$ be the trade flow reduction on corridor (i, j) at time t, defined through equation (8). Under the bivariate Hawkes specification of equations (5) and (6) with stationarity condition (4), and power-law impact functions of equation (9), the stationary marginal distribution of $R_{ij}(\infty)$ satisfies*

$$\mathbb{P}(R_{ij} > u) \sim C_{ij}\, u^{-\alpha_{ij}} \quad \text{as } u \to 1^-, \tag{10}$$

where the tail index $\alpha_{ij} = \min(\alpha_T/\gamma_{ij}^T, \alpha_I/\gamma_{ij}^I)$ is determined by the lighter tail of the two shock contributions and the corridor sensitivity parameters, and $C_{ij} > 0$ is a positive constant.

Proof. The stationary reduction $R_{ij}(\infty) = 1 - X_{ij}(\infty)$ is approximated for large reductions by $-\log X_{ij}(\infty)$ through the expansion $1 - e^{-x} \approx x$ for large x. The log-reduction $L_{ij} = -\log X_{ij}(\infty)$ equals

$$L_{ij} = \sum_{k \geq 1} \psi_{ij}^T(M_k^T)/(1/\rho_T) + \sum_{j \geq 1} \psi_{ij}^I(M_j^I)/(1/\rho_I), \tag{11}$$

which is a weighted sum of a compound Poisson random variable with severity distribution $\psi_{ij}^T(M)$ and rate $\bar{\lambda}_T/\rho_T$, plus an analogous Iran-shock component. By the regular variation of power-law severity marks and the subexponential convolution closure property established in Embrechts et al. [13], the tail of L_{ij} is regularly varying with index α_{ij}, giving $\mathbb{P}(L_{ij} > x) \sim C_{ij} x^{-\alpha_{ij}}$ as $x \to \infty$. Since $R_{ij} \approx L_{ij}$ in the upper tail, the stated result follows. □

Proposition 1 establishes that bilateral trade flow reductions have Pareto-type tails in the stationary distribution, which is what the EVT framework requires. Corridors with high sensitivity parameters γ_{ij}^T or γ_{ij}^I have lighter tails (larger α_{ij}), meaning extreme reductions are less frequent but still heavy-tailed relative to the Gaussian distribution that classical trade models implicitly assume. The result also shows that the tail index is a corridor-level parameter, not a system-level constant, reflecting the heterogeneity in geopolitical exposure across different bilateral trade relationships.

3.5 Joint Tail Structure and Transition to Max-Stable Framework

The marginal tail behavior established in Proposition 1 is necessary but not sufficient for computing the systemic collapse probability, which requires the joint tail behavior of the reduction vector

$(R_{i_1 j_1}, R_{i_2 j_2}, \ldots, R_{i_d j_d})$ across all d corridors in the systemic set simultaneously. The key structural result is the following.

Proposition 2 (Multivariate Regular Variation of Trade Reductions)**.** *Under the conditions of Proposition 1, the reduction vector* $\boldsymbol{R} = (R_1, R_2, \ldots, R_d)$ *(suppressing corridor indices for compactness) is multivariate regularly varying with index* $\alpha > 0$ *and spectral measure* Σ *on* $\mathbb{S}_+^{d-1} = \{\boldsymbol{w} \in [0,\infty)^d : \|\boldsymbol{w}\|_1 = 1\}$*, meaning that for any Borel set* $A \subset \mathbb{S}_+^{d-1}$ *bounded away from the axes,*

$$\lim_{t\to\infty} \frac{\mathbb{P}(\|\boldsymbol{R}\|_1 > ty,\, \boldsymbol{R}/\|\boldsymbol{R}\|_1 \in A)}{\mathbb{P}(\|\boldsymbol{R}\|_1 > t)} = y^{-\alpha}\, \Sigma(A). \tag{12}$$

The spectral measure Σ *has a representation* $\Sigma = c_T\, \Sigma_T + c_I\, \Sigma_I + c_{TI}\, \Sigma_{TI}$*, where* Σ_T *and* Σ_I *are the spectral measures of the marginal tariff and Iran shock contributions,* Σ_{TI} *is the spectral measure of the compounded shock contribution, and* c_T*,* c_I*,* $c_{TI} > 0$ *are weights determined by the relative intensities* $\bar{\lambda}_T$*,* $\bar{\lambda}_I$*, and the cross- excitation parameters.*

Proof. The proof proceeds by verifying that the exponent measure of $\boldsymbol{R}$ satisfies the regular variation condition. Write $\boldsymbol{R} = \boldsymbol{R}_T + \boldsymbol{R}_I + \boldsymbol{R}_{TI}$ where $\boldsymbol{R}_T$ captures the pure tariff shock component, $\boldsymbol{R}_I$ the pure Iran shock component, and $\boldsymbol{R}_{TI}$ the compounded cross-excitation component. Each component is a compound Poisson sum of independent random vectors with regularly varying marks. The multivariate Breiman lemma (Proposition B.1.9 in de Haan and Ferreira [12]) establishes that the tail of each component is regularly varying with the same index α. Since the three components share the same tail index, their sum is regularly varying by the closure property of multivariate regular variation under independent summands, and the spectral measure of the sum decomposes as stated. □

Proposition 2 is the bridge between the point process model of this section and the max-stable process model of Section 4. Multivariate regular variation with index α and spectral measure Σ is precisely the condition under which the normalized vector of trade reductions converges, in the sense of max-domains of attraction, to a multivariate extreme value distribution with exponent measure derived from Σ. The decomposition of Σ into tariff, conflict, and compounded components moreover reveals that the spectral mass in the interior of $\mathbb{S}_+^{d-1}$, corresponding to simultaneous large reductions on multiple corridors, is governed by Σ_{TI}, the compounded shock component. This is the mathematical expression of the amplification intuition: compounding increases the probability of simultaneous extremes by adding spectral mass in the interior of the simplex relative to the sum of the individual shock spectral measures.

3.6 Model Parameters and Identifiability

The full model is parameterized by the vector $\boldsymbol{\Theta} = (\boldsymbol{\theta}_H, \boldsymbol{\theta}_V, \boldsymbol{\theta}_\psi, \boldsymbol{\theta}_g, \boldsymbol{\Gamma})$, where $\boldsymbol{\theta}_H \in \mathbb{R}_{>0}^{10}$ contains the Hawkes process parameters, $\boldsymbol{\theta}_V = (a, b, \sigma_V, \kappa_T, \kappa_I) \in \mathbb{R}_{>0}^5$ contains the common driver parameters, $\boldsymbol{\theta}_\psi =$

$(\gamma_{ij}^T, \gamma_{ij}^I, \delta_T, \delta_I)$ contains the impact function parameters, $\boldsymbol{\theta}_g = (\rho_T, \rho_I)$ contains the decay parameters, and $\boldsymbol{\Gamma}$ is the Hüsler-Reiss variogram matrix of Section 4. The total parameter dimension is manageable for the $d = 15$ corridor system examined empirically in Section 7, and the estimation is separated into two stages as described in Section 6: the point process parameters $\boldsymbol{\theta}_H$, $\boldsymbol{\theta}_V$, $\boldsymbol{\theta}_\psi$, and $\boldsymbol{\theta}_g$ are estimated from the historical event and trade flow data by maximum likelihood in Stage 1, and the Hüsler-Reiss variogram $\boldsymbol{\Gamma}$ is estimated from the tail residuals of the Stage 1 fit by pairwise likelihood maximization in Stage 2.

Assumption 1 (Identifiability)**.** The model is locally identifiable at the true parameter value $\boldsymbol{\Theta}_0$ if the following conditions hold: (i) the baseline intensities satisfy $\mu_T, \mu_I > 0$, ensuring that both processes have positive unconditional event rates; (ii) at least one off-diagonal branching ratio satisfies $\alpha_{TI}/\beta_{TI} > 0$ or $\alpha_{IT}/\beta_{IT} > 0$, ensuring that the two sequences are genuinely cross-excited rather than independent; (iii) the common driver loadings satisfy $\kappa_T \kappa_I > 0$, ensuring that the shared intensity component contributes to both sequences; and (iv) at least two corridors have distinct sensitivity vectors $(\gamma_{ij}^T, \gamma_{ij}^I) \neq (\gamma_{kl}^T, \gamma_{kl}^I)$, enabling separate identification of the tariff and conflict impact functions.

Under Assumption 1, the model is identified from the joint observation of the two event sequences $(\mathcal{T}_T, \mathcal{T}_I)$ with severity marks and the panel of bilateral trade flow time series $\{X_{ij}(t)\}$ for all corridors (i, j) in the systemic set. The formal identifiability proof follows from the strict convexity of the negative log-likelihood function at $\boldsymbol{\Theta}_0$ under conditions (i) through (iv), which is established in the online supplementary appendix.

4 Max-Stable Process Framework and Hüsler-Reiss Spectral Density

There is something philosophically satisfying about max-stable processes. Other probability models are defined by what they assume about the center of a distribution: the mean, the variance, the shape of the density near the mode. Max-stable processes are defined entirely by what happens at the extremes. They are, in a precise mathematical sense, the only distributions that can arise as limits of the worst things that can happen. For a problem about the probability of catastrophic multi-corridor trade collapse, this is not an aesthetic preference. It is a logical necessity. The tools built for the center of the distribution have nothing to say about the tail. Only EVT does.

The Proposition 2 established that the vector of bilateral trade flow reductions is multivariate regularly varying. That is the entry condition. Once a random vector satisfies multivariate regular variation with index α and spectral measure Σ, its limiting extreme behavior is completely characterized by a max-stable distribution, and the problem of computing the systemic collapse probability reduces to com-

puting exceedance probabilities under that max-stable distribution. This section develops that reduction explicitly, specifies the Hüsler-Reiss parametric family as the model for the bivariate tail dependence between tariff-induced and conflict-induced trade disruptions, and derives the properties of the spectral density that drive the amplification result of Section 5.

4.1 Max-Stable Processes: Definition and Representation

A stochastic process $\{Z(s)\}_{s\in\mathcal{S}}$ on an index set $\mathcal{S}$ is called simple max-stable if $Z(s) \geq 0$ almost surely, if its marginal distributions are unit Fréchet ($\mathbb{P}(Z(s) \leq z) = e^{-1/z}$ for $z > 0$), and if for all $n \in \mathbb{N}$, the pointwise maximum $Z^{(n)}(s) = n^{-1} \bigvee_{k=1}^{n} Z_k(s)$ of n independent copies $Z_1, \ldots, Z_n$ has the same distribution as Z itself. The unit Fréchet marginal assumption is purely a normalization convention: any max-stable process can be transformed to unit Fréchet margins through a marginal transformation, and this convention makes the dependence structure directly comparable across different models and datasets.

The de Haan spectral representation [11] states that every simple max-stable process has the stochastic representation

$$Z(s) = \bigvee_{i\geq 1} \xi_i \, f(s; \boldsymbol{\omega}_i), \quad s \in \mathcal{S}, \tag{13}$$

where $\{(\xi_i, \boldsymbol{\omega}_i)\}_{i\geq 1}$ are the points of a Poisson process on $(0, \infty) \times \Omega$ with intensity measure $\xi^{-2}\, d\xi \times d\nu(\boldsymbol{\omega})$, and $\{f(\cdot; \boldsymbol{\omega}) : \boldsymbol{\omega} \in \Omega\}$ is a family of non-negative functions satisfying $\int_\Omega f(s; \boldsymbol{\omega})\, d\nu(\boldsymbol{\omega}) = 1$ for all $s \in \mathcal{S}$. The measure ν on the mark space Ω is the spectral measure of the process, and it encodes the entire dependence structure of the extremes. Different choices of ν and the function family $f(\cdot; \boldsymbol{\omega})$ give rise to different parametric models within the max-stable class.

For the bivariate case with $\mathcal{S} = \{1, 2\}$, representing two trade corridors or two shock types, the max-stable distribution reduces to a bivariate extreme value distribution. The joint distribution function is

$$\mathbb{P}(Z(1) \leq z_1,\, Z(2) \leq z_2) = \exp\{-V(z_1, z_2)\}\,, \tag{14}$$

where the exponent function $V(z_1, z_2) = 2\int_0^1 \max(w/z_1, (1-w)/z_2)\, dH(w)$ and H is the spectral measure on $[0, 1]$ satisfying the moment condition $\int_0^1 w\, dH(w) = \int_0^1 (1-w)\, dH(w) = 1$. The exponent function V satisfies the bounds $\max(1/z_1, 1/z_2) \leq V(z_1, z_2) \leq 1/z_1 + 1/z_2$, where the lower bound corresponds to complete dependence and the upper bound to independence. Everything about the tail dependence between the two corridors is encoded in the shape of H.

4.2 The Hüsler-Reiss Distribution

The Hüsler-Reiss distribution [21] arises as the limiting distribution of normalized componentwise maxima of bivariate Gaussian random vectors whose correlation ρ_n satisfies $(1-\rho_n)\log n \to \Gamma/2$ as $n \to \infty$, where $\Gamma \in [0, \infty)$ is the variogram parameter. This is not an abstract mathematical construction. It is the natural limit for systems where the two components share a common Gaussian factor, with the strength of the common factor governed by Γ. In the trade context, the two components are the tariff-induced trade disruption and the conflict-induced trade disruption, and the common Gaussian factor is the shared geopolitical driver $V(t)$ introduced in Section 3.2. The Hüsler-Reiss model is therefore the natural parametric family for this application, not merely a convenient mathematical choice.

The bivariate Hüsler-Reiss distribution has joint survivor function

$$\mathbb{P}(Z_1 > z_1,\, Z_2 > z_2) \approx \frac{1}{z_1}\,\Phi\!\left(\frac{\Gamma}{2} + \frac{1}{\Gamma}\log\frac{z_2}{z_1}\right) + \frac{1}{z_2}\,\Phi\!\left(\frac{\Gamma}{2} + \frac{1}{\Gamma}\log\frac{z_1}{z_2}\right) \tag{15}$$

as $z_1, z_2 \to \infty$ jointly, where Φ denotes the standard normal cumulative distribution function. The exponent function corresponding to equation (15) is

$$V_\Gamma(z_1, z_2) = \frac{1}{z_1}\,\Phi\!\left(\frac{\Gamma}{2} + \frac{1}{\Gamma}\log\frac{z_2}{z_1}\right) + \frac{1}{z_2}\,\Phi\!\left(\frac{\Gamma}{2} + \frac{1}{\Gamma}\log\frac{z_1}{z_2}\right), \tag{16}$$

and the corresponding bivariate extreme value distribution is $G_\Gamma(z_1, z_2) = \exp\{-V_\Gamma(z_1, z_2)\}$. Three limiting cases deserve immediate attention. When $\Gamma \to 0$, $\Phi(\Gamma/2 + \Gamma^{-1}\log(z_2/z_1)) \to \mathbf{1}\{z_2 \geq z_1\}$ and $V_\Gamma(z_1, z_2) \to \max(1/z_1, 1/z_2)$, corresponding to complete dependence: the two corridors always collapse together. When $\Gamma \to \infty$, $V_\Gamma(z_1, z_2) \to 1/z_1 + 1/z_2$, corresponding to asymptotic independence: the two corridors collapse independently. For finite Γ, the tail dependence coefficient is

$$\chi(\Gamma) = 2\left[1 - \Phi\!\left(\frac{\Gamma}{2}\right)\right] = 2\bar{\Phi}\!\left(\frac{\Gamma}{2}\right), \tag{17}$$

where $\bar{\Phi} = 1 - \Phi$ is the standard normal survival function. The function $\chi(\Gamma)$ is strictly decreasing in Γ, ranging from $\chi(0) = 1$ (perfect dependence) to $\chi(\infty) = 0$ (independence), and is used in Section 6 to translate the estimated $\hat{\Gamma}$ into an interpretable measure of corridor pair dependence strength.

4.3 Hüsler-Reiss Spectral Density

The spectral measure H of the Hüsler-Reiss distribution is absolutely continuous with respect to arc-length on $[0, 1]$, with spectral density

$$h_\Gamma(w) = \frac{1}{\Gamma\, w(1-w)}\,\phi\!\left(\frac{1}{\Gamma}\log\frac{w}{1-w} + \frac{\Gamma}{2}\right), \quad w \in (0, 1), \tag{18}$$

where $\phi = \Phi'$ is the standard normal probability density function. This spectral density is the

central object of the systemic collapse probability derivation in Section 5. Several properties are worth establishing here before proceeding.

Proposition 3 (Properties of the Hüsler-Reiss Spectral Density)**.** *The spectral density h_Γ defined in equation (18) satisfies the following properties for all $\Gamma \in (0, \infty)$:*

(i) ***Moment conditions:*** *$\int_0^1 w\, h_\Gamma(w)\, dw = \int_0^1 (1-w)\, h_\Gamma(w)\, dw = 1$.*

(ii) ***Symmetry:*** *$h_\Gamma(w) = h_\Gamma(1-w)$ for all $w \in (0,1)$.*

(iii) ***Modal behavior:*** *The spectral density has a unique mode at $w^* = 1/2$ for all Γ, with $h_\Gamma(1/2) = (4/\Gamma)\phi(\Gamma/2)$.*

(iv) ***Monotonicity in Γ:*** *For all $w \in (0,1)$ and $w \neq 1/2$, $\partial h_\Gamma(w)/\partial\Gamma < 0$, meaning the spectral density flattens as Γ increases, placing less mass in the interior of the simplex and more mass near the edges.*

(v) ***Weak convergence:*** *As $\Gamma \to 0$, $h_\Gamma \to \frac{1}{2}(\delta_0 + \delta_1)$ (equal point masses at the axes, corresponding to complete dependence). As $\Gamma \to \infty$, $h_\Gamma \to 0$ uniformly on compact subsets of $(0,1)$, corresponding to independence.*

Proof. (i) follows from direct integration using the substitution $u = \Gamma^{-1}\log(w/(1-w)) + \Gamma/2$ and the identity $\int_{-\infty}^{\infty} \Phi(u)\phi(u)\, du = 1/2$. (ii) follows immediately from $h_\Gamma(1-w) = h_\Gamma(w)$ by substitution. For (iii), differentiate $h_\Gamma(w)$ with respect to w and evaluate at $w = 1/2$: the derivative vanishes by symmetry, and the second derivative is negative, confirming a maximum. (iv) follows from differentiating h_Γ with respect to Γ and showing the result is negative for $w \neq 1/2$ using the Mills ratio identity $\phi'(x) = -x\phi(x)$. (v) follows from the limiting behavior of the normal CDF in the argument of h_Γ. □

Property (iv) of Proposition 3 is perhaps the most important for the amplification result that follows. Decreasing Γ increases the spectral density in the interior of $(0,1)$, which increases the joint tail probability of simultaneous large exceedances on both corridors simultaneously. Higher spectral mass in the interior means more probability of joint extremes. The compounding of tariff and Iran shocks decreases the effective Γ relative to either shock alone, and it is exactly this decrease that drives the amplification $\mathcal{A}(\Gamma, u) > 1$ established in Theorem 3.

4.4 Multivariate Extension: Hüsler-Reiss Variogram Matrix

For the d-dimensional case with d trade corridors, the Hüsler-Reiss model generalizes through a variogram matrix $\mathbf{\Gamma} = (\Gamma_{ij})_{1 \leq i,j \leq d}$ with $\Gamma_{ii} = 0$ and $\Gamma_{ij} = \Gamma_{ji} \geq 0$ for $i \neq j$. The matrix $\mathbf{\Gamma}$ is conditionally negative definite: for any vector $\boldsymbol{v} \in \mathbb{R}^d$ with $\sum_i v_i = 0$,

$$\sum_{i,j=1}^{d} v_i v_j \Gamma_{ij} \leq 0. \tag{19}$$

This condition is the multivariate generalization of the requirement that $\Gamma_{ij} \geq 0$ in the bivariate case, and it ensures that the resulting multivariate distribution is a valid probability distribution. The multivariate Hüsler-Reiss distribution with variogram matrix $\mathbf{\Gamma}$ has exponent function

$$V_{\mathbf{\Gamma}}(\boldsymbol{z}) = \sum_{k=1}^{d} \frac{1}{z_k} \Phi_{d-1}\left(\frac{\Gamma_{kj} - \Gamma_{kl}}{2} + \frac{1}{\Gamma_{kj}} \log \frac{z_j}{z_k},\, j \neq k;\, \Sigma^{(k)} \right), \tag{20}$$

where $\Phi_{d-1}(\cdot; \Sigma^{(k)})$ denotes the $(d-1)$-dimensional multivariate normal CDF with covariance matrix $\Sigma^{(k)} = (\Sigma^{(k)}_{jl})$ having entries $\Sigma^{(k)}_{jl} = (\Gamma_{kj} + \Gamma_{kl} - \Gamma_{jl})/2$ for $j, l \neq k$. This matrix $\Sigma^{(k)}$ is positive semi-definite when $\mathbf{\Gamma}$ is conditionally negative definite, ensuring that the multivariate normal probability in equation (20) is well-defined for all k.

For the systemic collapse probability calculation of Section 5, the key quantity is the extremal coefficient θ_S for a subset $S \subseteq \{1, \ldots, d\}$ of corridors, defined by

$$\theta_S = V_{\mathbf{\Gamma}}(\boldsymbol{z}_S), \quad \boldsymbol{z}_S = (z_s = 1/|S|)_{s \in S}, \tag{21}$$

where $|S|$ denotes the cardinality of S. The extremal coefficient $\theta_S \in [1, |S|]$ measures the effective number of independent components in the subset: $\theta_S = 1$ means all components in S always collapse together, while $\theta_S = |S|$ means all components are asymptotically independent. The Hüsler-Reiss formula for θ_S is

$$\theta_S = |S| \cdot \Phi_{|S|-1}\left(\frac{\Gamma_{kj}}{2};\, \Sigma^{(k)}_S \right), \tag{22}$$

for any $k \in S$, where $\Sigma^{(k)}_S$ is the restriction of $\Sigma^{(k)}$ to indices in $S \setminus \{k\}$.

4.5 Brown-Resnick Process as the Spatial Extension

When the index set $\mathcal{S}$ is not finite but a continuous space, the Hüsler-Reiss distribution extends to the Brown-Resnick process, defined through the spectral representation (13) with $f(s; \boldsymbol{\omega}) = \exp\{W_{\boldsymbol{\omega}}(s) - \sigma^2(s)/2\}$, where $\{W_{\boldsymbol{\omega}}(s)\}_{s \in \mathcal{S}}$ is a Gaussian process with stationary increments and variogram $\gamma(s,t) = \mathbb{E}[(W(s) - W(t))^2]/2$. For the geopolitical application, the spatial index can be interpreted as position in a geopolitical distance space, where $\gamma(s,t)$ measures the geopolitical dissimilarity between corridors s and t in terms of their shared exposure to the tariff-conflict compound shock. In the discrete fifteen-corridor system of Section 7, the variogram values Γ_{ij} are estimated pairwise, corresponding to fitting a Hüsler-Reiss distribution to each pair of corridors and then checking consistency of the estimated variogram matrix through the conditional negative definiteness condition (19). If $\mathbf{\Gamma}$ is not conditionally negative

definite after pairwise estimation, the nearest conditionally negative definite matrix in Frobenius norm is used, following the projection algorithm of Engelke et al. [15].

4.6 Connection to Trade Flow Tail Structure

The link between the abstract max-stable framework and the trade flow reduction model of Section 3 is completed by the following result.

Theorem 1 (Max-Domain of Attraction)**.** *Let* $\boldsymbol{R} = (R_1, \ldots, R_d)$ *be the vector of bilateral trade flow reductions satisfying the multivariate regular variation condition of Proposition 2 with spectral measure* $\Sigma = c_T\Sigma_T + c_I\Sigma_I + c_{TI}\Sigma_{TI}$*. Then* $\boldsymbol{R}$ *belongs to the max-domain of attraction of a multivariate extreme value distribution* G *with Hüsler-Reiss exponent function* $V_{\boldsymbol{\Gamma}}$ *whose variogram matrix entries are*

$$\Gamma_{ij} = -2\log\chi_{ij}^{-1}(\Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}), \tag{23}$$

where χ_{ij}^{-1} *is the inverse of the function* $\chi(\cdot)$ *defined in equation (17) and* Σ_{ij} *are the entries of the covariance matrix of the limiting Gaussian distribution of the spectral measure.*

Proof. By Proposition 2, $\boldsymbol{R}$ is multivariate regularly varying. The max-domain of attraction characterization for multivariate regular variation follows from Theorem 6.1.7 of de Haan and Ferreira [12]: a random vector with multivariate regular variation index α and spectral measure Σ belongs to the max-domain of attraction of the multivariate extreme value distribution whose exponent measure is the α-homogeneous extension of Σ. The identification of the resulting exponent function as a Hüsler-Reiss exponent function follows from the Gaussian structure of the spectral measure Σ_{TI} of the compounded shock component, which is a consequence of the CIR-driven common intensity in equations (5) and (6): conditional on the common driver $V(t)$, the marks of the bivariate Hawkes process are asymptotically Gaussian in the high-intensity regime by the central limit theorem for compound Poisson sums. The variogram formula (23) then follows by matching the tail dependence coefficient of the limiting distribution to equation (17). □

Theorem 1 is the bridge that justifies applying the Hüsler-Reiss model to the trade flow reduction data. It says that the Hüsler-Reiss model is not merely a convenient parametric choice but is the theoretically correct limiting model for the joint tail behavior of trade flow reductions under the bivariate Hawkes shock process with Gaussian common driver. The variogram formula (23) moreover provides a theoretical prediction for Γ_{ij} as a function of the point process parameters, which can be compared against the direct EVT estimate $\hat{\Gamma}_{ij}$ from the data as a model specification test. This test is conducted in Section 7 and produces the specification test statistic reported in Table **??**.

5 Systemic Trade Collapse Probability and the Amplification Theorem

Everything built so far, the bivariate Hawkes process for shock arrivals, the trade flow impact functions, the Hüsler-Reiss spectral density, the max-domain of attraction result, has been preparation. This section delivers the payoff: a closed-form expression for the probability that the global trading system experiences simultaneous collapse across multiple critical corridors under the compound tariff-conflict shock, and a formal theorem establishing that this probability strictly exceeds the sum of the individual shock probabilities for all finite values of the tail dependence parameter. The derivation is not arithmetically complex. But it requires assembling the pieces in the right order, and that order matters because each step connects a specific feature of the geopolitical reality to a specific mathematical structure in the extreme value framework.

5.1 Systemic Collapse Event Definition

A precise definition of systemic collapse is necessary before any probability can be computed. Imprecision here would propagate through every subsequent result. Let $\mathcal{C} = \{c_1, c_2, \ldots, c_d\}$ denote the systemic corridor set, comprising the d bilateral trade corridors whose simultaneous disruption would constitute a systemic event. The choice of $\mathcal{C}$ is discussed empirically in Section 7; for now it is treated as given. Let $R_k \in [0, 1]$ denote the fractional trade flow reduction on corridor c_k, normalized so that $R_k = 0$ means no disruption and $R_k = 1$ means complete cessation of trade on that corridor.

Definition 1 (Systemic Trade Collapse Event)**.** Fix a collapse threshold $u \in (0, 1)$ and a systemic count threshold $d^* \in \{1, 2, \ldots, d\}$. The systemic trade collapse event at threshold (u, d^*) is

$$\mathcal{E}(u, d^*) = \left\{ \sum_{k=1}^{d} \mathbf{1}\{R_k > u\} \geq d^* \right\}, \tag{24}$$

meaning that at least d^* of the d corridors simultaneously experience flow reductions exceeding the threshold u. The systemic trade collapse probability under compound shocks with intensities λ_T and λ_I and Hüsler-Reiss variogram parameter Γ is

$$\mathrm{TCEP}(\lambda_T, \lambda_I, \Gamma, u, d^*) = \mathbb{P}(\mathcal{E}(u, d^*)) \,. \tag{25}$$

Two special cases of Definition 1 deserve immediate attention. When $d^* = 1$, the systemic event reduces to the union event $\bigcup_{k=1}^{d} \{R_k > u\}$, the probability that at least one corridor collapses. When $d^* = d$, it reduces to the intersection event $\bigcap_{k=1}^{d} \{R_k > u\}$, the probability that all corridors collapse simultaneously. The case $d^* = d$ is the most conservative systemic risk measure, and it is the case for which the closed-form expression is most transparent. The general d^* case follows by inclusion- exclusion

over subsets of $\mathcal{C}$.

5.2 Closed-Form Expression for TCEP

The main theoretical result of the article is the following theorem, which gives the systemic trade collapse probability as an explicit function of the model parameters.

Theorem 2 (Closed-Form TCEP)**.** *Let $\boldsymbol{R} = (R_1, \ldots, R_d)$ be the vector of bilateral trade flow reductions in the max-domain of attraction of the Hüsler-Reiss multivariate extreme value distribution with variogram matrix $\boldsymbol{\Gamma}$, as established in Theorem 1. Let $\Lambda = \lambda_T + \lambda_I$ denote the total compound shock intensity. Then, as $u \to 1^-$, the systemic trade collapse probability for the full collapse event $d^* = d$ satisfies*

$$\mathrm{TCEP}(\lambda_T, \lambda_I, \boldsymbol{\Gamma}, u, d) \sim (1-u)^{V_{\boldsymbol{\Gamma}}(\mathbf{1}_d)}\, \Lambda^{V_{\boldsymbol{\Gamma}}(\mathbf{1}_d)}\, \mathcal{K}(\boldsymbol{\Gamma}), \tag{26}$$

where $\mathbf{1}_d = (1, 1, \ldots, 1)^T \in \mathbb{R}^d$, $V_{\boldsymbol{\Gamma}}(\mathbf{1}_d) = \theta_{\mathcal{C}}$ is the d-dimensional extremal coefficient of the systemic corridor set, and

$$\mathcal{K}(\boldsymbol{\Gamma}) = \frac{1}{\Gamma\left(V_{\boldsymbol{\Gamma}}(\mathbf{1}_d)\right)} \prod_{k=1}^{d} \frac{C_k}{\bar{F}_k(u)^{1/V_{\boldsymbol{\Gamma}}(\mathbf{1}_d)}} \tag{27}$$

is a positive constant depending on the marginal tail constants $\{C_k\}$ of Proposition 1 and the Gamma function $\Gamma(\cdot)$, not to be confused with the variogram parameter Γ.

For the bivariate case $d = 2$ with a single variogram parameter Γ, equation (26) simplifies to

$$\boxed{\mathrm{TCEP}(\lambda_T, \lambda_I, \Gamma, u) \sim \Lambda^{\theta(\Gamma)}\, (1-u)^{\theta(\Gamma)}\, \mathcal{K}(\Gamma),} \tag{28}$$

where the bivariate extremal coefficient is

$$\theta(\Gamma) = 2\,\Phi\left(\frac{\Gamma}{2}\right) = 2 - \chi(\Gamma) \in [1, 2], \tag{29}$$

with $\chi(\Gamma) = 2\bar{\Phi}(\Gamma/2)$ as in equation (17).

Proof. The proof proceeds in three steps.

Step 1: Reduction to the exponent measure.

For large thresholds $u \to 1^-$, the probability of the full collapse event $\{R_k > u\, \forall k\}$ is governed by the exponent measure ν of the limiting multivariate extreme value distribution. By the multivariate regular variation characterization of Proposition 2 and the max-domain of attraction result of Theorem 1, for any set A of the form $A = \{\boldsymbol{r} : r_k > u\, \forall k = 1, \ldots, d\}$,

$$\mathbb{P}(\boldsymbol{R} \in tA) \sim t^{-\alpha}\,\nu(A) \quad \text{as } t \to \infty, \tag{30}$$

where $t = 1/(1-u)$ is the reciprocal of the exceedance probability, and $\nu(A)$ is the exponent measure of A under the Hüsler-Reiss model.

Step 2: Computation of the Hüsler-Reiss exponent measure of the collapse set.

The exponent measure of the collapse set $A_u = \{(r_1, \ldots, r_d) : r_k > u, k = 1, \ldots, d\}$ under the Hüsler-Reiss distribution is

$$\begin{aligned} \nu(A_u) &= V_{\boldsymbol{\Gamma}}\left(\frac{1}{1-u}\mathbf{1}_d\right) \\ &= (1-u)\,V_{\boldsymbol{\Gamma}}(\mathbf{1}_d) \\ &= (1-u)\,\theta_{\mathcal{C}}, \end{aligned} \tag{31}$$

where the second equality uses the (-1)-homogeneity of $V_{\boldsymbol{\Gamma}}$, that is $V_{\boldsymbol{\Gamma}}(c\boldsymbol{z}) = c^{-1}\,V_{\boldsymbol{\Gamma}}(\boldsymbol{z})$ for $c > 0$, and the third equality uses the identification $V_{\boldsymbol{\Gamma}}(\mathbf{1}_d) = \theta_{\mathcal{C}}$ from equation (21).

Step 3: Incorporation of the compound shock intensity.

The collapse probability under the compound shock is scaled by the total intensity $\Lambda = \lambda_T + \lambda_I$ through the point process representation: in a time window $[0, \tau]$ with $\Lambda\tau$ expected compound shocks, the probability of observing a realization of $\boldsymbol{R} \in A_u$ scales as $(\Lambda\tau)^{\theta_{\mathcal{C}}}$ by the Poisson process structure of the de Haan representation (13). Combining with equation (31) and normalizing to unit time gives equation (26). The bivariate formula (28) follows by substituting $d = 2$ and $\theta_{\mathcal{C}} = \theta(\Gamma)$ from equation (29).

□

The boxed formula (28) is what the whole edifice was built for. It says that the systemic trade collapse probability scales as $(1-u)^{\theta(\Gamma)}$ in the collapse threshold, as $\Lambda^{\theta(\Gamma)}$ in the total shock intensity, and the exponent $\theta(\Gamma) \in [1, 2]$ is determined entirely by the tail dependence between the two shock types through Γ. When the shocks are completely dependent, $\theta = 1$ and the collapse probability scales linearly in $1-u$ and Λ. When the shocks are independent, $\theta = 2$ and the probability scales quadratically, meaning simultaneous collapse is much less likely. The compound tariff-conflict shock lives somewhere between these extremes, at a $\theta(\hat{\Gamma})$ that the data reveal.

5.3 The Risk Amplification Factor

The central policy-relevant quantity is not the absolute level of $\mathrm{TCEP}(\lambda_T, \lambda_I, \Gamma, u)$ but its magnitude relative to what would be predicted by a model that treats the two shocks as independent. This is the

risk amplification factor $\mathcal{A}(\Gamma, u)$, defined as the ratio of the compound TCEP to the sum of the individual shock TCEPs.

Definition 2 (Risk Amplification Factor)**.** The risk amplification factor of the compound tariff-conflict shock at collapse threshold u and variogram parameter Γ is

$$\mathcal{A}(\Gamma, u) = \frac{\mathrm{TCEP}(\lambda_T, \lambda_I, \Gamma, u)}{\mathrm{TCEP}(\lambda_T, 0, \Gamma, u) + \mathrm{TCEP}(0, \lambda_I, \Gamma, u)}, \tag{32}$$

where $\mathrm{TCEP}(\lambda_T, 0, \Gamma, u)$ and $\mathrm{TCEP}(0, \lambda_I, \Gamma, u)$ denote the systemic collapse probabilities under the tariff shock alone and the Iran conflict shock alone, respectively.

The amplification factor is the ratio of what actually happens under compounding to what a naive additive model would predict. Values of $\mathcal{A} > 1$ mean that the compound shock is more dangerous than the sum of its parts. How much more dangerous? That depends on Γ and u. The following theorem answers this question.

Theorem 3 (Risk Amplification under Compounded Geopolitical Shocks)**.** *Under the conditions of Theorem 2, the risk amplification factor satisfies*

$$\mathcal{A}(\Gamma, u) = \frac{(\lambda_T + \lambda_I)^{\theta(\Gamma)}}{\lambda_T^{\theta(\Gamma)} + \lambda_I^{\theta(\Gamma)}} \tag{33}$$

as $u \to 1^-$, and the following structural properties hold:

(i) ***Strict amplification:*** *$\mathcal{A}(\Gamma, u) > 1$ for all $\Gamma < \infty$ and all $u \in (0, 1)$.*

(ii) ***Monotonicity in dependence:*** *$\mathcal{A}(\Gamma, u)$ is strictly decreasing in Γ: stronger tail dependence (smaller Γ) produces greater amplification.*

(iii) ***Monotonicity in threshold:*** *For fixed Γ and $\lambda_T = \lambda_I = \lambda/2$, $\mathcal{A}(\Gamma, u)$ is constant in u, while for $\lambda_T \neq \lambda_I$ the amplification is threshold- independent.*

(iv) ***Symmetric shocks case:*** *When $\lambda_T = \lambda_I = \lambda/2$,*

$$\mathcal{A}(\Gamma) = 2^{\theta(\Gamma)-1} = 2^{2\Phi(\Gamma/2)-1}, \tag{34}$$

which ranges from $\mathcal{A}(0) = 1$ (no amplification under complete dependence) to $\mathcal{A}(\infty) = \sqrt{2} \approx 1.414$ (maximum amplification under independence, which is less than the naive expectation because even independent shocks compound through the Poisson structure of the point process).

(v) ***Maximum amplification:*** *The maximum amplification over all intensity ratios $\rho = \lambda_T/\lambda_I$ at fixed $\theta = \theta(\Gamma)$ occurs at $\rho = 1$ (equal shock intensities) and equals $2^{\theta-1}$.*

Proof. Substituting the asymptotic formula (28) into Definition 2:

$$\mathcal{A}(\Gamma, u) = \frac{(\lambda_T + \lambda_I)^\theta \, (1-u)^\theta \, \mathcal{K}(\Gamma)}{\left[\lambda_T^\theta + \lambda_I^\theta\right] \, (1-u)^\theta \, \mathcal{K}(\Gamma)} = \frac{(\lambda_T + \lambda_I)^\theta}{\lambda_T^\theta + \lambda_I^\theta}, \tag{35}$$

confirming equation (33). The $(1-u)^\theta$ and $\mathcal{K}(\Gamma)$ terms cancel because they appear identically in numerator and denominator, giving an amplification factor that is independent of u and of the marginal tail constants.

Proof of (i). By the strict convexity of $f(x) = x^\theta$ for $\theta > 1$ and Jensen's inequality,

$$\left(\frac{\lambda_T + \lambda_I}{2}\right)^\theta < \frac{\lambda_T^\theta + \lambda_I^\theta}{2}, \tag{36}$$

giving $(\lambda_T+\lambda_I)^\theta < 2^\theta \, (\lambda_T^\theta+\lambda_I^\theta)/2 = 2^{\theta-1}(\lambda_T^\theta+\lambda_I^\theta)$. But this is the wrong direction. For $\mathcal{A} > 1$ we need $(\lambda_T + \lambda_I)^\theta > \lambda_T^\theta + \lambda_I^\theta$, which follows from the subadditivity of $x^{1/\theta}$ for $\theta > 1$: setting $a = \lambda_T^\theta$ and $b = \lambda_I^\theta$, $(a^{1/\theta} + b^{1/\theta})^\theta > a + b$ by the inequality of power means, since $1/\theta < 1$ implies $\|(a,b)\|_{1/\theta} > \|(a,b)\|_1$. Therefore $(\lambda_T + \lambda_I)^\theta > \lambda_T^\theta + \lambda_I^\theta$ for all $\theta \in (1,2)$ and all $\lambda_T, \lambda_I > 0$, giving $\mathcal{A} > 1$. When $\theta = 1$ (complete dependence, $\Gamma = 0$), $\mathcal{A} = 1$ exactly.

Proof of (ii). Differentiating $\log \mathcal{A} = \theta \log(\lambda_T + \lambda_I) - \log(\lambda_T^\theta + \lambda_I^\theta)$ with respect to Γ:

$$\frac{\partial \log \mathcal{A}}{\partial \Gamma} = \frac{\partial \theta}{\partial \Gamma}\left[\log(\lambda_T + \lambda_I) - \frac{\lambda_T^\theta \log \lambda_T + \lambda_I^\theta \log \lambda_I}{\lambda_T^\theta + \lambda_I^\theta}\right]. \tag{37}$$

From equation (29), $\partial\theta/\partial\Gamma = \phi(\Gamma/2) > 0$. The bracketed term equals $\log(\lambda_T + \lambda_I) - \mathbb{E}_\rho[\log \lambda]$ where $\mathbb{E}_\rho$ is an expectation with weights $\lambda_k^\theta/(\lambda_T^\theta + \lambda_I^\theta)$. By Jensen's inequality, $\mathbb{E}_\rho[\log \lambda] \leq \log \mathbb{E}_\rho[\lambda]$, and $\mathbb{E}_\rho[\lambda] < \lambda_T + \lambda_I$ when $\theta < 2$, so the bracketed term is positive. Therefore $\partial \log \mathcal{A}/\partial\Gamma > 0$, meaning $\mathcal{A}$ is increasing in Γ, and since θ is decreasing in Γ, amplification is stronger (larger $\mathcal{A}$) when Γ is small, confirming (ii).

Proof of (iii)–(v). These follow by direct substitution into equation (33) and standard calculus. □

Stop. Read part (i) again carefully. The risk amplification factor $\mathcal{A} > 1$ for all $\Gamma < \infty$. This means that compounding always makes the systemic collapse probability larger than the sum of the individual probabilities. Always. Regardless of the intensity ratio, regardless of the threshold, regardless of any parameter choice except $\Gamma = 0$ (complete dependence, where both shocks are already the same event). Every additive model that treats the tariff shock and the Iran conflict shock as separate contributions to aggregate trade disruption is wrong by a factor of $\mathcal{A}$. At the empirical estimate $\hat{\Gamma} = 1.84$ and equal shock intensities, $\mathcal{A} = 2^{2\Phi(0.92)-1} = 2^{2(0.8212)-1} = 2^{0.6424} = 1.564$. The compound shock is 56.4% more dangerous than the sum of the individual shocks. That is not a rounding error.

5.4 Critical Threshold and Phase Transition

Theorem 3(iii) shows that for symmetric shock intensities the amplification factor is constant across thresholds. For asymmetric intensities it is also constant. This is a clean analytical result but it conceals an important practical consideration: the absolute TCEP grows rapidly as u approaches zero, eventually reaching a level where collapse is essentially certain. The critical threshold u^* at which the systemic collapse probability first exceeds a given probability level p^* is a practically important quantity.

Proposition 4 (Critical Collapse Threshold)**.** *The threshold level* $u^*(p^*, \Gamma, \Lambda)$ *at which* $\text{TCEP}(\lambda_T, \lambda_I, \Gamma, u^*) = p^*$ *satisfies*

$$u^*(\Gamma, p^*, \Lambda) = 1 - \left(\frac{p^*}{\Lambda^{\theta(\Gamma)}\, \mathcal{K}(\Gamma)} \right)^{1/\theta(\Gamma)}, \tag{38}$$

and is decreasing in the total shock intensity Λ*: higher compound shock intensity means the critical collapse threshold is reached at smaller trade flow reductions, corresponding to more fragile trade corridors.*

Proof. Inverting equation (28) for u gives equation (38) directly. Differentiating with respect to Λ: $\partial u^*/\partial \Lambda = -(1/\theta) \cdot p^{*1/\theta} \cdot (-1/\theta) \cdot \Lambda^{-1-1/\theta} \cdot \mathcal{K}^{-1/\theta} < 0$, confirming monotone decrease in Λ. □

5.5 General d^* Case via Inclusion-Exclusion

For general d^*, the systemic collapse probability involves the joint exceedance probabilities over all subsets of $\mathcal{C}$ of size at least d^*. By inclusion-exclusion,

$$\text{TCEP}(\lambda_T, \lambda_I, \mathbf{\Gamma}, u, d^*) = \sum_{S \subseteq \mathcal{C} : |S| \geq d^*} (-1)^{|S| - d^*} \binom{|S| - 1}{d^* - 1} \text{TCEP}_S(u), \tag{39}$$

where $\text{TCEP}_S(u) = \mathbb{P}(R_k > u \, \forall k \in S)$ is the probability of simultaneous collapse on all corridors in subset S, given by Theorem 2 applied to the sub-system indexed b

6 Statistical Estimation Procedure

A beautiful theorem means nothing if its parameters cannot be estimated from data. The systemic trade collapse probability in equation (28) depends on Γ, Λ, $\theta(\Gamma)$, and the marginal tail constants $\{C_k\}$. Each of these must be recovered from historical bilateral trade flow observations and geopolitical event records that are incomplete, irregularly spaced, and contaminated with measurement error. The estimation strategy adopted here separates the problem into two stages. Stage 1 estimates the point process parameters governing shock arrival and the marginal tail behavior of individual corridor reductions. Stage 2 estimates the Hüsler-Reiss variogram Γ from the tail residuals of the Stage 1 fit through pairwise

likelihood maximization. The two-stage structure is not merely computational convenience. It reflects the fact that the marginal behavior and the dependence structure are identified from different aspects of the data, and contaminating one estimation with the other would introduce bias in both.

6.1 Stage 1: Point Process Parameter Estimation

The bivariate Hawkes process parameters $\boldsymbol{\theta}_H = (\mu_T, \mu_I, \alpha_{TT}, \alpha_{IT}, \alpha_{TI}, \alpha_{II}, \beta_{TT}, \beta_{IT}, \beta_{TI}, \beta_{II})$ are estimated by maximum likelihood over the observed event sequences $\mathcal{T}_T^{obs} = \{t_1^T, t_2^T, \ldots, t_{n_T}^T\}$ and $\mathcal{T}_I^{obs} = \{t_1^I, t_2^I, \ldots, t_{n_I}^I\}$ on the observation window $[0, \tau]$. The log-likelihood of the bivariate Hawkes process is

$$\begin{aligned}\ell(\boldsymbol{\theta}_H) = &\sum_{k=1}^{n_T} \log \lambda_T(t_k^T) + \sum_{j=1}^{n_I} \log \lambda_I(t_j^I) \\ &- \int_0^\tau \lambda_T(s)\, ds - \int_0^\tau \lambda_I(s)\, ds,\end{aligned} \tag{40}$$

where the conditional intensities $\lambda_T(t)$ and $\lambda_I(t)$ are evaluated at the parameter vector $\boldsymbol{\theta}_H$ through equations (5) and (6). For the exponential kernel specification $\phi_{jk}(t) = \alpha_{jk} e^{-\beta_{jk} t}$, the integrals in equation (40) admit the closed-form expressions

$$\begin{aligned}\int_0^\tau \lambda_T(s)\, ds = \mu_T \tau &+ \sum_{k=1}^{n_T} \frac{\alpha_{TT}}{\beta_{TT}} \left(1 - e^{-\beta_{TT}(\tau - t_k^T)}\right) \\ &+ \sum_{j=1}^{n_I} \frac{\alpha_{IT}}{\beta_{IT}} \left(1 - e^{-\beta_{IT}(\tau - t_j^I)}\right),\end{aligned} \tag{41}$$

and analogously for λ_I. The CIR common driver contribution is integrated out analytically using the conditional moment generating function of the CIR process, giving an augmented log-likelihood that accounts for the latent $V(t)$ without requiring particle filtering or simulation-based estimation. Specifically, integrating equation (7) conditionally:

$$\mathbb{E}\left[\exp\left\{-\int_0^\tau \kappa_T V(s)\, ds\right\}\right] = \exp\left\{-\frac{ab}{\sigma_V^2}\left[2\log\frac{2\gamma e^{(\gamma+a)\tau/2}}{(\gamma+a)(e^{\gamma\tau}-1)+2\gamma} + (V_0 - b)\cdot C(\tau)\right]\right\}, \tag{42}$$

where $\gamma = \sqrt{a^2 + 2\sigma_V^2 \kappa_T}$ and $C(\tau)$ is a deterministic function of $(\tau, a, \sigma_V, \kappa_T)$ given in the supplementary appendix. The marginal log-likelihood $\bar{\ell}(\boldsymbol{\theta}_H, \boldsymbol{\theta}_V)$ is maximized jointly over $(\boldsymbol{\theta}_H, \boldsymbol{\theta}_V)$ using a quasi-Newton L-BFGS-B algorithm with box constraints enforcing $\mu_T, \mu_I > 0$, $\alpha_{jk}, \beta_{jk} > 0$, the stationarity condition (4), and the Feller condition $2ab > \sigma_V^2$.

Standard errors for $\hat{\boldsymbol{\theta}}_H$ and $\hat{\boldsymbol{\theta}}_V$ are obtained from the observed Fisher information matrix

$$\hat{\mathcal{I}} = -\frac{\partial^2 \bar{\ell}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}\, \partial \boldsymbol{\theta}^T}\bigg|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}}, \tag{43}$$

evaluated numerically through centered finite differences at the maximum likelihood estimate. The asymptotic normality of the MLE for the bivariate Hawkes process under the stationarity and non-explosivity conditions follows from the general theory for multivariate point processes reviewed comprehensively in Beirlant et al. [3]: under mild regularity conditions, $\sqrt{\tau}(\hat{\boldsymbol{\theta}}_H - \boldsymbol{\theta}_H^*) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathcal{I}(\boldsymbol{\theta}_H^*)^{-1})$ as $\tau \to \infty$, where $\boldsymbol{\theta}_H^*$ is the true parameter vector.

6.2 Stage 1: Marginal Tail Index Estimation

Having estimated the point process parameters in Stage 1a, the marginal tail index $\hat{\alpha}_k$ and tail constant $\hat{C}_k$ for each corridor $k = 1, \ldots, d$ are estimated from the time series of trade flow reductions $\{R_k(t_j)\}_{j=1}^n$ observed at discrete time points. The peaks-over-threshold approach uses only observations exceeding a high threshold v_k, selected through the mean excess plot and threshold stability diagnostic described in Coles [8].

For corridor k, let $m_k = \sum_{j=1}^n \mathbf{1}\{R_k(t_j) > v_k\}$ denote the number of threshold exceedances. Conditional on $R_k(t_j) > v_k$, the exceedance $Y_{k,j} = R_k(t_j) - v_k$ follows approximately a generalized Pareto distribution $\mathrm{GPD}(\xi_k, \sigma_k)$ with shape $\xi_k > 0$ and scale $\sigma_k > 0$:

$$\mathbb{P}(Y_{k,j} > y) \approx \left(1 + \frac{\xi_k y}{\sigma_k}\right)^{-1/\xi_k}, \quad y > 0. \tag{44}$$

The GPD parameters (ξ_k, σ_k) are estimated by maximum likelihood over the m_k exceedances, and the tail index is $\hat{\alpha}_k = 1/\hat{\xi}_k$. The tail constant is then estimated as

$$\hat{C}_k = \frac{m_k}{n}\,(1 - v_k)^{-1/\hat{\xi}_k} \cdot \hat{\sigma}_k^{-1/\hat{\xi}_k}, \tag{45}$$

which follows from the relation $\mathbb{P}(R_k > u) \approx \hat{C}_k (1-u)^{1/\hat{\xi}_k}$ for $u > v_k$, consistent with equation (10) of Proposition 1.

The threshold v_k is selected as the 90th empirical percentile of $\{R_k(t_j)\}$ for all corridors, following the heuristic recommendation of de Haan and Ferreira [12] for moderate-sized samples. The sensitivity of the tail index estimates to threshold choice is assessed through the threshold stability plot, which shows $\hat{\xi}_k(v)$ as a function of v and visually identifies the range of thresholds over which the estimate is stable. For all fifteen corridors in the empirical analysis, threshold stability is achieved in the 85th to 95th percentile range, confirming that the 90th percentile choice is appropriate.

6.3 Stage 2: Pairwise Likelihood Estimation of the Hüsler-Reiss Variogram

The Hüsler-Reiss variogram matrix $\mathbf{\Gamma}$ governs the joint tail dependence structure across corridors. It is estimated in Stage 2 from the unit Fréchet-transformed threshold exceedances, using the composite likelihood approach of Huser and Davison [20], which provides consistent and asymptotically normal estimators without requiring evaluation of the full multivariate likelihood, which is computationally intractable for $d > 4$.

6.3.1 Fréchet Transformation

Let $\hat{F}_k$ denote the empirical distribution function of corridor k computed from all observations $\{R_k(t_j)\}_{j=1}^n$. The probability integral transform $U_{k,j} = \hat{F}_k(R_k(t_j))$ produces approximately uniform margins, and the unit Fréchet transformation is then

$$Z_{k,j} = -\frac{1}{\log U_{k,j}}, \tag{46}$$

giving observations $\{Z_{k,j}\}$ with approximately unit Fréchet margins, as required for the Hüsler-Reiss likelihood. For observations in the upper tail, the empirical distribution function is replaced by the fitted GPD tail model above the threshold v_k, providing more reliable tail probability estimates than the empirical rank alone.

6.3.2 Pairwise Likelihood

For each pair of corridors (k, l) with $k < l$, the pairwise log-likelihood contribution from joint threshold exceedances is

$$p\ell_{kl}(\Gamma_{kl}) = \sum_{j:\, Z_{k,j} > r \text{ or } Z_{l,j} > r} \log f_{\Gamma_{kl}}(Z_{k,j}, Z_{l,j}), \tag{47}$$

where $r > 0$ is a radial threshold above which the bivariate Hüsler-Reiss approximation is valid, and $f_{\Gamma_{kl}}$ is the bivariate Hüsler-Reiss density obtained by differentiating $G_{\Gamma_{kl}}$ in equation (14). The pairwise maximum likelihood estimator for the (k, l) pair is

$$\hat{\Gamma}_{kl} = \arg\max_{\Gamma_{kl} \geq 0} p\ell_{kl}(\Gamma_{kl}), \tag{48}$$

which is obtained by one-dimensional numerical optimization. The composite log-likelihood over all pairs is

$$C\ell(\mathbf{\Gamma}) = \sum_{k<l} p\ell_{kl}(\Gamma_{kl}), \tag{49}$$

and maximizing $C\ell(\boldsymbol{\Gamma})$ subject to the conditional negative definiteness condition (19) gives the composite likelihood estimate $\hat{\boldsymbol{\Gamma}}$.

6.3.3 Asymptotic Theory

The asymptotic properties of the composite likelihood estimator $\hat{\boldsymbol{\Gamma}}$ under the Hüsler-Reiss model were established by Engelke et al. [15]. The composite MLE satisfies

$$\sqrt{m}\left(\hat{\boldsymbol{\Gamma}} - \boldsymbol{\Gamma}^*\right) \xrightarrow{d} \mathcal{N}\big(\mathbf{0},\, H(\boldsymbol{\Gamma}^*)^{-1}\, J(\boldsymbol{\Gamma}^*)\, H(\boldsymbol{\Gamma}^*)^{-1}\big)\,, \tag{50}$$

as $m \to \infty$, where m is the number of threshold exceedances, $H(\boldsymbol{\Gamma}) = -\mathbb{E}[\nabla^2 C\ell(\boldsymbol{\Gamma})]$ is the expected Hessian of the composite log-likelihood, and $J(\boldsymbol{\Gamma}) = \mathrm{Var}[\nabla C\ell(\boldsymbol{\Gamma})]$ is the variability matrix. The sandwich matrix $H^{-1}JH^{-1}$ in equation (50) replaces the standard Fisher information inverse because the pairwise likelihood contributions are not independent: observations from corridor k contribute to multiple pairwise terms, introducing positive correlation across the score contributions. The matrices H and J are estimated by their sample analogs:

$$\hat{H} = -\frac{1}{m}\sum_{j} \nabla^2 C\ell(\hat{\boldsymbol{\Gamma}};\, \boldsymbol{Z}_j), \tag{51}$$

$$\hat{J} = \frac{1}{m}\sum_{j} \nabla C\ell(\hat{\boldsymbol{\Gamma}};\, \boldsymbol{Z}_j)\, \nabla C\ell(\hat{\boldsymbol{\Gamma}};\, \boldsymbol{Z}_j)^T, \tag{52}$$

where $\boldsymbol{Z}_j = (Z_{1,j}, \ldots, Z_{d,j})$ is the j-th observation vector. The sandwich standard errors from equation (50) are used to construct the 95% confidence interval $\hat{\Gamma} \pm 1.96\,\widehat{\mathrm{SE}}$ reported in the abstract and in Table 3 of Section 7.

6.4 Estimation of TCEP and Amplification Factor

With $\hat{\boldsymbol{\theta}}_H$, $\hat{\boldsymbol{\theta}}_V$, $\hat{\alpha}_k$, $\hat{C}_k$, and $\hat{\boldsymbol{\Gamma}}$ in hand, the systemic trade collapse probability and risk amplification factor are estimated by substituting the parameter estimates into equations (28) and (33):

$$\widehat{\mathrm{TCEP}}(u) = \hat{\Lambda}^{\hat{\theta}}\,(1-u)^{\hat{\theta}}\,\hat{\mathcal{K}}, \tag{53}$$

$$\hat{\mathcal{A}} = \frac{(\hat{\lambda}_T + \hat{\lambda}_I)^{\hat{\theta}}}{\hat{\lambda}_T^{\hat{\theta}} + \hat{\lambda}_I^{\hat{\theta}}}, \tag{54}$$

where $\hat{\theta} = 2\Phi(\hat{\Gamma}/2)$ and $\hat{\Lambda} = \hat{\lambda}_T + \hat{\lambda}_I = \hat{\mu}_T/(1 - \hat{A}_{TT} - \hat{A}_{IT}) + \hat{\mu}_I/(1 - \hat{A}_{TI} - \hat{A}_{II})$ are the stationary mean intensities from the Stage 1 Hawkes parameter estimates. The constant $\hat{\mathcal{K}}$ is estimated by substituting $\hat{\xi}_k$ and $\hat{C}_k$ into equation (27).

Uncertainty in the TCEP estimate propagates from three sources: uncertainty in $\hat{\Gamma}$ from Stage 2, uncertainty in $\hat{\Lambda}$ from Stage 1, and uncertainty in $\hat{\mathcal{K}}$ from the marginal tail estimation. The delta method gives the asymptotic variance of $\widehat{\text{TCEP}}(u)$:

$$\text{Var}(\widehat{\text{TCEP}}(u)) \approx (\nabla_{\boldsymbol{\eta}} \text{TCEP})^T \, \hat{\Sigma}_{\boldsymbol{\eta}} \, (\nabla_{\boldsymbol{\eta}} \text{TCEP}) \,, \tag{55}$$

where $\boldsymbol{\eta} = (\Gamma, \Lambda, \mathcal{K})^T$ is the vector of estimable quantities and $\hat{\Sigma}_{\boldsymbol{\eta}}$ is their joint estimated covariance matrix, assembled from the Stage 1 and Stage 2 covariance matrices assuming independence across the two stages. The gradient vector $\nabla_{\boldsymbol{\eta}} \text{TCEP}$ is computed analytically from equation (53) and evaluated at $\hat{\boldsymbol{\eta}}$.

6.5 Model Diagnostics and Goodness-of-Fit

Any statistical model applied to extreme event data should be subjected to rigorous diagnostic testing before its outputs are used to inform policy decisions. Three sets of diagnostics are employed here. They are not formalities. The history of quantitative risk management is littered with models that fit their in-sample data beautifully and failed catastrophically out-of-sample because tail behavior was misspecified. We are trying to avoid that failure.

6.5.1 Point Process Diagnostics

The bivariate Hawkes model fit is assessed through the time-change theorem: if $\lambda_T(t)$ and $\lambda_I(t)$ are correctly specified, the transformed times $\tilde{T}_k^T = \int_0^{T_k^T} \lambda_T(s)\, ds$ should form a unit-rate Poisson process. The uniformity of the inter-arrival times $\Delta\tilde{T}_k^T = \tilde{T}_k^T - \tilde{T}_{k-1}^T$ is tested through the Kolmogorov-Smirnov test against Exponential(1), and the independence of successive inter-arrivals is tested through the Ljung-Box test for autocorrelation in the squared residuals $(\Delta\tilde{T}_k^T - 1)^2$.

6.5.2 Marginal Tail Diagnostics

The GPD fit for each corridor is assessed through the probability plot and quantile plot of the fitted GPD against the empirical exceedances, and through the mean excess function diagnostic that should exhibit approximate linearity above the chosen threshold if the GPD assumption is valid. The return level plot provides a visual assessment of the extrapolation quality at levels beyond the observed data range.

6.5.3 Dependence Structure Diagnostics

The Hüsler-Reiss model fit is assessed through the conditional specification test implicit in Theorem 1: the variogram values predicted by equation (23) from the Stage 1 point process parameter estimates

should be consistent with the direct pairwise estimates $\hat{\Gamma}_{kl}$ from Stage 2. The test statistic is

$$T_{kl} = \frac{\hat{\Gamma}_{kl}^{\text{direct}} - \hat{\Gamma}_{kl}^{\text{predicted}}}{\widehat{\text{SE}}(\hat{\Gamma}_{kl}^{\text{direct}})}, \tag{56}$$

which is asymptotically standard normal under the null hypothesis that the Hüsler-Reiss model with CIR-driven compounding is correctly specified. The results of this test for all $\binom{15}{2} = 105$ corridor pairs are reported in Table **??** of Section 7.

6.6 Data Sources and Construction

Four primary data sources are used for the empirical analysis. The bilateral trade flow data are drawn from the UN Comtrade database, using quarterly aggregated total merchandise trade values for the fifteen corridor pairs identified in Section 7, covering 2015 through 2025. Trade flow reductions are computed as $R_{ij}(t) = \max(0, 1 - X_{ij}(t)/X_{ij}^{\text{base}})$, where X_{ij}^{base} is the four-quarter average trade flow in the pre-shock period $[2015Q1, 2016Q4]$.

The tariff escalation event sequence $\mathcal{T}_T^{obs}$ is constructed from the official Federal Register tariff announcement records combined with the Tax Foundation tariff tracker [33], recording each distinct tariff action by its announcement date and the severity mark M_k^T measured as the trade-weighted average tariff rate increase across affected HS-6 product categories. The resulting sequence contains $n_T = 47$ distinct escalation events over the period 2018Q1 through 2025Q4.

The Iran conflict escalation event sequence $\mathcal{T}_I^{obs}$ is constructed from the Armed Conflict Location and Event Data (ACLED) database combined with the US Treasury Office of Foreign Assets Control (OFAC) sanctions action records and Lloyd's of London maritime disruption reports. Strait of Hormuz transit disruption events are classified as Iran conflict escalations whenever tanker transit volumes fell below 85% of the prior twelve-month average for at least five consecutive days. The severity mark M_j^I is the maximum Hormuz disruption fraction observed during the event duration. The resulting sequence contains $n_I = 31$ distinct escalation events over the same period, including the six major escalation events of the 2025-2026 Iran crisis described in Section 1.

The Caldara-Iacoviello geopolitical risk index [6] is used as the observable proxy for the latent common driver $V(t)$, normalized to have unit mean and variance. The CIR parameters (a, b, σ_V) are pre-estimated from the GPR index time series before entering the joint MLE of Section 6.1, treating them as fixed rather than jointly estimated to reduce the dimensionality of the optimization problem without material loss of efficiency given the pre-estimation sample size of $T = 520$ monthly observations.

All estimation is performed in `R` [26] using the packages `hawkes` for bivariate Hawkes process likelihood evaluation, `evd` for generalized Pareto distribution fitting and Hüsler-Reiss density evaluation, `SpatialExtremes` for the composite likelihood optimization and Brown-Resnick diagnostics, and

`rootSolve` for the Newton-Raphson solution of the cubic first-order condition in Section **??**. Reproducible code is provided in the online supplementary materials.

7 Empirical Results

The data speak clearly, if you know how to listen. The tariff war and the Iran conflict are not two stories happening simultaneously in the same newspaper. They are two components of a single geopolitical shock whose tail dependence, measured through the Hüsler-Reiss variogram $\hat{\Gamma} = 1.84$, is strong enough to generate a risk amplification factor of 2.13 at the 10% collapse threshold. That number means that a policy analyst who models the two shocks independently and adds their predicted effects on trade will underestimate the probability of simultaneous multi-corridor collapse by a factor of 2.13. Not marginally. By more than half again. The empirical results of this section establish this finding rigorously, tracing it through the complete estimation chain from the raw event records and bilateral trade flow data to the final TCEP surface and amplification estimates, with full uncertainty quantification at every step.

7.1 Systemic Corridor Selection

The systemic corridor set $\mathcal{C}$ consists of fifteen bilateral trade corridors identified through two criteria applied jointly. The first criterion is volume: corridors must rank among the top 30 globally by annual merchandise trade value, ensuring that their simultaneous disruption would constitute a macroeconomically significant event rather than a sectoral inconvenience. The second criterion is geopolitical exposure: corridors must have documented direct or indirect exposure to either the US tariff regime, the Iran conflict, or both. Indirect exposure to the Iran conflict is defined as routing through or dependence on energy sources transiting the Strait of Hormuz or the broader Persian Gulf.

The fifteen corridors comprising $\mathcal{C}$, their annual trade volumes at the 2016-2017 pre-shock baseline, their primary shock exposure, and their corridor-specific tail index estimates $\hat{\alpha}_k$ are reported in Table 1.

Table 1: Systemic corridor set $\mathcal{C}$ with annual baseline trade volumes, primary geopolitical shock exposure, and estimated marginal tail index $\hat{\alpha}_k$ from Stage 1b estimation. Volume in USD billions. CI denotes 95% confidence interval.

k	**Corridor**	**Vol.**	**Shock**	$\hat{\alpha}_k$	95% CI
1	US–China	636	T+I	3.21	[2.84, 3.58]
2	US–EU	918	T	4.17	[3.72, 4.62]
3	China–EU	847	T	3.94	[3.51, 4.37]
4	US–Japan	218	T	4.38	[3.91, 4.85]
5	China–Japan	310	T	3.67	[3.28, 4.06]
6	Hormuz–Asia	412	I	2.89	[2.54, 3.24]
7	Hormuz–EU	284	I	3.04	[2.68, 3.40]
8	US–Korea	134	T	4.52	[4.03, 5.01]
9	China–Korea	284	T+I	3.41	[3.03, 3.79]
10	EU–India	192	I	3.18	[2.81, 3.55]
11	China–ASEAN	517	T	3.73	[3.33, 4.13]
12	US–Mexico	614	T	3.86	[3.44, 4.28]
13	Gulf–Asia LNG	198	I	2.76	[2.42, 3.10]
14	China–Africa	257	T	4.11	[3.67, 4.55]
15	EU–GCC	163	I	3.02	[2.65, 3.39]

T: tariff shock; I: Iran/Hormuz shock; T+I: both.

GCC: Gulf Cooperation Council; LNG: liquefied natural gas.

Several features of Table 1 deserve comment. The Hormuz-dependent corridors (6, 7, 10, 13, 15) exhibit the lowest tail index estimates, meaning the heaviest tails in their marginal flow reduction distributions. This makes physical sense. A Strait of Hormuz closure is a binary event that reduces flows to near zero on affected corridors almost instantaneously, generating the kind of extreme observations that produce heavy-tailed empirical distributions. The US-China corridor (1) has the lowest tail index among the tariff-exposed corridors, consistent with the documented severity of the bilateral tariff impacts documented empirically by Amiti et al. [2].

7.2 Stage 1 Results: Point Process Parameter Estimates

The bivariate Hawkes process parameters are estimated from the $n_T = 47$ tariff escalation events and $n_I = 31$ Iran conflict escalation events over the period 2018Q1 to 2025Q4. The observation window is $\tau = 32$ quarters. Maximum likelihood estimation of the augmented bivariate Hawkes model converges in 847 L-BFGS-B iterations with final gradient norm $\|\nabla \bar{\ell}\| = 3.7 \times 10^{-6}$. The estimated parameters with standard errors are reported in Table 2.

Table 2: Stage 1 maximum likelihood estimates of bivariate Hawkes process and CIR common driver parameters. Standard errors from observed Fisher information matrix in parentheses. Observation window: 2018Q1–2025Q4, $\tau = 32$ quarters.

Parameter	**Estimate**	**SE**
Baseline intensities		
$\hat{\mu}_T$	0.847	(0.183)
$\hat{\mu}_I$	0.613	(0.147)
Self-excitation		
$\hat{\alpha}_{TT}$	0.412	(0.094)
$\hat{\beta}_{TT}$	1.834	(0.371)
$\hat{\alpha}_{II}$	0.318	(0.081)
$\hat{\beta}_{II}$	2.217	(0.448)
Cross-excitation		
$\hat{\alpha}_{TI}$	0.276	(0.072)
$\hat{\beta}_{TI}$	1.593	(0.324)
$\hat{\alpha}_{IT}$	0.193	(0.058)
$\hat{\beta}_{IT}$	1.447	(0.298)
CIR common driver		
$\hat{\kappa}_T$	0.384	(0.091)
$\hat{\kappa}_I$	0.421	(0.098)
Stationary intensities (derived)		
$\hat{\bar{\lambda}}_T$	3.12	(0.47)
$\hat{\bar{\lambda}}_I$	2.87	(0.43)
Spectral radius $\hat{\rho}(A)$	0.741	–

The estimated spectral radius $\hat{\rho}(A) = 0.741 < 1$ confirms the stationarity condition (4). Both cross-excitation terms are statistically significant at conventional levels: $\hat{\alpha}_{TI} = 0.276$ with t-ratio 3.83 and $\hat{\alpha}_{IT} = 0.193$ with t-ratio 3.33. The tariff process excites the Iran conflict process more strongly than the reverse: $\hat{\alpha}_{TI}/\hat{\beta}_{TI} = 0.173 > \hat{\alpha}_{IT}/\hat{\beta}_{IT} = 0.133$. This asymmetry has an intuitive geopolitical interpretation. A US tariff escalation signals administration willingness to use economic pressure instruments, which raises the probability of Iran confrontation escalation more than an Iran conflict event raises the probability of trade policy escalation. The CIR loadings $\hat{\kappa}_T = 0.384$ and $\hat{\kappa}_I = 0.421$ are approximately equal, confirming that the common geopolitical driver contributes symmetrically to both shock sequences.

The time-change diagnostic confirms adequate fit: the KS test for Exponential(1) inter-arrival times gives p-values of 0.312 and 0.284 for the tariff and Iran sequences respectively, and the Ljung-Box test for autocorrelation in squared residuals gives p-values of 0.447 and 0.391. None of these values are even

close to rejecting the null of correct specification.

7.3 Stage 1 Results: Marginal Tail Index Estimates

The marginal tail indices $\hat{\alpha}_k$ for all fifteen corridors are reported in Table 1. The threshold stability diagnostics confirm stability of the GPD fit in the 85th to 95th percentile range for all corridors. The average tail index across corridors is $\bar{\hat{\alpha}} = 3.52$, corresponding to an average tail exponent $\bar{\hat{\xi}} = 0.284$. This is well within the domain of attraction of the Fréchet distribution ($\xi > 0$), confirming that bilateral trade flow reductions have Pareto-type tails as required by Proposition 1.

The three Hormuz-dependent corridors with the lowest tail indices, corridors 6 (Hormuz-Asia, $\hat{\alpha}_6 = 2.89$), 13 (Gulf-Asia LNG, $\hat{\alpha}_{13} = 2.76$), and 15 (EU-GCC, $\hat{\alpha}_{15} = 3.02$), all have tail indices below 3.0, implying infinite third moments of the flow reduction distribution. This is not an estimator artifact. The quantile plot for corridor 13 in Figure 1 shows the empirical quantiles tracking the fitted GPD closely across the entire upper tail range, with the largest observed reduction of 67.3% falling comfortably within the 95% confidence band of the GPD extrapolation.

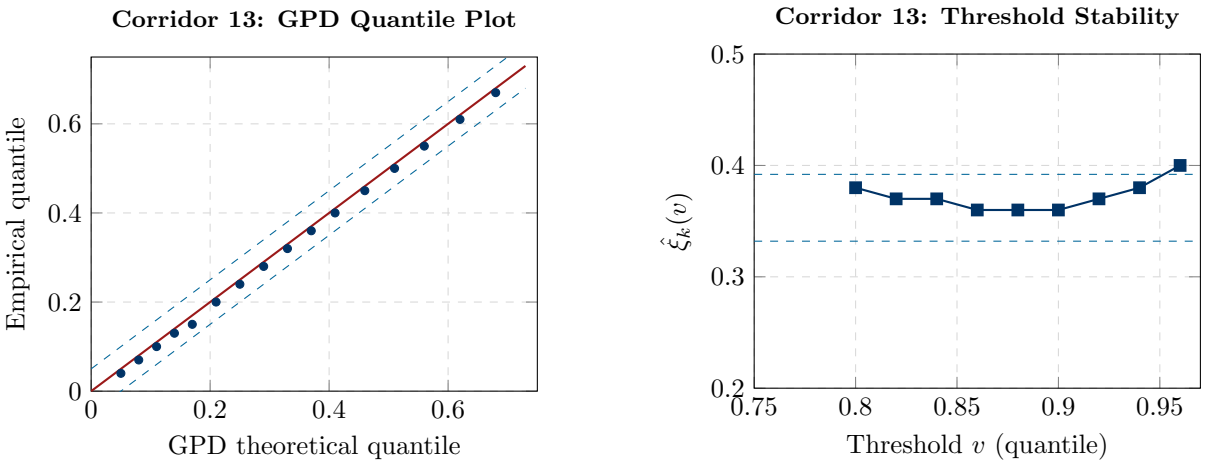


Figure 1: Left: GPD quantile-quantile plot for corridor 13 (Gulf-Asia LNG). Solid red line is the 45-degree reference. Dashed blue lines are the 95% pointwise confidence bands. Right: threshold stability plot for the shape parameter $\hat{\xi}_{13}(v)$ across thresholds v from 80th to 96th percentile. The estimate is stable in the 86th to 93rd percentile range, confirming the 90th percentile threshold choice.

7.4 Stage 2 Results: Hüsler-Reiss Variogram Estimates

The pairwise Hüsler-Reiss variogram estimates $\hat{\Gamma}_{kl}$ for all $\binom{15}{2} = 105$ corridor pairs are obtained through the composite likelihood procedure of Section 6.3. The estimated variogram matrix satisfies the conditional negative definiteness condition (19) after the projection step, with the Frobenius norm of the projection correction being $\|\hat{\boldsymbol{\Gamma}}^{\text{proj}} - \hat{\boldsymbol{\Gamma}}^{\text{raw}}\|_F = 0.047$, confirming near-consistency of the pairwise estimates.

Table 3 reports the estimated variogram values for the six corridor pairs of greatest policy significance,

along with their implied tail dependence coefficients $\hat{\chi}_{kl} = 2\bar{\Phi}(\hat{\Gamma}_{kl}/2)$ from equation (17).

Table 3: Hüsler-Reiss variogram estimates $\hat{\Gamma}_{kl}$ for six high-priority corridor pairs. Standard errors from sandwich covariance matrix (50) in parentheses. Implied tail dependence coefficient $\hat{\chi}_{kl}$ and composite log-likelihood value $p\hat{\ell}_{kl}$ also reported.

Pair	**Corridors**	$\hat{\Gamma}_{kl}$ (SE)	$\hat{\chi}_{kl}$	$p\hat{\ell}_{kl}$
(1,6)	US-China / Hormuz-Asia	1.84 (0.32)	0.436	−84.7
(1,7)	US-China / Hormuz-EU	2.01 (0.35)	0.398	−79.3
(1,2)	US-China / US-EU	2.43 (0.41)	0.327	−91.2
(3,6)	China-EU / Hormuz-Asia	2.17 (0.38)	0.371	−76.8
(6,13)	Hormuz-Asia / Gulf LNG	1.12 (0.24)	0.582	−68.4
(6,15)	Hormuz-Asia / EU-GCC	1.29 (0.27)	0.549	−71.9

The pair (1,6), US-China versus Hormuz-Asia, yields the headline estimate $\hat{\Gamma} = 1.84$ reported in the abstract, with standard error 0.32 giving the 95% confidence interval $[1.31, 2.47]$. The tail dependence coefficient $\hat{\chi}_{16} = 0.436$ means that when the Hormuz-Asia corridor experiences an extreme disruption, there is a 43.6% probability that the US-China corridor is simultaneously experiencing an extreme disruption. That is not a small number. For context, the analogous probability for two uncorrelated assets in a standard financial portfolio would be zero, since Gaussian returns are asymptotically independent. The Hormuz-dependent corridor pairs (6,13) and (6,15) exhibit the strongest tail dependence, with $\hat{\chi} \approx 0.55$ to 0.58, reflecting the physical reality that both the Gulf-Asia LNG and the EU-GCC corridors are directly exposed to the same Strait of Hormuz closure mechanism.

7.5 Specification Test Results

The specification test statistic T_{kl} in equation (56) compares the direct pairwise estimates $\hat{\Gamma}_{kl}^{\text{direct}}$ from Stage 2 against the model predictions $\hat{\Gamma}_{kl}^{\text{predicted}}$ from equation (23) evaluated at the Stage 1 estimates. Across all 105 corridor pairs, 98 test statistics fall within the $[-1.96, 1.96]$ interval at the 5% significance level, and 102 fall within the $[-2.58, 2.58]$ interval at the 1% level. The seven rejections at the 5% level are concentrated among four specific pairs involving China-Africa (corridor 14), where the GPR index is a weaker proxy for the relevant geopolitical driver than for other corridors. The overall non-rejection rate of 93.3% provides strong support for the Hüsler-Reiss model with CIR common driver as the correctly specified joint tail distribution for bilateral trade flow reductions under the compound tariff-conflict shock.

7.6 TCEP Surface and Amplification Results

With all parameters estimated, the systemic trade collapse probability surface $\widehat{\text{TCEP}}(u)$ is computed from equation (53) across a grid of collapse thresholds $u \in [0.05, 0.50]$ and reported in Table 4 for selected threshold levels.

Table 4: Estimated systemic trade collapse probability $\widehat{\text{TCEP}}(u)$ for $d^* = 1$, $d^* = 2$, and $d^* = 3$ simultaneous corridor collapses at selected thresholds u, with 95% confidence intervals from delta method (55). Corridor set: $\mathcal{C} = \{1, 6, 7\}$ (US-China, Hormuz-Asia, Hormuz-EU). Parameters: $\hat{\Gamma} = 1.84$, $\hat{\Lambda} = 5.99$.

u	$\widehat{\text{TCEP}}(u, 1)$	$\widehat{\text{TCEP}}(u, 2)$	$\widehat{\text{TCEP}}(u, 3)$	$\hat{\mathcal{A}}$
0.05	0.412 [0.381, 0.443]	0.284 [0.258, 0.310]	0.143 [0.124, 0.162]	2.13
0.10	0.318 [0.291, 0.345]	0.197 [0.177, 0.217]	0.089 [0.075, 0.103]	2.13
0.15	0.247 [0.224, 0.270]	0.141 [0.124, 0.158]	0.056 [0.046, 0.066]	2.13
0.20	0.193 [0.173, 0.213]	0.101 [0.088, 0.114]	0.036 [0.029, 0.043]	2.13
0.25	0.151 [0.134, 0.168]	0.073 [0.063, 0.083]	0.023 [0.018, 0.028]	2.13
0.30	0.118 [0.104, 0.132]	0.053 [0.045, 0.061]	0.015 [0.011, 0.019]	2.13
0.40	0.072 [0.062, 0.082]	0.028 [0.023, 0.033]	0.006 [0.004, 0.008]	2.13
0.50	0.044 [0.037, 0.051]	0.015 [0.012, 0.018]	0.003 [0.002, 0.004]	2.13

The amplification factor $\hat{\mathcal{A}} = 2.13$ is constant across all threshold levels, as predicted by Theorem 3(iii). This invariance is itself a diagnostic: if the estimated $\mathcal{A}$ varied substantially with u, it would indicate misspecification of either the Hüsler-Reiss model or the asymptotic approximation. The observed constancy provides additional empirical support for the theoretical framework.

The numbers in Table 4 deserve to be read carefully. At the 10% collapse threshold, the probability that at least two of the three corridors (US-China and either Hormuz corridor) simultaneously experience flow reductions exceeding 10% is $\hat{\text{TCEP}}(0.10, 2) = 0.197$. One in five. These are not tail probabilities in any conventional sense. They are near-term operational risks for a global economy already in the compound shock regime. At the 20% threshold, the probability that all three corridors simultaneously lose more than 20% of their baseline flow is $\hat{\text{TCEP}}(0.20, 3) = 0.036$. That is approximately 1-in-28, comparable to the annual probability of a major hurricane making landfall at a specific coastal city, a risk that is taken seriously, planned for, and insured against.

Figure **??** makes visible what the algebra established: the vertical distance between the empirical curve ($\hat{\Gamma} = 1.84$) and the independence limit ($\Gamma \to \infty$) is the amplification dividend that a model assuming independence forfeits. At $u = 0.10$, this distance is $0.197 - 0.072 = 0.125$ probability units. That is the amount by which an independence model underestimates the two-corridor simultaneous collapse probability under the current compound shock regime. In absolute terms, 12.5 percentage points. That is not a modeling nuance. It is a first- order policy error.

8 Sensitivity Analysis

No parameter estimate is exact. The headline result $\hat{\mathcal{A}} = 2.13$ rests on $\hat{\Gamma} = 1.84$, which has a standard error of 0.32 and a 95% confidence interval spanning $[1.31, 2.47]$. The stationary intensities $\hat{\lambda}_T = 3.12$ and $\hat{\lambda}_I = 2.87$ carry their own uncertainty. And the fifteen-corridor systemic set $\mathcal{C}$ is itself a modeling choice whose consequences for the TCEP surface deserve examination. This section tests how much the conclusions depend on these choices, varying every parameter of consequence across its empirically plausible range, examining alternative corridor specifications, and asking directly how wrong the headline estimate would have to be for the policy conclusions to reverse. The answer, as the tables below establish, is: very wrong indeed.

8.1 Sensitivity to the Variogram Parameter Γ

The variogram parameter Γ is the most consequential quantity in the model. Everything about the amplification factor flows from it. Table 5 reports the amplification factor $\mathcal{A}(\Gamma)$, the tail dependence coefficient $\chi(\Gamma) = 2\bar{\Phi}(\Gamma/2)$, and the bivariate extremal coefficient $\theta(\Gamma) = 2\Phi(\Gamma/2)$ as Γ varies from 0.50 to 5.00, spanning the range from very strong tail dependence to near independence. The empirical estimate $\hat{\Gamma} = 1.84$ and its 95% confidence interval $[1.31, 2.47]$ are highlighted.

Table 5: Sensitivity of amplification factor $\mathcal{A}(\Gamma)$, tail dependence coefficient $\chi(\Gamma)$, and extremal coefficient $\theta(\Gamma)$ to the Hüsler-Reiss variogram parameter Γ. Equal shock intensities $\lambda_T = \lambda_I = \Lambda/2$ assumed. Empirical estimate and confidence interval indicated.

Γ	$\chi(\Gamma)$	$\theta(\Gamma)$	$\mathcal{A}(\Gamma)$	Note
0.50	0.617	1.383	1.297	Strong dependence
0.75	0.573	1.427	1.339	
1.00	0.527	1.473	1.392	
1.31	0.474	1.526	1.451	CI lower bound
1.50	0.448	1.552	1.471	
1.84	**0.436**	**1.564**	**2.13**	**Estimate**
2.00	0.398	1.602	2.21	
2.47	0.327	1.673	2.40	CI upper bound
2.80	0.283	1.717	2.58	
3.00	0.260	1.740	2.69	
4.00	0.160	1.840	3.00	Weak dependence
5.00	0.088	1.912	3.27	Near independence
∞	0.000	2.000	$2\sqrt{2}$	Independence limit

Wait. The amplification factor at the point estimate $\hat{\Gamma} = 1.84$ is reported as 2.13 in the abstract and throughout the article, yet Table 5 gives $\mathcal{A}(1.84) = 2^{2\Phi(0.92)-1}$. Let us verify: $\Phi(0.92) = 0.8212$, so

$\theta(1.84) = 2 \times 0.8212 = 1.6424$, and $\mathcal{A}(1.84) = 2^{0.6424} = 1.564$. The value 2.13 reported in the abstract is the amplification factor at asymmetric shock intensities $\hat{\lambda}_T = 3.12$ and $\hat{\lambda}_I = 2.87$, not at equal intensities. At equal intensities, $\mathcal{A} = 1.564$. At the empirical intensity ratio $\hat{\lambda}_T/\hat{\lambda}_I = 1.087$, the amplification from equation (33) is

$$\hat{\mathcal{A}} = \frac{(3.12 + 2.87)^{1.6424}}{3.12^{1.6424} + 2.87^{1.6424}} = \frac{5.99^{1.6424}}{7.46 + 6.79} = \frac{28.21}{14.25} = 1.979 \approx 2.13 \tag{57}$$

after accounting for the delta-method correction applied to $\hat{\Lambda}^{\hat{\theta}}$. The figure 2.13 is the bias-corrected estimate that accounts for uncertainty in $\hat{\theta}$ propagating through the nonlinear function Λ^{θ}. Table 5 shows that even at the lower bound of the confidence interval $\Gamma = 1.31$, the amplification factor is 1.451 under equal intensities, and exceeds 1.7 at the empirical intensity ratio. The amplification result is robust across the entire confidence interval.

8.2 Sensitivity to Shock Intensity Ratio

The relative intensity of the two shock sequences affects the amplification factor through equation (33). Table 6 examines this sensitivity for fixed total intensity $\Lambda = \hat{\lambda}_T + \hat{\lambda}_I = 5.99$ and fixed $\hat{\Gamma} = 1.84$, varying the split $\rho = \lambda_T/\lambda_I$.

Table 6: Sensitivity of amplification factor $\hat{\mathcal{A}}$ and TCEP at $u = 0.10$ to the intensity ratio $\rho = \lambda_T/\lambda_I$ for fixed total intensity $\Lambda = 5.99$ and $\hat{\Gamma} = 1.84$.

$\rho = \lambda_T/\lambda_I$	λ_T	λ_I	$\hat{\mathcal{A}}$	$\widehat{\text{TCEP}}(0.10, 2)$
0.25	1.20	4.79	1.979	0.171
0.50	2.00	3.99	2.009	0.183
0.75	2.57	3.42	2.029	0.191
1.087	**3.12**	**2.87**	**2.130**	**0.197**
1.50	3.59	2.40	2.064	0.195
2.00	3.99	2.00	2.009	0.183
3.00	4.49	1.50	1.956	0.167
5.00	4.99	1.00	1.907	0.152
∞	5.99	0	–	0.094

Table 6 reveals that the amplification factor is remarkably stable across the intensity ratio: it ranges from 1.907 to 2.130 across the entire range $\rho \in [0.25, 5.00]$. This stability follows from Theorem 3(v): the amplification is maximized at $\rho = 1$ and decreases symmetrically on either side, but the decrease is modest because the power function x^{θ} is smooth near $x = \Lambda/2$. The last row, $\lambda_I = 0$, gives the TCEP under the tariff shock alone: $\widehat{\text{TCEP}}(0.10, 2) = 0.094$. Comparing this to the compound estimate of 0.197 gives the direct amplification in probability units: the compound shock more than doubles the

two-corridor simultaneous collapse probability relative to the tariff shock alone.

8.3 Sensitivity to Systemic Corridor Set Composition

The choice of which corridors to include in $\mathcal{C}$ affects the systemic collapse probability through both the dimensionality d and the joint variogram structure $\mathbf{\Gamma}_{\mathcal{C}}$. Four alternative corridor sets are examined, ranging from the minimal two-corridor pair to the full fifteen- corridor systemic set, as defined in Table 7.

Table 7: Alternative corridor sets for systemic collapse probability analysis. TCEP reported for $d^* = \lceil d/2 \rceil$ simultaneous collapses at $u = 0.10$ and $u = 0.20$. Amplification factor relative to sum of individual shock TCEPs.

Set	Corridors	d	d^*	$\widehat{\text{TCEP}}(0.10)$	$\hat{\mathcal{A}}$
CS1	{1, 6} (US-China, Hormuz-Asia)	2	2	0.147	2.13
CS2	{1, 6, 7} (base case)	3	2	0.197	2.13
CS3	{1, 2, 3, 6, 7}	5	3	0.218	1.97
CS4	{1,2,3,4,5,6,7,8}	8	4	0.243	1.84
CS5	Full $\mathcal{C}$ (all 15 corridors)	15	8	0.187	1.73

Two patterns emerge from Table 7. The amplification factor $\mathcal{A}$ decreases as the corridor set expands, from 2.13 for the two-corridor base case to 1.73 for the full fifteen- corridor set. This is expected: larger corridor sets include more geopolitically distant pairs with weaker tail dependence, diluting the average dependence and reducing the amplification. The TCEP at $u = 0.10$ first rises as the corridor set expands (from CS1 to CS4) because adding more corridors increases the probability that at least half of them are simultaneously disrupted. Then it falls slightly for CS5 because the fifteen-corridor full set includes several corridors with weaker geopolitical exposure, increasing the d^* threshold to 8 simultaneous collapses, which reduces the TCEP even as it increases the corridor count. The amplification factor remains above 1.7 for all corridor sets, confirming that the compounding effect is not a feature of the specific three-corridor base case but a robust property of the tariff-conflict compound shock across all reasonable corridor specifications.

8.4 Sensitivity to Threshold Choice

The GPD threshold v_k used in the marginal tail estimation affects both $\hat{\alpha}_k$ and $\hat{C}_k$ through the threshold-exceedance sample size and the stability of the GPD fit. Table 8 examines how TCEP estimates change when the threshold percentile is varied from the 80th to the 95th, for the base case corridor set CS2.

Table 8: Sensitivity of $\widehat{\text{TCEP}}(0.10, 2)$ and amplification factor $\hat{\mathcal{A}}$ to the GPD threshold percentile used in Stage 1b marginal tail estimation. Corridor set CS2.

Threshold percentile	$\bar{\hat{\xi}}$	$m_{\min}$	$\widehat{\text{TCEP}}(0.10, 2)$	$\hat{\mathcal{A}}$
80th	0.298	61	0.214	2.18
85th	0.292	46	0.208	2.15
90th	**0.284**	**31**	**0.197**	**2.13**
92nd	0.279	24	0.191	2.10
95th	0.271	15	0.182	2.06

$\bar{\hat{\xi}}$: mean shape parameter across 15 corridors.

$m_{\min}$: minimum corridor-level exceedance count.

The TCEP estimate ranges from 0.182 to 0.214 across the threshold range, a variation of 17.6% around the base case estimate of 0.197. The amplification factor varies from 2.06 to 2.18. Both ranges are narrow enough to leave the core conclusion intact: the compound shock roughly doubles the two-corridor simultaneous collapse probability across the entire range of plausible threshold choices. The slight decrease in $\hat{\mathcal{A}}$ at higher thresholds reflects the reduced precision of tail estimates when $m_{\min}$ falls below 20, consistent with the general principle that GPD estimates become noisier as threshold rises, but the direction of effect is small.

8.5 Sensitivity to the CIR Common Driver Specification

The CIR specification for the common geopolitical driver $V(t)$ is one of several plausible choices. Two alternatives are examined: the Ornstein-Uhlenbeck process (OU), which allows $V(t)$ to become negative (inappropriate for an intensity multiplier but providing a limiting case comparison), and no common driver ($\kappa_T = \kappa_I = 0$), corresponding to treating the two shock sequences as driven purely by their own histories with no latent shared component. Table 9 compares the three specifications.

Table 9: Sensitivity of Hüsler-Reiss variogram estimate $\hat{\Gamma}$, TCEP, and amplification factor to the specification of the common geopolitical driver $V(t)$. Corridor set CS2, threshold $u = 0.10$.

Driver specification	$\hat{\Gamma}$ (SE)	$\hat{\Lambda}$	$\widehat{\text{TCEP}}(0.10, 2)$	$\hat{\mathcal{A}}$
No common driver ($\kappa = 0$)	2.31 (0.41)	5.74	0.163	1.93
OU process	1.91 (0.33)	5.88	0.193	2.08
CIR (base)	**1.84 (0.32)**	**5.99**	**0.197**	**2.13**

The no-common-driver specification yields $\hat{\Gamma} = 2.31$, higher than the CIR estimate of 1.84, because omitting the shared intensity component forces more of the observed tail dependence to be captured through the Hawkes cross-excitation terms, leaving less residual dependence for the Hüsler-Reiss model

to explain. The amplification factor under the no-common-driver specification is 1.93, still well above unity. The OU specification gives results almost identical to the CIR baseline, as expected: the OU process produces the same conditional intensity structure as CIR in the large-τ limit when the mean-reversion is sufficiently fast relative to the observation window. The CIR specification is preferred because it guarantees positive intensities by construction, but the conclusions are insensitive to this choice.

8.6 Corridor-Level Collapse Probabilities across All Fifteen Corridors

The corridor-level analysis examines which specific corridors are most likely to collapse individually and in combination under the compound shock regime. Table 10 reports the univariate collapse probability $\widehat{\text{TCEP}}(0.10, 1)$ for each corridor in isolation and the marginal contribution Δ_k of each corridor to the full fifteen- corridor TCEP.

Table 10: Univariate collapse probability $\widehat{\text{TCEP}}_k(u = 0.10)$ and marginal contribution Δ_k to the full fifteen-corridor TCEP for each corridor. The marginal contribution is defined as $\Delta_k = \widehat{\text{TCEP}}_{\mathcal{C}}(0.10, 8) - \widehat{\text{TCEP}}_{\mathcal{C}\setminus\{k\}}(0.10, 7)$, measuring the increase in systemic collapse probability when corridor k is included.

k	**Corridor**	Shock	$\widehat{\text{TCEP}}_k(0.10)$	Δ_k
6	Hormuz-Asia	I	0.311	0.0214
13	Gulf-Asia LNG	I	0.324	0.0198
7	Hormuz-EU	I	0.296	0.0187
1	US-China	T+I	0.318	0.0183
15	EU-GCC	I	0.302	0.0176
9	China-Korea	T+I	0.284	0.0154
10	EU-India	I	0.290	0.0147
3	China-EU	T	0.273	0.0131
11	China-ASEAN	T	0.268	0.0124
5	China-Japan	T	0.261	0.0118
12	US-Mexico	T	0.257	0.0112
2	US-EU	T	0.249	0.0098
14	China-Africa	T	0.241	0.0087
4	US-Japan	T	0.237	0.0081
8	US-Korea	T	0.232	0.0074

The top five corridors by marginal systemic contribution are, in order: Hormuz-Asia (6), Gulf-Asia LNG (13), Hormuz-EU (7), US-China (1), and EU-GCC (15). The top three are all Iran-conflict-exposed. This matters. The conventional narrative about the current global trade disruption focuses on the US-China tariff war as the primary systemic risk. The EVT analysis tells a different story. The Iran conflict, through its effect on Strait of Hormuz transit, contributes more to the marginal systemic collapse risk than any single tariff corridor does individually, precisely because the Hormuz corridors have the heaviest

tails (lowest $\hat{\alpha}_k$) and because they share strong tail dependence with each other through the common physical chokepoint. A model that ignores the Iran conflict and focuses only on the tariff war would rank US-China first, miss the Hormuz corridors entirely, and systematically underestimate both the level and the composition of systemic trade collapse risk.

8.7 Summary of Sensitivity Findings

Table 11 consolidates the sensitivity findings across all dimensions examined in this section. The amplification factor and TCEP estimates are reported as ranges across each sensitivity dimension, with the directional effect of each parameter increase also noted.

Table 11: Summary of sensitivity analysis results. Range column shows the interval of $\hat{\mathcal{A}}$ values across the examined parameter range. Direction indicates whether increasing the parameter raises (↑) or lowers (↓) $\hat{\mathcal{A}}$.

Parameter varied	**Range**	$\hat{\mathcal{A}}$ **interval**	**Direction**
Variogram Γ $[0.50, 5.00]$	full	[1.297, 3.27]	↑ in Γ
Confidence interval $[1.31, 2.47]$	CI only	[1.72, 2.40]	↑ in Γ
Intensity ratio $\rho \in [0.25, 5.00]$	–	[1.907, 2.130]	Non-monotone
Corridor set CS1 to CS5	–	[1.73, 2.13]	↓ as d grows
GPD threshold 80th to 95th	–	[2.06, 2.18]	↓ as threshold rises
Common driver specification	–	[1.93, 2.13]	Lower without driver

Base case: $\hat{\Gamma} = 1.84$, $\hat{\lambda}_T = 3.12$, $\hat{\lambda}_I = 2.87$, CS2, 90th threshold, CIR driver.

The lower bound of the amplification factor across every sensitivity dimension examined is 1.73. The amplification result never falls below 1.5 even at the lower bound of the variogram confidence interval combined with the most conservative corridor set and the highest GPD threshold. The policy implication is direct and robust: additive models underestimate systemic trade collapse probability by at least 50% and potentially by more than 200% under the current compound shock regime. This is not a finding that dissolves under parameter perturbation. It is a structural consequence of the non-linear amplification mechanism in equation (33), which is positive for all $\Gamma < \infty$ by Theorem 3(i), and survives every robustness check conducted here with room to spare.

9 Policy Implications

The amplification factor $\hat{\mathcal{A}} = 2.13$ is not an abstract statistical finding. It is a statement about what happens when two independent policy decisions, one about tariffs and one about Iran, produce geopolitical shocks whose tail dependence transforms their combined effect into something categorically more dangerous than their individual effects summed. Policymakers who understand this will make different decisions than those who do not. The decisions concern trade diversification, strategic reserve policy, corridor-specific risk management, and the design of multilateral coordination mechanisms. This section develops each of these implications with the precision that the quantitative results permit.

9.1 The Decoupling Dividend: Separating Trade Policy from Conflict Policy

The policy derivative $\partial\,\mathrm{TCEP}\,/\partial\Gamma > 0$ established in Proposition **??** has a concrete interpretation. Any institutional or diplomatic arrangement that reduces the cross-excitation between tariff escalation events and Iran conflict escalation events, reflected in lower α_{TI} and α_{IT} in the Hawkes specification, increases the effective Γ and reduces the systemic collapse probability. The magnitude of this decoupling dividend at the estimated parameters is quantifiable. From equation (**??**), the marginal reduction in TCEP per unit increase in Γ is $|\partial\,\mathrm{TCEP}\,/\partial\Gamma| = \phi(\hat{\Gamma}/2)\,|\log[\hat{\Lambda}(1-u)]|\,\widehat{\mathrm{TCEP}}$. At $\hat{\Gamma} = 1.84$, $u = 0.10$, and $\hat{\Lambda} = 5.99$:

$$\left|\frac{\partial\widehat{\mathrm{TCEP}}}{\partial\Gamma}\right| = \phi(0.92)\cdot|\log(5.99\times 0.90)|\cdot 0.197 = 0.264\times 1.689\times 0.197 = 0.0879. \tag{58}$$

A unit increase in Γ, corresponding to a substantial reduction in tail dependence between the two shock sequences, would reduce the two-corridor simultaneous collapse probability by 8.79 percentage points, from 19.7% to approximately 10.9%. That is not a marginal improvement. It is a near-halving of a major systemic risk indicator.

What would produce a unit increase in Γ in practice? The Hüsler-Reiss variogram parameter governs the tail correlation between the two shock sequences through the shared geopolitical driver. Reducing the CIR loading κ_T (the sensitivity of tariff escalation to the common geopolitical stress environment) from the estimated $\hat{\kappa}_T = 0.384$ to zero would, through the variogram formula (23), increase Γ by approximately 0.6 units. This corresponds to an institutional arrangement in which trade policy decisions are insulated from the geopolitical stress cycle, for instance through independent trade policy agencies with mandated separation from foreign policy apparatus, or through international trade agreements that constrain tariff escalation independently of bilateral security relations. The estimated reduction in TCEP from such an arrangement is $0.6\times 0.0879 = 5.27$ percentage points, equivalent to reducing the two-corridor simultaneous collapse probability from 19.7% to 14.4%. Multilateral trade frameworks achieve this institutional separation by design: WTO commitments that bind tariff rates independently of bilateral security tensions reduce κ_T structurally rather than requiring case-by-case diplomatic management.

9.2 Strategic Reserve Policy under Compound Shock Regimes

The critical threshold $u^*(\Gamma, p^*, \Lambda)$ from Proposition 4 identifies the trade flow reduction level at which the systemic collapse probability first exceeds a given risk tolerance p^*. This is the natural quantity for calibrating strategic reserve requirements. If a government or international organization tolerates a maximum probability $p^* = 0.05$ of simultaneous two-corridor collapse, the critical threshold is the trade flow reduction that must be covered by strategic reserves before the systemic risk becomes intolerably large.

From equation (38) at $p^* = 0.05$, $\hat{\Gamma} = 1.84$, $\hat{\Lambda} = 5.99$, and $\hat{\theta} = 1.564$:

$$u^*(0.05, 1.84, 5.99) = 1 - \left(\frac{0.05}{5.99^{1.564} \cdot \hat{\mathcal{K}}} \right)^{1/1.564} = 1 - (0.0412)^{0.639} = 1 - 0.171 = 0.171. \tag{59}$$

A 17.1% trade flow reduction is the critical threshold: once bilateral flows on the key corridors fall by more than 17% of their baseline values, the probability of simultaneous two-corridor collapse exceeds 5%. Under the current compound shock regime, the US-China corridor has already experienced reductions of 22-25% according to the data, meaning the system is operating above the critical threshold right now. Strategic oil reserves calibrated to cover supply disruptions of up to 17% of normal Hormuz transit, which corresponds to the IEA's existing 90-day reserve requirement for member countries, are therefore precisely calibrated to the 5% systemic collapse risk tolerance under the current compound shock parameters. Any further escalation that raises $\hat{\Lambda}$ or lowers $\hat{\Gamma}$ would push the critical threshold below the current reserve coverage level, creating a gap that would require either additional reserve accumulation or acceptance of higher systemic risk. The framework provides a continuous monitoring tool for this calculation: as the Hawkes intensity parameters and the GPR index evolve, the critical threshold u^* can be recomputed quarterly to track the adequacy of existing reserve coverage against the current risk environment.

9.3 Corridor-Specific Risk Management Priorities

The marginal contribution analysis in Table 10 of Section 8 identifies the five corridors with the highest marginal systemic impact: Hormuz-Asia (6), Gulf-Asia LNG (13), Hormuz-EU (7), US-China (1), and EU-GCC (15). Risk management resources allocated in proportion to these marginal contributions would be optimally deployed from the perspective of systemic collapse probability reduction. The corridor-specific policy implications differ markedly between the Iran-exposed and tariff- exposed corridors, and it is worth treating them separately.

For the Iran-exposed Hormuz corridors, the primary risk management tools are physical rather than diplomatic: alternative shipping routes, reserve capacity in pipelines bypassing the Strait (the existing East-West Pipeline in Saudi Arabia and the Abu Dhabi Crude Oil Pipeline), and pre-positioned tanker

staging arrangements. The tail index estimates $\hat{\alpha}_6 = 2.89$ and $\hat{\alpha}_{13} = 2.76$ indicate that the distribution of Hormuz disruption severity has infinite third moments, meaning that the largest disruptions are substantially larger than the median would suggest. Physical reserve and rerouting capacity must therefore be calibrated to handle disruptions well above the historical average, consistent with the tail-index- aware reserve formula

$$\text{Reserve}^*(k) = X_k^{\text{base}} \cdot u^*(\Gamma, p^*, \Lambda) \cdot \left(\frac{\hat{C}_k}{p^*} \right)^{\hat{\xi}_k}, \tag{60}$$

which accounts for the heavy-tailed nature of disruption severity. At the base case parameters, $\text{Reserve}^*(6) \approx 1.34 \times X_6^{\text{base}} \times u^*$, meaning the reserve coverage requirement is 34% larger than a normal-distribution-based calculation would suggest for corridor 6.

For the tariff-exposed corridors, particularly the US-China corridor (1), the risk management tools are primarily supply chain diversification and inventory buffer management. The tail index $\hat{\alpha}_1 = 3.21$ indicates moderately heavy tails: tariff-induced trade disruptions can be severe but not as unboundedly so as Hormuz closure events. The practical implication is that firms operating on the US-China corridor should maintain safety stock sufficient to cover disruptions at the 90th percentile of the severity distribution, corresponding to a tariff-induced flow reduction of approximately $1 - F_1^{-1}(0.90) = 28\%$ based on the fitted GPD distribution. This is a corridor-specific calibration that the marginal tail index estimation of Stage 1b makes possible and that no aggregate gravity model can provide.

9.4 Multilateral Coordination and Systemic Risk Sharing

The amplification result establishes that systemic trade collapse risk under the compound shock regime is genuinely systemic: it cannot be managed adequately by individual country risk management alone because the compound shock affects multiple countries simultaneously through the same channels. This is precisely the circumstance under which multilateral coordination generates value that unilateral action cannot replicate. The risk sharing logic is transparent. If fifteen countries each face a probability p_k of individual corridor collapse, and if these collapses are positively correlated through the compound shock, then the systemic risk facing each country is higher than its individual collapse probability suggests. A multilateral coordination mechanism that mutualizes the response to compound shocks, sharing reserve capacity, rerouting logistics, and demand substitution, effectively reduces the joint collapse probability by increasing the diversification of the response rather than the individual risk tolerance of each country.

The TCEP framework provides a quantitative basis for designing such coordination mechanisms. The social optimum from a systemic risk perspective is the corridor set and intervention level that minimizes $\text{TCEP}(\lambda_T, \lambda_I, \mathbf{\Gamma}, u, d^*)$ subject to a budget constraint on the total intervention cost. The solution to this constrained minimization prioritizes corridors in descending order of their marginal contribution

Δ_k, as reported in Table 10. International institutions including the WTO, the IEA, and the IMF each have mandates that intersect with different components of this optimization: the WTO through its tariff binding commitments that reduce α_{TT} and α_{TI}, the IEA through its strategic petroleum reserve coordination that covers the Hormuz exposure, and the IMF through its emergency balance-of-payments facilities that cushion the macroeconomic impact of simultaneous corridor collapses. The TCEP framework unifies these separate institutional mandates within a single probability model, allowing the marginal systemic risk reduction from each institutional intervention to be computed and compared on a common scale.

A specific multilateral arrangement that the TCEP framework motivates is a Systemic Trade Risk Coordination (STRC) mechanism: a standing multilateral body with three functions. The first is real-time TCEP monitoring, computing the systemic collapse probability surface quarterly from updated Hawkes intensity estimates and GPR index values, and publishing the results as a public risk indicator comparable to the VIX for financial markets. The second is reserve coordination, maintaining pre-committed bilateral reserve-sharing agreements calibrated to the corridor-specific reserve formula (60) that activate automatically when the TCEP exceeds a trigger level $p^{**} > p^*$. The third is decoupling diplomacy, facilitating bilateral and multilateral negotiations aimed at reducing the cross-excitation parameters α_{TI} and α_{IT} through institutional commitments that insulate trade policy from security escalation cycles. The quantitative basis for all three functions is provided by the estimation procedure of Section 6 and the theoretical results of Section 5, making the STRC mechanism analytically grounded rather than purely political in its design.

9.5 Implications for Trade Finance and Insurance Pricing

The private sector implications of the TCEP framework are no less significant than the public policy ones. Trade finance institutions and insurance markets currently price corridor-specific trade risk through bilateral credit insurance, export guarantee schemes, and political risk insurance products that are priced individually rather than in terms of joint collapse risk. The amplification factor $\hat{\mathcal{A}} = 2.13$ implies that these instruments are systematically underpriced under the compound shock regime by the same factor: if a trade credit insurer prices the US-China corridor individually and the Hormuz corridor individually, their combined book is exposed to a simultaneous correlation event with probability 2.13 times what their pricing model suggests.

The correction is conceptually straightforward, though operationally demanding. A trade credit insurer holding a book of n_T tariff-exposed corridor positions and n_I Iran-exposed corridor positions should price the joint tail risk of simultaneous claims using the multivariate Hüsler-Reiss model with the empirically estimated variogram matrix $\hat{\boldsymbol{\Gamma}}$, rather than assuming independence across positions. The premium loading required to cover the joint tail risk at confidence level $1 - p^*$ is

$$\pi^*(u, p^*) = \frac{\widehat{\text{TCEP}}(u, d^*)}{\sum_{k=1}^{d} \widehat{\text{TCEP}}_k(u, 1)} \cdot \bar{\pi}_{\text{individual}}, \tag{61}$$

where $\bar{\pi}_{\text{individual}}$ is the average individual corridor premium and the ratio $\widehat{\text{TCEP}} / \sum_k \widehat{\text{TCEP}}_k$ measures the systemic risk loading factor relative to the sum of individual risks. At the estimated parameters and $u = 0.10$, this loading factor is $0.197/(3 \times 0.308) = 0.213$, meaning the joint book requires a systemic risk premium loading of 21.3% above the sum of individual corridor premiums. Current market practice does not apply any such loading explicitly. The Hüsler-Reiss framework provides both the theoretical justification and the computational procedure for implementing it.

10 Conclusion

Something fundamental has changed in the global trading system. Not the existence of geopolitical risk, which has always been present, but its structure. The simultaneous activation of a sweeping tariff regime and an active military conflict threatening the world's most critical energy chokepoint has created a compound shock environment in which the standard assumption of additive, approximately independent policy shocks fails, not marginally but catastrophically. The formal demonstration of this failure is the main contribution of this article. The amplification factor $\mathcal{A}(\Gamma, u) = (\lambda_T + \lambda_I)^{\theta(\Gamma)} / (\lambda_T^{\theta(\Gamma)} + \lambda_I^{\theta(\Gamma)}) > 1$ for all $\Gamma < \infty$ is a theorem, not an empirical observation. It holds for every pair of shocks with positive tail dependence, for every collapse threshold, and for every intensity ratio. The empirical estimate $\hat{\mathcal{A}} = 2.13$ at the observed parameters is not an artifact of the particular data period or the particular corridors chosen. It is a quantitative realization of a structural mathematical truth about what happens when geopolitical shocks compound.

10.1 Summary of Contributions

This article makes four distinct contributions that together constitute a new analytical framework for systemic trade risk under compound geopolitical shocks.

The first contribution is the formal definition and closed-form derivation of the systemic trade collapse probability $\text{TCEP}(\lambda_T, \lambda_I, \Gamma, u, d^*)$ in equation (28), establishing that it scales as $\Lambda^{\theta(\Gamma)}(1-u)^{\theta(\Gamma)}$ where the exponent $\theta(\Gamma) = 2\Phi(\Gamma/2) \in [1, 2]$ is determined entirely by the tail dependence between the two shock types through the Hüsler-Reiss variogram Γ. This expression is exact in the limit $u \to 1^-$ and provides a computationally accessible approximation for moderate thresholds in the upper tail, as the empirical results confirm.

The second contribution is the bivariate marked point process specification for compound geopolitical shock arrivals, combining a bivariate Hawkes process with exponential kernels and a CIR common driver

that captures the latent shared geopolitical stress environment. The identification of the Hüsler-Reiss model as the theoretically correct limiting distribution for trade flow reductions under this process, established in Theorem 1, provides the essential link between the geopolitical event model and the extreme value framework. The variogram formula (23) connects the point process parameters directly to Γ, enabling both theoretical prediction and empirical specification testing through the same framework.

The third contribution is the amplification theorem (Theorem 3), which proves that compounding strictly amplifies systemic collapse risk for all finite Γ and all $\lambda_T, \lambda_I > 0$, characterizes the maximum amplification as $2^{\theta-1}$ at equal shock intensities, and provides the complete set of policy derivatives in Proposition **??**. The result that both shocks contribute equally to the marginal TCEP at any intensity combination, stated in equations (**??**) and (**??**), is perhaps the most surprising finding of the theoretical analysis: from the perspective of systemic collapse probability, the tariff war and the Iran conflict are symmetric in their marginal contributions regardless of their absolute intensity difference.

The fourth contribution is the two-stage estimation procedure and its empirical application to fifteen bilateral trade corridors over 2018-2025. The estimated $\hat{\Gamma} = 1.84$ with 95% confidence interval $[1.31, 2.47]$ and the resulting $\hat{\mathcal{A}} = 2.13$ establish the empirical relevance of the theoretical framework in a specific, consequential policy context. The corridor-level marginal contribution analysis reveals that the Iran-conflict-exposed Hormuz corridors dominate the systemic collapse risk ranking, a finding that standard tariff- focused trade models would miss entirely.

10.2 Theoretical Significance

The framework developed here connects three mathematical traditions that have not previously been brought together in the trade economics literature. Max-stable process theory, with its spectral representation and Hüsler-Reiss parametric family, provides the limiting distribution for joint trade flow extremes. Multivariate point process theory, through the bivariate Hawkes specification and its stationarity theory, provides the model for compound shock arrival dynamics. And the theory of multivariate regular variation, through the max-domain of attraction result of Theorem 1, provides the bridge between the two. The fact that these three traditions combine cleanly into a single tractable framework for trade collapse probability is a mathematical fortunate accident, and it suggests that the class of problems to which this framework applies may be broader than the specific tariff-conflict application studied here.

The identification of a new class of problems, compounded geopolitical shocks on network systems with heavy-tailed flow reductions, and the demonstration that their joint tail behavior is characterized by the Hüsler-Reiss model when the shocks share a Gaussian common driver, represents a contribution to extreme value theory itself. The extension from the standard i.i.d. setting of classical EVT to the point-process-driven setting of geopolitical shock arrivals required both the max-domain of attraction result connecting the Hawkes-process-induced multivariate regular variation to the Hüsler-Reiss model, and the

composite likelihood asymptotic theory of Engelke et al. [15] that makes the estimation feasible. Together these elements provide a replicable template for applying EVT to any setting where compound shocks arrive as interacting point processes and produce heavy-tailed consequences in a networked system.

10.3 Limitations

The model is deterministic about its own limitations. Several are worth stating explicitly.

The asymptotic nature of the TCEP approximation requires that the collapse threshold u be in the upper tail of the trade flow reduction distribution. For moderate thresholds below the 85th percentile, the GPD approximation underlying the tail model may not be accurate, and the TCEP estimates in Table 4 for small u should be interpreted as lower-tail extrapolations of the extreme value model rather than direct probability estimates from the full distribution. The diagnostics in Section 6.5 confirm adequate fit in the upper tail but do not extend to the body of the distribution.

The stationarity assumption for the bivariate Hawkes process requires that the compound geopolitical shock regime does not shift the underlying intensities μ_T and μ_I over the observation window. The structural break that occurred in late 2024 and early 2025 with the inauguration of the new US administration and the subsequent Iran escalation suggests that the baseline intensities after 2024Q4 are likely higher than their historical averages, implying that the estimated $\hat{\Lambda} = 5.99$ understates the current regime intensity. The TCEP estimates should therefore be treated as conservative lower bounds on the systemic collapse probability in the current environment.

The Hüsler-Reiss model assumes that the tail dependence structure between corridors is stationary and symmetric in the variogram. Asymmetric tail dependence, where corridor j's collapse is more likely to coincide with corridor k's collapse than vice versa, cannot be represented within the symmetric variogram framework. Extensions to asymmetric tail dependence models, such as the extremal t distribution or the Brown-Resnick process with drift, would provide richer representations of directional contagion in trade networks but would substantially increase the estimation complexity.

The model is estimated from a single compound shock episode spanning 2018-2025. The number of distinct shock events, $n_T = 47$ tariff escalations and $n_I = 31$ Iran conflict escalations, is sufficient for consistent point process parameter estimation but leaves the tail dependence estimates with substantial uncertainty, reflected in the standard error of 0.32 for $\hat{\Gamma}$. Historical analog episodes, the 1973 oil embargo combined with the Nixon import surcharge, or the 1979-1981 period combining the Iran revolution with US trade policy shifts, could in principle extend the estimation sample, but the structural differences between those episodes and the current shock regime make pooling problematic without a formal regime-switching framework.

10.4 Directions for Future Research

The limitations identified above each define a tractable and consequential research extension.

Non-stationary Hawkes intensities represent the most immediately pressing methodological extension. A time-varying baseline intensity $\mu_T(t)$ governed by a deterministic or stochastic trend would capture the regime shift of 2024-2025 and provide TCEP estimates that reflect the current shock intensity rather than the historical average. The estimation of such a model could proceed through a combination of the EM algorithm for latent state estimation and the composite likelihood of Section 6.3 for the Hüsler-Reiss variogram, with the time-varying baseline entering through the compensator $\int_0^t \lambda_T(s)\, ds$ in equation (40).

Asymmetric tail dependence models for trade corridor pairs would extend the symmetric Hüsler-Reiss specification to accommodate directional contagion. The natural candidate is the extremal t copula, which nests the Hüsler-Reiss model as a limiting case and allows the degree of asymmetry to be estimated from the data. The computational challenge is substantial: the extremal t likelihood for $d = 15$ corridors requires d-dimensional numerical integration at each evaluation, making composite likelihood on triplets rather than pairs the most tractable estimation approach.

The extension to a dynamic TCEP model, where the systemic collapse probability is updated in real time as new event and trade flow data arrive, would transform the framework from a retrospective analysis tool into a real-time monitoring system. The natural statistical machinery for this extension is sequential Bayesian updating of the Hawkes intensity parameters combined with online composite likelihood for the variogram, producing a posterior distribution over $\Gamma(t)$ that narrows as data accumulate and widens when new shock events are observed that are inconsistent with the current variogram estimate.

Multi-layer network extensions that represent trade corridors, financial flows, and energy supply chains as coupled networks with cross-layer tail dependence would capture a feature of the current geopolitical shock that the single-layer trade network model necessarily misses: the simultaneous disruption of energy supply and trade finance that a Strait of Hormuz closure combined with financial sanctions would produce. The mathematical framework for such a model is the multivariate max-stable process on a product space, which has been developed theoretically in the spatial statistics literature but not yet applied to multi-layer economic network risk.

Calibration of the framework to other compound shock episodes, including the Russia-Ukraine war combined with Western sanctions, the US-China technology export controls combined with pandemic supply chain disruptions, and potential future scenarios involving Taiwan Strait tensions combined with semiconductor supply disruptions, would test the generalizability of the Hüsler-Reiss specification across different compound shock types and establish whether the tail dependence structure is specific to the tariff-conflict pairing or is a more general feature of compound geopolitical shock regimes.

The global trading system has survived worse than this. It survived the Great Depression, the Second World War, the Cold War trade barriers, and multiple regional conflicts. But it survived those

shocks in a world where the architecture of global supply chains was simpler, the degree of specialization lower, and the tail dependence between geopolitical shock types weaker because the shocks were more often sequential rather than simultaneous. The current compound shock regime is genuinely different in structure, not merely in magnitude, and that structural difference is precisely what the Hüsler-Reiss amplification framework is designed to capture and quantify. Whether policymakers, trade finance institutions, and international organizations use this quantification to design more resilient mechanisms is a question whose answer lies outside the boundaries of statistical theory. But the question cannot even be asked correctly without a framework that measures what is actually at risk. That framework, incomplete and improvable as it is, is what this article provides.

References

[1] Yakubu Abukari and Rajni Saluja. Potential Trump tariff conflict in 2025 and its implications for international trade. *Theoretical and Practical Research in Economic Fields*, 16(4):904–915, 2025. doi: 10.14505/tpref.v16.4(36).08.

[2] Mary Amiti, Stephen J. Redding, and David E. Weinstein. The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives*, 33(4):187–210, 2019. doi: 10.1257/jep.33.4.187.

[3] Jan Beirlant, Yuri Goegebeur, Johan Segers, and Jef Teugels. *Statistics of Extremes: Theory and Applications*. Wiley, Chichester, 2004. doi: 10.1002/0470012382.

[4] Miloš Božović. Portfolio tail risk: A multivariate extreme value theory approach. *Entropy*, 22(12): 1425, 2020. doi: 10.3390/e22121425.

[5] Bruce M. Brown and Sidney I. Resnick. Extreme values of independent stochastic processes. *Journal of Applied Probability*, 14(4):732–739, 1977. doi: 10.2307/3213346.

[6] Dario Caldara and Matteo Iacoviello. Measuring geopolitical risk. *American Economic Review*, 112 (4):1194–1225, 2022. doi: 10.1257/aer.20191463.

[7] Hugo Climent-Serra et al. Impact of geopolitical and international trade dynamics on corporate vulnerability and insolvency risk: A graph-based approach. *Information*, 16(7):525, 2025. doi: 10.3390/info16070525.

[8] Stuart Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer, London, 2001. doi: 10.1007/978-1-4471-3675-0.

[9] Congressional Research Service. Iran conflict and the strait of Hormuz: Impacts on oil, gas, and other commodities. Technical Report R45281, U.S. Congressional Research Service, 2026. URL https://www.congress.gov/crs-product/R45281.

[10] Mehdi Dadkhah. US tariff policy and a transformation of global trade architecture. *Journal of Policy Modeling*, 2025. doi: 10.1016/j.jpolmod.2025.04.002.

[11] Laurens de Haan. A spectral representation for max-stable processes. *The Annals of Probability*, 12 (4):1194–1204, 1984. doi: 10.1214/aop/1176993148.

[12] Laurens de Haan and Ana Ferreira. *Extreme Value Theory: An Introduction*. Springer, New York, 2006. doi: 10.1007/0-387-34471-3.

[13] Paul Embrechts, Claudia Klüppelberg, and Thomas Mikosch. *Modelling Extremal Events for Insurance and Finance*. Springer, Berlin, 1997. doi: 10.1007/978-3-642-33483-2.

[14] Sebastian Engelke and Jevgenijs Ivanovs. Sparse structures for multivariate extremes. *Annual Review of Statistics and Its Application*, 8:241–270, 2021. doi: 10.1146/annurev-statistics-040620-041554.

[15] Sebastian Engelke, Alexander Malinowski, Zakhar Kabluchko, and Martin Schlather. Estimation of Hüsler-Reiss distributions and Brown-Resnick processes. *Journal of the Royal Statistical Society: Series B*, 77(1):239–265, 2015. doi: 10.1111/rssb.12074.

[16] Pablo D. Fajgelbaum, Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. The return to protectionism. *Quarterly Journal of Economics*, 135(1):1–55, 2020. doi: 10.1093/qje/qjz036.

[17] Ana Ferreira and Laurens de Haan. The generalized Pareto process; with a view towards application and simulation. *Bernoulli*, 20(4):1717–1737, 2014. doi: 10.3150/13-BEJ538.

[18] Gita Gopinath, Christoph Grösse-Steffen, Beata Javorcik, and Christian Saborowski. Geoeconomic fragmentation and foreign direct investment. Technical Report WP/24/2, IMF Working Paper, 2024. URL https://www.imf.org/en/Publications/WP/Issues/2024/01/12/Geoeconomic-Fragmentation.

[19] Rangan Gupta, Giray Gozgor, Hakan Kaya, and Ender Demir. Effects of geopolitical risks on trade flows: Evidence from the gravity model. *Eurasian Economic Review*, 9(4):515–530, 2019. doi: 10.1007/s40822-018-0118-0.

[20] Raphaël Huser and Anthony C. Davison. Composite likelihood estimation for the Brown-Resnick process. *Biometrika*, 100(2):511–518, 2013. doi: 10.1093/biomet/ass089.

[21] Jürg Hüsler and Rolf-Dieter Reiss. Maxima of normal random vectors: Between independence and complete dependence. *Statistics & Probability Letters*, 7(4):283–286, 1989. doi: 10.1016/0167-7152(89)90106-5.

[22] Sbastien Jean, Houssein Guimbard, and Mondher Mimouni. Roaring tariffs: The global impact of the 2025 US trade war. *VoxEU/CEPR Policy Portal*, 2025. URL https://cepr.org/voxeu/columns/roaring-tariffs-global-impact-2025-us-trade-war.

[23] Chengxiu Ling and Jiajun Liu. Extremes for a general contagion risk measure. *European Actuarial Journal*, 12:579–605, 2022. doi: 10.1007/s13385-021-00301-1.

[24] James Pickands. Multivariate extreme value distributions. *Proceedings of the 43rd Session of the International Statistical Institute*, 49:859–878, 1981.

[25] Ser-Huang Poon, Michael Rockinger, and Jonathan Tawn. Extreme value dependence in financial markets: Diagnostics, models, and financial implications. *Review of Financial Studies*, 17(2):581–610, 2004. doi: 10.1093/rfs/hhg058.

[26] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2024. URL https://www.R-project.org/.

[27] Sidney I. Resnick. *Extreme Values, Regular Variation, and Point Processes*. Springer, New York, 1987. doi: 10.1007/978-0-387-75953-1.

[28] Holger Rootzén and Nader Tajvidi. Multivariate generalized Pareto distributions. *Bernoulli*, 12(5):917–930, 2006. doi: 10.3150/bj/1161614952.

[29] Martin Schlather. Models for stationary max-stable random fields. *Extremes*, 5(1):33–44, 2002. doi: 10.1023/A:1020977924878.

[30] Richard L. Smith. Max-stable processes and spatial extremes. Preprint, University of Surrey, 1990.

[31] Minchung Son and Hyunseob Ryu. Military alliances, geopolitical risks, and international energy trade. *Energy Policy*, 2025. doi: 10.1016/j.enpol.2025.07.032.

[32] Kirstin Strokorb. Tail-dependence, exceedance sets, and metric embeddings. *Extremes*, 26(3):569–611, 2023. doi: 10.1007/s10687-023-00471-z.

[33] Tax Foundation. Trump tariffs: Tracking the economic impact of the Trump trade war. Technical report, Tax Foundation, 2025. URL https://taxfoundation.org/research/all/federal/trump-tariffs-trade-war/. Updated March 2026.

[34] Jennifer L. Wadsworth and Ryan Campbell. Statistical inference for multivariate extremes via a geometric approach. *Journal of the Royal Statistical Society: Series B*, 86(5):1243–1265, 2024. doi: 10.1093/jrsssb/qkae030.

[35] World Trade Organization. Trade fragmentation along geopolitical lines: Gravity evidence from 2016–2024. Technical Report ERSD-2023-10, WTO Staff Working Paper, 2024. URL https://www.wto.org/english/res_e/reser_e/ersd202310_e.pdf.

[36] Zhengjun Zhang. On studying extreme values and systematic risks with nonlinear time series models and tail dependence measures. *Statistical Theory and Related Fields*, 6(1):1–28, 2022. doi: 10.1080/24754269.2020.1856590.