Master Degree
in Economics, Finance and Sustainability

Final Thesis

Volatility Connectedness and Network Topology in the Reinsurance Sector

Supervisor
Prof. Giulia Iori

Graduand
Andrea De Bortoli
Matriculation Number 886277

Academic Year
2025/2026

Table of contents

Abstract

This thesis analyzes systemic risk and volatility spillovers in the global reinsurance market by utilizing Bayesian Vector Autoregressive (BVAR) modeling with network analysis. The financial impact of the reinsurance sector in the overall stability remains relatively unexplored in systemic risk literature, with uncertain and discordant findings. We addressed this gap by studying 27 publicly listed global reinsurance companies over the period October 2015 to May 2025, using daily stock prices to construct Yang-Zhang volatility series. We then employed a BVAR framework with hierarchical Minnesota, SUR and SOC priors, estimated through Metropolis-Hastings MCMC sampling. The adjacency matrix is derived using Generalized Forecast Error Variance Decomposition (GFEVD) and is used as the basis for graph theoretic analysis and the pairwise volatility association of the companies. The static analysis identified a highly integrated reinsurance network, with a Total Connectedness of 74.28, low modularity, and disassortative structure, with European companies identified as the main source of variance contributors, and Asian firms as principal recipients. The dynamic analysis was based on a 250-day rolling window BVAR estimation across 349 windows, uncovering a bimodal regime structure of standard equilibrium and high connectivity. We find that asset-side shocks of financial and geopolitical origin tend to amplify the overall reinsurance market connectedness. Finally we cross-validate the network centrality rankings against independent market-based tail-risk measures and BVAR stress-test propagation. The topological hubs identified by the network coincides with the firms flagged as systemically important by these independent measures, confirming that centrality captures an economically meaningful dimensions of systemic risk rather than an abstract topological property. At the same time, no firm exhibits a positive capital shortfall under any scenario, indicating that systemic risk in reinsurance arises from interconnection and shock propagation rather than from financial leverage or capital insufficiency.

Introduction

Reinsurance and Systemic Risk

The global reinsurance market, with dedicated capital exceeding $766 billion and written premiums approaching $1.75 trillion, constitutes a critical layer of the financial system's risk-absorption infrastructure (Gallagher Re 2024; International Association of Insurance Supervisors 2025). By accepting risk from primary insurers and redistributing it through retrocession chains and insurance-linked securities, reinsurers create a dense web of bilateral obligations that spans continents, asset classes, and regulatory jurisdictions. While this interconnectedness enhances capital efficiency and enables the underwriting of risks that would otherwise be uninsurable, it simultaneously creates channels through which localized distress can propagate across the network.

The systemic relevance of the insurance sector remains contested. Traditional analyses, grounded in the pre-funded, loss-indemnity nature of insurance liabilities, argue that core insurance activities pose limited systemic risk because insurers do not engage in maturity transformation and policyholder claims are contingent rather than demand-payable (Cummins and Weiss 2014; The Geneva Association 2010). However, this assessment breaks down when institutions engage in non-traditional activities and, critically, when the interconnectedness of reinsurance and retrocession networks creates contagion channels that are opaque to outside observers (Battiston et al. 2012; Chen, Iyengar, and Moallemi 2013). The failure of a single node in such a network can trigger cascading reassessments of counterparty exposure, amplifying the initial shock beyond what aggregate statistics would predict (Park and Xie 2014).

The reinsurance sector is particularly suited to network-based systemic risk analysis for three reasons. First, unlike banking, where bilateral exposures are increasingly reported to trade repositories, reinsurance cession data remain proprietary and unavailable to researchers, motivating the use of market-based proxies to infer connectedness (Billio et al. 2012). Second, the sector's exposure to both asset-side risks (investment portfolio losses during equity or bond market declines) and liability-side risks (catastrophe claims, pandemic losses) creates distinct transmission channels whose systemic implications differ qualitatively. Third, the concentration of global reinsurance capacity in a small number of firms raises the question of whether these hubs constitute Systemically Important Financial Institutions (SIFIs) whose distress would impose externalities on the broader financial system (Battiston et al. 2016; V. V. Acharya et al. 2016).

Despite a growing body of work on systemic risk in banking and insurance (Billio et al. 2012; Adrian and Brunnermeier 2016; Brownlees and Engle 2016;

V. Acharya, Engle, and Richardson 2012), the reinsurance sector has received comparatively limited attention. Existing studies tend to treat reinsurers as a subset of the broader insurance industry rather than analyzing the sector's unique network structure, or focus on single measures of systemic risk without cross-validating across complementary methodologies. This thesis addresses that gap.

Research Questions and Contributions

This thesis investigates three interrelated questions:

1. **Network structure**: What is the topology of the global reinsurance volatility network, and which firms occupy structurally critical positions as variance transmitters, receivers, or bridges?
2. **Dynamic evolution**: How does the network's connectedness, topology, and firm-level hierarchy evolve over time, and do structural breaks in the system align with identifiable economic events?
3. **Systemic importance**: Do the firms identified as topological hubs through the network analysis also rank as the most systemically important when assessed through independent tail-risk measures (ΔCoVaR, MES) and shock propagation scenarios?

The methodological contribution is the cross-validation of Generalized Forecast Error Variance Decomposition (GFEVD) derived network centrality against independent market-based tail-risk measures, applied to a single, consistently defined sample of 27 global reinsurers over the period October 2015 to May 2025. The GFEVD, embedded within a Bayesian Vector Autoregression (BVAR) model with hierarchical Minnesota priors, produces a directed, weighted adjacency matrix that approximates the network of volatility spillovers (Diebold and Yılmaz 2014; Giannone, Lenza, and Primiceri 2015). This network is analyzed using graph-theoretic methods to identify centrality, community structure, and topological properties both statically and through a rolling-window estimation of 349 overlapping windows. The BVAR-derived network measures are then cross-validated against market-based systemic risk indicators to test whether topological prominence translates into genuine systemic importance.

The analysis confirms this hypothesis: network centrality and the independent tail-risk measures converge on the same European core, establishing that topological prominence is a genuine marker of systemic importance. It further reveals a dissociation between systemic centrality and financial fragility, locating the sector's systemic risk in its interconnection rather than in balance-sheet leverage.

Thesis Outline

The thesis is organized as follows.

Chapter 1 provides the institutional background on the reinsurance industry, covering the mechanics of proportional and non-proportional contracts, co-reinsurance and layering structures, and the current state of the global reinsurance market. It concludes by establishing the theoretical link between reinsurance interconnectedness and systemic risk, motivating the empirical analysis that follows.

Chapter 2 presents the methodology and data. The sample of 27 publicly listed reinsurers is described, including sample selection criteria, coverage, and limitations such as survivorship bias and non-synchronous trading. The chapter details the Yang-Zhang drift-independent volatility estimator used to construct the input series, the BVAR specification with Minnesota, Sum-of-Coefficients, and Single-Unit-Root priors estimated via Metropolis-Hastings MCMC, and the GFEVD framework that transforms the estimated model into a connectedness matrix.

Chapter 3 reports the static network analysis. The GFEVD-derived adjacency matrix is analyzed for global topology, node-level centrality, and community structure via Louvain detection. Robustness is assessed across multiple sparsification thresholds. The chapter identifies European reinsurers as the dominant variance contributors and documents a disassortative, globally integrated network with low modularity.

Chapter 4 extends the analysis to a dynamic setting through rolling-window BVAR estimation. It traces the evolution of the Total Connectedness Index across 349 windows, identifies structural breaks via Bai-Perron and OLS-CUSUM tests, and introduces a persistent-versus-reactive centrality classification that uncovers systemically important firms invisible to the static analysis. Five event studies test how different transmission channels (asset-side versus liability-side) reshape the network hierarchy.

Chapter 5 complements the topological analysis with market-based systemic risk measures. ΔCoVaR quantifies each firm's marginal contribution to system-wide tail risk; dynamic MES captures vulnerability under extreme market stress; SRISK combines market sensitivity with balance-sheet leverage to estimate expected capital shortfall. A BVAR-based stress testing exercise via Generalized Impulse Response Functions produces a Systemic Impact Index for each firm. The chapter concludes with a cross-validation of the network centrality measures against these independent systemic risk indicators, testing the central hypothesis that topological hubs are also the most systemically important institutions.

Note that all source code, data, and scripts for this thesis are available on the project's GitHub repository: adb-coding/network-connectedness-volatility-spillover.

An interactive HTML version of this thesis is available online at the project website.

Part I

Thesis

Chapter 1

Introduction to Reinsurance

1.1 Reinsurance

Financial theory suggests that the use of reinsurance, under the conditions of capital market equilibrium, is redundant and not necessary for diversified and widely held firms. These entities are able to pool risks internally and access capital markets for financing when needed, representing a reinsurance specific application of the broader principle of corporate hedging irrelevance under frictionless capital markets (Modigliani and Miller 1958; Doherty and Tinic 1981). However, as shown in Mayers and Smith (1982), market imperfections such as bankruptcy costs, tax convexity, and agency conflicts create a rational corporate demand for insurance and risk transfer. Froot, Scharfstein, and Stein (1993) further demonstrate that costly external financing makes hedging strategies value-enhancing: by smoothing internal cash flows, reinsurance enables insurers to fund investment opportunities without resorting to expensive capital markets. These theoretical foundations establish reinsurance as a rational response to real-world frictions. In this context, reinsurance plays an instrumental role in enhancing diversification beyond the boundaries of solely insurance, providing financial and operational flexibility in addition to the transfer of risk.

From a financial perspective, reinsurance impacts an insurer's surplus management and financing capacity. Since insurers must establish unearned premium reserves upon collecting premiums, their ability to underwrite new policies is limited by surplus levels. Reinsurance mitigates such constraint by reducing the amount of unearned premium reserves requirements, increasing the underwriting capacity, both in terms of single-risk exposure and the volume of contracts written. This allows insurers to underwrite policies that might otherwise exceed their risk tolerance or capital base, and to remain competitive in high-volume lines of business.

An additional benefit is the stabilization of loss experience over time. Due to economic, environmental, and portfolio-specific factors, underwriting results can vary significantly on a yearly basis. Reinsurance enables insurers to smooth these fluctuations, functioning similarly to a contingent capital mechanism whereby reinsurers absorb part of the losses in bad years and receive premiums in better years (Webb, Property, and Underwriters 1984). This rationale is formalized by Plantin and Rochet (2007), who model the optimal demand for reinsurance as a function of the insurer's liability structure and regulatory capital constraints, showing that reinsurance serves as a risk transfer mechanism rather than a generic contingent capital facility.

Catastrophe protection is another key driver for reinsurance usage. In the event of natural disasters or other extreme events, reinsurance shields the primary insurer from potentially severe financial consequences. Furthermore, reinsurance provides underwriting assistance through access to specialized expertise, advanced actuarial data, and claims management capabilities. This knowledge is especially valuable for insurers that expand into new product lines, geographic territories, or unconventional risk classes. Reinsurance also facilitates the entry and exit from specific markets, facilitating the cessation of operations in certain lines of business, allowing the transfer of entire portfolios without sudden contract cancellations, which could result in reputational and regulatory backlash.

Before proceeding further, the next sections briefly describe the types of reinsurance protections and contracts that characterize the business model, to better understand the advantages they offer, but also the complexities that the underwriting process includes. We start by introducing the different types of agreements in reinsurance.

1.2 Facultative and Treaty Reinsurance

Reinsurance agreements can be broadly categorized as *facultative* or *treaty* reinsurance, which differ in application, risk exposure, contractual complexity, and strategic purpose. Facultative reinsurance covers individual risks on a case-by-case basis, while treaty reinsurance involves the cession of all risks within a defined class of policies through a predetermined agreement (Outreville 1998). The distinction reflects different risk-management strategies and degrees of control retained by the primary insurer: facultative provides highly customized coverage for large or unusual risks, whereas treaty emphasizes automation and scalability, transferring homogeneous risks with reduced transactional friction.

1.2.1 Facultative Reinsurance

Facultative reinsurance is a transactional form in which the ceding insurer offers a specific individual risk to a reinsurer, who has the discretion to accept or decline all or part of it (Thompson 1966).

Each contract is negotiated separately and each risk individually underwritten, with no binding obligation on either party: the primary insurer is not required to cede the risk, nor the reinsurer to accept it. This mutual optionality makes facultative reinsurance a flexible tool for non-standard exposures or risks falling outside existing treaty agreements.

Since each risk is addressed individually, the reinsurer can tailor terms, pricing, and conditions to the specific exposure, enabling accurate risk selection and premium adequacy in high-severity or complex underwriting situations such as industrial facilities, large infrastructures, or natural catastrophes.

This granularity comes at a cost: each contract entails extensive underwriting review and detailed documentation, often across several rounds of negotiation with multiple reinsurers. Facultative placement is therefore labor-intensive and administratively burdensome, making it less suitable for high-frequency or commoditized lines of business. It also supplements treaty arrangements, covering exposures that exceed treaty limits or fall outside their scope and thereby reducing risk concentration within the insurer's book

In practice, facultative agreements are generally short-term and often structured as single-event transactions tied to the duration of the insurance policy, providing transaction-specific coverage for exceptional risks.

1.2.2 Treaty Reinsurance

Treaty reinsurance is a blanket arrangement whereby the reinsurer agrees in advance to accept a defined portion of all risks within a specified portfolio of policies written by the ceding company. This automatic acceptance makes it the pivot for efficient, scalable risk transfer in the industry (Rejda and McNamara 2021).

The treaty specifies the business and risks covered, attachment points, limits, premium-sharing quotas, and claim procedures, and is legally binding over a fixed period. It is designed to improve operational efficiency and to foster long-term partnerships supporting the growth of the insurer's underwriting activity.

By dispensing with risk-by-risk negotiation, treaty arrangements reduce administrative cost and transaction time, letting primary insurers focus on core underwriting. They smooth loss exposure over time and foster deeper cedant–reinsurer relationships, often bundling value-added services such as actuarial support, product development, and portfolio analysis.

The main drawback is adverse selection - the cedant may load disproportionately risky policies into the treaty - which is why reinsurers require detailed portfolio analytics and due diligence before committing. A second limitation is the lack of flexibility: once in force, the treaty applies uniformly across the portfolio, leaving little scope to exclude or tailor individual risks. Despite this, treaty reinsurance remains the preferred method for managing routine, high-volume exposures.

Reinsurance agreements, whether in facultative or in treaty form, can be categorized as proportional or non-proportional type.

1.3 Proportional and Non-proportional Reinsurance

Under proportional reinsurance, one or more reinsurers take a pre-agreed percentage share of the premiums and liabilities of each policy, so the reinsurer's premium share is directly proportional to its claims obligation; the reinsurer also pays the cedant a commission, expressed as a percentage of the original premium, toward administrative and acquisition costs.

Non-proportional forms of reinsurance - meaning coverage with no fixed or pre-determined split of premiums and losses - are mainly used to provide protection against the biggest and accumulated losses that can potentially jeopardize insurer's solvency. Non-proportional reinsurance requires that all losses up to a certain amount are covered by the primary insurer, while claims exceeding that amount are covered by the reinsurer up to a pre-agreed limit.

1.3.1 Proportional

The most frequently used proportional agreements are the *quota share* and the *surplus share* agreement.

1.3.1.1 Quota-share reinsurance

Quota share contracts are a form of proportional reinsurance, in which a fixed percentage of the primary insurer's risk is transferred to one or more reinsurers. This agreement typically applies to the entire portfolio under the treaty. The reinsurer receives a premium rate proportional to the percentage of claims covered, minus a ceding commission paid to the insurer as a way to cover the acquisition and administrative costs resulting from the acquisition of the business (Outreville 1998).

Assuming a payout function Y for a policy p, $Y^p = P - X^P$, where P represents the written premium, and X^P is the random loss with a maximum claim of x_{max}

$$X^P = \begin{cases} x & \text{in the case of claim} \quad 0 < x \leq x_{max} \\ 0 & \text{otherwise} \end{cases}$$

the payout functions can be represented as:

$$Y^i = (1 - \Lambda)P - (1 - \Lambda)X^P + c\Lambda P \tag{1.1}$$

for the primary insurer and

$$Y^r = \Lambda P - \Lambda X^P - c\Lambda P \tag{1.2}$$

for the reinsurer, where $\Lambda \in (0, 1)$ is the retention ratio and $(1 - \Lambda)$ is the ceded ratio by the primary insurer, while c is the commission rate paid by the reinsurer to the insurer as a percentage of the written premiums.

Quota-share smooths underwriting results and reduces volatility in the insurer's retained loss distribution, making it well suited to high-volume, homogeneous lines such as property, liability, or business interruption insurance. Its drawback is the loss of profitability from ceding a significant share of premiums.

1.3.1.2 Surplus-share reinsurance

Surplus share reinsurance is another form of proportional reinsurance but differs in its conditional application. Instead of applying a fixed percentage to all policies, the insurer retains a fixed monetary amount, known as the *retention limit*. Any amount exceeding the retention limit is ceded to the reinsurer. The proportion of intervention for each policy is determined by the ratio of the surplus (the amount exceeding the retention), to the total amount written by the policy. This implies that the ceded proportion varies by policy, depending on the size of the exposure.

This type of agreement is advantageous for insurers that want to retain full control over smaller risks while sharing larger exposures that could potentially undermine their solvency or capital. Surplus-share, compared to quota-share reinsurance, allows for more detailed control over which risks are transferred and greater management of the loss distribution's volatility.

Figure 1.1 provides a graphical representation of the differences between the two types of proportional agreement: figure a shows the scenario of the quota-share, with $\Lambda = 40\%$; figure b shows the amounts ceded and retained, for the same policies, in the case of a surplus-share agreement. Only policies above the retention line will see the intervention of the reinsurer.

Figure 1.1: Comparison of the Quota-share and Surplus-share reinsurance contracts.

1.3.2 Non-proportional

Non-proportional reinsurance agreements typically take one of these three forms:

- Excess-of-loss
- Catastrophe
- Stop-Loss

1.3.2.1 Excess-of-loss

Excess-of-loss (XoL) reinsurance is a type of non-proportional reinsurance, in which the reinsurer provides coverage only for losses that exceed a pre-specified threshold, known as the *priority* or *retention*, up to an agreed *limit*. This arrangement is typically used to mitigate the financial impact of large infrequent losses.

The reinsurer bears the portion of claims x that lies between the priority d and the limit M, resulting in the following payout structure:

$$\text{Reinsurer Payout} = \min(M, \max(x - d, 0)) \tag{1.3}$$

For instance, in a "€10M XS €5M" contract, the primary insurer retains losses up to €5M, while the reinsurer covers the following €10M in losses. Losses greater than €15M are uncovered by the reinsurer, and will be borne by the primary insurer.

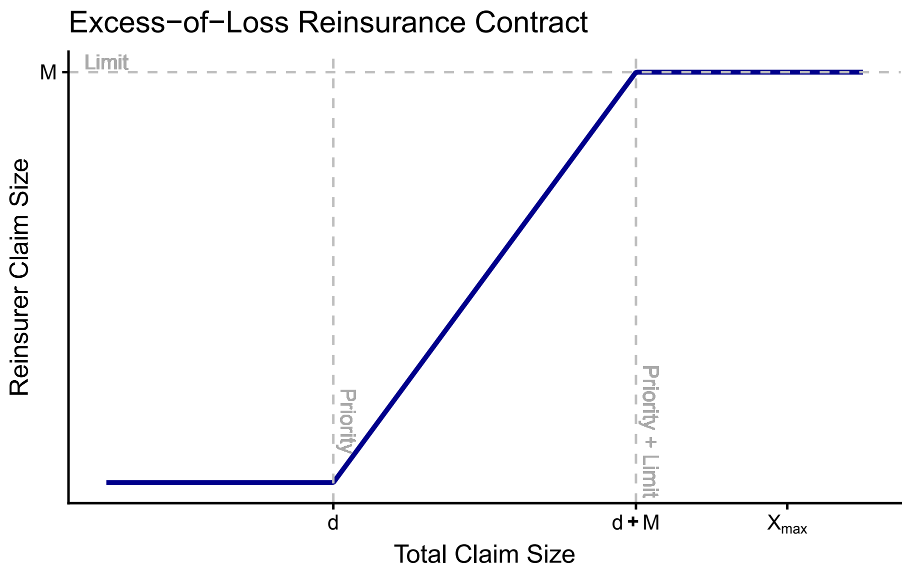

Figure 1.2: Function of the losses faced by the reinsurer when the total loss crosses the priority level. Reinsurer is obliged to repay up to the limit amount agreed with the insurer.

XoL contracts diverge into two different categories, to serve two separate types of needs: excess of loss *per-risk*, and excess of loss *per-event*. The former is the traditional reinsurance product in which the reinsurer indemnifies the primary insurer for the loss amount of all the individual policies affected defined in the treaty agreement who exceed the priority amount. The reinsurer's contribution is limited by its cover obligations, and typically an annual limit of the total amount refundable is established. The per risk agreement is effective to mitigate risks against large single losses, but is inadequate for protection against high frequency or cumulative losses. To address this lack of capability, excess of loss per-event contracts are more suitable, as the unit of loss is not the individual loss per policy, but rather the aggregate loss caused by an event within the insurance portfolio covered by the treaty. This arrangement represents an effective risk reduction against large losses made up of the sum of potentially hundreds of small losses (SwissRe 2015).

Excess-of-loss contracts are preferred for their efficiency in capital management, especially in portfolios with high exposure volatility (Boyer and Dupont-Courtade 2015), and are used especially to provide coverage against events or risks of an unusual nature, which involve potential losses that can undermine the solvency of the primary insurer.

1.3.2.2 Catastrophe Reinsurance

Catastrophe reinsurance (Cat-XL), is another type of non-proportional reinsurance arrangement that aims at covering the portfolio of insurers against claims arising from a single accident, such as natural disasters (earthquakes, hurricanes, floods) or other large-scale man-made accidents. Catastrophe reinsurance addresses the correlation of risks where a single event triggers the simultaneous claims across multiple policies, potentially threatening the insurer's solvency. A catastrophe is usually defined in terms of a given number of claims within a time interval of a given duration, or based on the amount of total claims caused by the single event.

Transfer of risk arising from catastrophic event among market participants is essential if we look at the total economic losses caused in 2024: according to the Banerjee et al. (2025) report, the amount was \$328 billion, for a total insured losses of \$146 billion. Around 43% of the total losses, and insured losses arising from natural catastrophes totaled \$137 billion, with a year-on-year increase from 2023 of \$22 billion, in line with the 5-7% annual growth rate (in real terms) of the recent years. Major causes of insurance loss are still the "major perils" events, such as earthquakes and hurricanes, that typically mark the "peak loss" years, like in 2011 due to the Japan earthquake (\$169 B insured losses) or in 2005 for the Katrina hurricane (\$166 B insured losses). However, recent years have seen an increase in the impact of "secondary perils" events, such as severe convective storms (SCS), floodings, or wildfires, due to the increase in severity and frequency of such events, with studies providing clear evidence for the role of man-made climate change in this transition (Munich Re 2025).

1.3.2.3 Stop-loss

Stop-loss reinsurance is another form of non-proportional reinsurance, in which the reinsurer covers the excess of the aggregate losses of a portfolio over a pre-agreed threshold. Stop-loss reinsurance provides protection against the accumulation of numerous small or medium-sized losses that could undermine the financial capability of the insurer.

In a stop-loss agreement, the retention level is defined as a fixed monetary amount or as a loss ratio, that is, a percentage of the earned premiums from the covered business. The reinsurer only becomes liable if the total incurred losses exceed the defined retention level over the contract period. Typically, the contract also identifies a limit over which payments are no longer provided. This type of protection is suitable especially for lines of business with high frequency but typically low losses (like motor insurance), new or volatile lines of business, or for insurers seeking to stabilize underwriting results in adverse years.

Mathematically, the repayment function of the insurer X^i and of the reinsurer X^r can be expressed as follows:

$$X^i = \begin{cases} X^p, & \text{if } X^p \leq \Lambda \\ \Lambda, & \text{if } \Lambda < X^p \leq \Lambda + L \\ X^p - L, & \text{if } X^p \geq \Lambda + L \end{cases} \tag{1.4}$$

$$X^r = \begin{cases} 0 & \text{if } X^p \leq \Lambda \\ X^p - \Lambda & \text{if } \Lambda < X^p < \Lambda + L \\ L & \text{if } X^p \geq \Lambda + L \end{cases} \tag{1.5}$$

where X^p denotes the total aggregate losses of the primary insurer's portfolio under the contract, Λ is the retention level (defined here as a fixed monetary amount) of the primary insurer, and L is the maximum coverage limit provided by the reinsurer. Note that here Λ denotes a monetary retention level, distinct from the dimensionless retention ratio $\Lambda \in (0, 1)$ used in equations 1.1 and 1.2. The same symbol is retained following actuarial convention.

Stop-loss is among the most complex arrangements to negotiate: the covered business must be precisely defined to monitor the retention level, and losses must be aggregated across multiple policies and lines. Reinsurers therefore require detailed actuarial projections, robust reserving, and frequent audit monitoring (SwissRe 2015).

1.4 Co-reinsurance and Layering

Reinsurance contracts involve high capital commitments, uncertainty regarding potential losses, and complex negotiations due to information asymmetries, agency problems, and moral hazard concerns. To mitigate concentration risk and improve capital efficiency, insurers typically adopt strategies to spread risks among multiple reinsurers. That is why most reinsurance agreements involve the use of *co-reinsurance* and *layering*.

1.4.0.1 Co-reinsurance

Co-reinsurance refers to the practice in which multiple reinsurers underwrite the same layer of a given risk, each assuming a specified percentage share of the liability. The reinsurers involved in the agreement participate collectively in both claim settlement and premium sharing, based on the proportion of the contract.

This approach allows the primary insurer to reduce its counterparty risk, by not relying on a single reinsurer for a specific layer, and is particularly used in the coverage for high-severity portfolios with large aggregate exposures. For example, consider an insurer ceding any losses exceeding €100,000 to a co-reinsurance agreement involving three reinsurers. If the loss incurred by the insurer is €280,000, the €180,000 in excess will be split equally between reinsurers, with a €60,000 claim each. However, co-reinsurance can also be structured with asymmetric participation, in which the share of the excess loss underwritten by each reinsurer varies. Figure 1.3 provides a graphical representation of the different contributions of reinsurers in the symmetric and asymmetric agreements for the same claim size.

1.4.0.2 Multi-layering

In contrast to co-reinsurance, multi-layering involves structuring the reinsurance program in distinct tranches to multiple reinsurers, who cover different layers of loss. These contracts are typically of non-proportional form, such as excess of loss or stop-loss reinsurance.

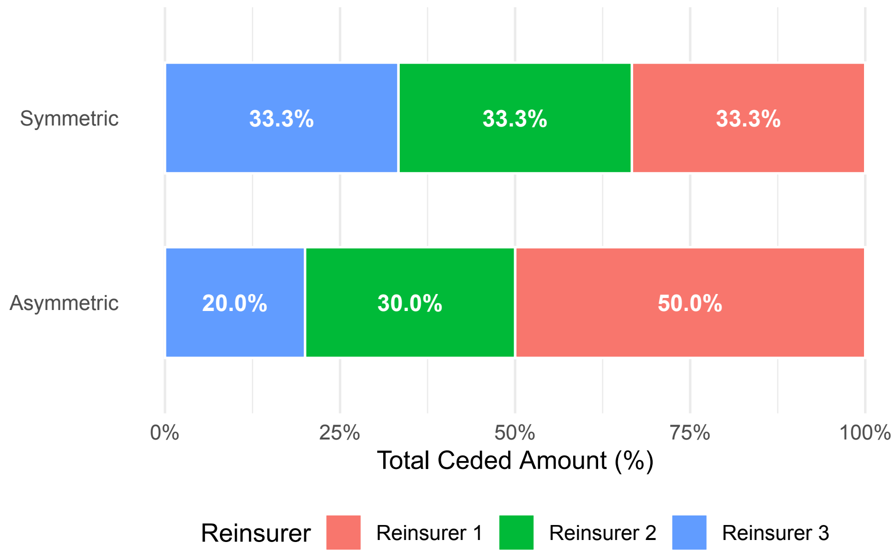

Figure 1.3: Representation of the division of claims between the reinsurers in the case of symmetric and asymmetric co-reinsurance.

Under a layering strategy, the primary insurer retains losses up to a predefined retention level. Losses that exceed this level are passed sequentially to reinsurers responsible for specific ranges or layers. Thus, each reinsurer faces different risk exposures, with different severity and probability, depending on the covered layer. This setup is ideal to construct a tailored risk-transfer structure, with low-severity claims being managed internally or through lower layers, while high-severity, but with low-frequency losses are transferred to reinsurers willing to assume tail risks, as represented in Figure 1.4. Separating total liability into distinct layers increases the attractiveness of the entire contract, since it allows each reinsurer to choose the degree of volatility and exposure it seeks, and also reduces each reinsurer's risk, since they are only responsible for the part of the loss included in the layer (Boyer and Dupont-Courtade 2015).

1.5 The Reinsurance Market

The previous sections have outlined the functional mechanics of reinsurance, from facultative placement to complex treaty layers. However, to fully grasp the role of reinsurance in the global economy, one must view it not as a collection of isolated bilateral agreements, but as a dense, interconnected market structure, absorbing volatility for primary insurers and redistributing it across a global network of capital.

1.5.1 Market Volume and Capitalization

As of 2025, the global reinsurance market has continued its trajectory of robust growth. The total dedicated reinsurance capital has reached approximately $766

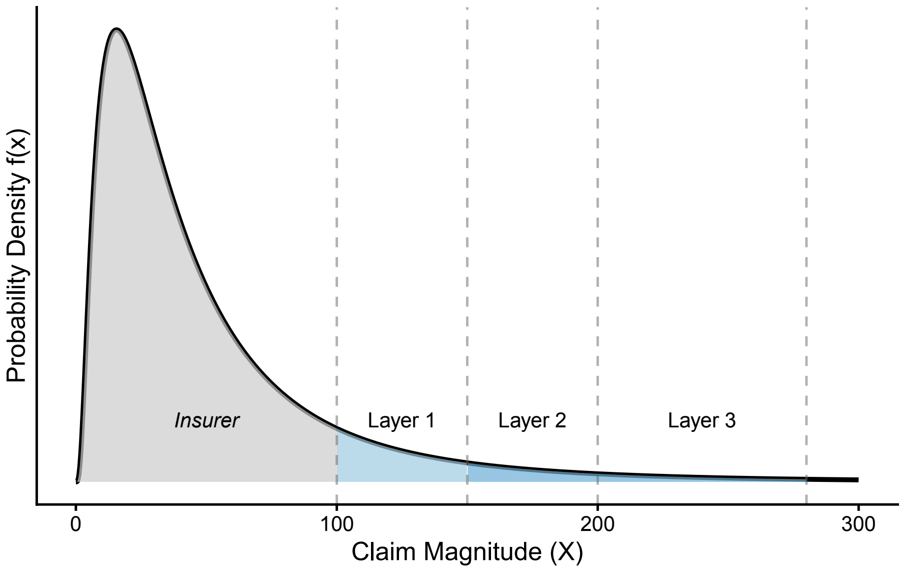

Figure 1.4: Probability density of the random claim payment of a policy with multiple layers covered by different reinsurers. Higher layers correspond to higher claim size and higher risk taken.

billion, driven by both traditional equity capital and a surging alternative capital structure (Gallagher Re 2024). The market has seen a particular expansion in third-party capital, which is estimated to exceed $144 billion, largely fueled by record issuance of Catastrophe Bonds (Cat Bonds) and other Insurance-Linked Securities (ILS) (Banerjee et al. 2025). This influx of capital has led to a "softening" of pricing in certain retrocession markets during the mid-2025 renewals, creating a favorable environment for buyers despite the increase in frequency of loss events.

With the rise of ILS, the distinction between traditional reinsurance and capital market has blurred. By securitizing insurance risks into tradeable assets like Cat Bonds, the industry has tapped into a deeper pool of institutional capital. While this collateralized retrocession mitigates some counterparty risk, it also links the insurance market more directly to broader financial market trends, potentially correlating insurance losses with financial market liquidity shocks.

Also the "retrocession spiral" phenomenon, in which reinsurers retroceed the risk to multiple reinsurers, can obscure the true accumulation of risk. Although retrocession is vital for capital efficiency, it introduces significant counterparty risk and opacity. As noted by SwissRe (2015) and expanded in recent network literature (Battiston et al. 2012), the failure of major retrocessionaires could trigger a cascade of defaults, transmitting financial distress across the entire network, as happened in the London Market Spiral in the late 1980s. This interconnectedness makes the reinsurance market a complex financial network, where the stability of the whole depends also on the resilience of its most connected and exposed nodes.

It is worth noting, despite the new business model and financial products that insurance and reinsurance companies have begun to offer, their core model remains substantially different from that of the typical banking sector. The consequences for the regulatory regimes, the supervisory authorities, and the systemic risk that characterize the industry therefore differ significantly. As banks take in short-term callable deposits and convert them into long-term, illiquid loans, borrowing short and lending long, insurance and reinsurance companies do the reverse. They collect premiums up front and pay claims later, which makes these institutions pre-funded by construction. Moreover, their liabilities are not callable on demand, as banks' deposits are, but are paid only in the event of a claim. Substitutability is also higher than in the banking sector: if an underwriter fails, its capacity can usually be replaced by competitors.

These differences between the banking and (re)insurance business models imply that the two sectors are governed by different regulatory regimes, with different systemic measures and supervisory activities. The next section briefly introduces the regulatory architecture of the (re)insurance sector.

1.5.2 Regulatory Framework

Reinsurance is not governed by a dedicated "reinsurance regulator"; reinsurers are instead supervised as insurance undertakings under the same prudential regimes that apply to primary insurers, with an additional systemic risk overlay imposed when a firm is large and internationally active. Four overlapping layers can be distinguished:

- The **Financial Stability Board (FSB)** determines whether and how the insurance sector is treated as a source of global systemic risk, and it owns the cross-sector resolution standard embodied in the *Key Attributes of Effective Resolution Regimes*;
- The **International Association of Insurance Supervisors (IAIS)** builds the technical framework to assess the systemic importance and draft the Insurance Core Principles;
- Regional and jurisdictional prudential regimes, such as Solvency II in the European Union and the UK, or the NAIC[1] in the US, set the binding capital and supervisory requirements that apply to firms;
- Finally, a layer of group-wide supervision for **internationally active insurance groups (IAIG)**, which sits above solo-entities and is the specific target of the Insurance Capital Standards.

The regulatory assessment of the insurance sector and its systemic risk has evolved over the years. Initial investigations of key systemic institution in the reinsurance sector were performed using an entity-based approach. The FSB, applying the IAIS methodology, designated nine Global Systemically Important Insurers (G-SIIs), keeping this list updated from 2013 to 2016. However, no reinsurer was ever identified as G-SII. Systemic importance was scored against five categories under a weighting scheme that assigned 45% of the weight to

[1]The transposition was necessary to comply with the usage of the R package `igraph`, used to perform the presented analysis.

Non-Traditional Non-Insurance activities (NTNI) and a further 40% to interconnectedness, with size accounting only for 5% of the weight, inverting the G-SIB methodology used to identify systemically important banks.

As the entity-based approach received numerous criticisms, since 2019 the IAIS adopted an *activity-based* approach, recognizing that systemic risk can arise from the collective exposures of many insurers rather than only from the disorderly failure of one. It rests on an enhanced set of macroprudential supervisory measures embedded in the Insurance Core Principles; on an annual, data-driven Global Monitoring Exercise that scans the sector for build-ups of risk; and an assessment of whether supervisors can apply the measures.

Running alongside the systemic-risk overlay, the European Union and the UK adopted the Solvency II to constrain firms. It is based on three pillars - quantitative requirements calibrated to a 99.5% one-year Value-at-Risk, internal control and Own Risk and Solvency Assessment, and mandatory disclosure - and subjects reinsurers to the same requirements of primary insurers. Reinsurance serves as a risk mitigation tool, allowing a cedant to reduce its capital requirements, while simultaneously creating a new counterparty risk, captured in the counterparty-default-risk module. The system becomes safer as the risk is not eliminated, but spread into the system.

However, these regulatory frameworks focus on the ability of reinsurers to sustain an extremely rare loss, calibrating their measures to solvency rather than to contagion. The work that follows tries to measure the contagion effects.

1.5.3 Future Challenges

Before proceeding further, it is notable to identify the challenges that the reinsurance sector faces in 2026 and in the coming years, as the multi-crisis environment with traditional natural catastrophe risks are compounded by emerging and evolving threats:

- **Climate Change and Secondary Perils**: The industry is witnessing a structural shift where secondary perils are becoming primary drivers of loss. Global insured losses from natural catastrophes are projected to reach $145 billion in 2025, with climate change acting as a "permanent stress test" on portfolio resilience (Banerjee et al. 2025).

- **Cyber Risk**: Cyber coverage represents a unique systemic threat due to its lack of geographical boundaries and rapidly evolving technologies. The cyber insurance market is expanding rapidly, with global written premiums amounting to $15 billion for the year 2024, an increase of 7% from the previous year (National Association of Insurance Commissioners 2025). Between 50 to 65% of the premiums coming from cyber security policies are ceded to reinsurers, highlighting the crucial role of the reinsurers' sector in sustaining market capacity, but increasing the potential for accumulation risk — where a single event (e.g., a cloud service outage) triggers simultaneous claims globally — and creating difficulties in monitoring the portfolios of primary insurers, posing concrete challenges to control aggregation and capacity risks.

- **Artificial Intelligence (AI)**: AI presents a dual-edged challenge. While it enhances underwriting efficiency, it also introduces new liability landscapes (e.g., algorithmic bias, deepfakes) and empowers more sophisticated cyber-attacks, creating volatile risk profiles that lack historical data for accurate pricing.

The complexity described above challenges traditional "siloed" analysis methods. The debate regarding the systemic nature of reinsurance is ongoing. Traditional views, grounded in the pre-funded, loss-indemnity nature of insurance liabilities, argue that core insurance liabilities pose limited systemic risk (Cummins and Weiss 2014; The Geneva Association 2010). However, as Cummins and Weiss (2014) demonstrate, this assessment breaks down when insurers engage in non-traditional activities and, critically, when interconnectedness of reinsurance and retrocession networks creates channels through which localized distress can propagate.

The reinsurance structures described in the previous sections translate directly into such contagion channels. When a primary insurer cedes risk via a proportional treaty, its solvency becomes partially contingent on the reinsurer's ability to honor claims. In layered excess-of-loss programs, the failure of a reinsurer covering a layer of the agreement exposes the cedant to unhedged tail risk. The retrocession spiral compounds this: risk is not eliminated but redistributed through bilateral obligations, where the ultimate bearer of risk is often unidentifiable from publicly available data (Battiston et al. 2012; Park and Xie 2014). As Chen, Iyengar, and Moallemi (2013) show, this network opacity means that the failure of a node can trigger cascading reassessment of counterparty exposure, amplifying an initial shock beyond what aggregate statistics would predict.

The lack of transparency of retrocession networks poses a fundamental empirical challenge: bilateral reinsurance exposures are proprietary and unavailable by researchers. This motivates the use of market-based data to infer connectedness. Following Billio et al. (2012), who demonstrate that statistical associations in stock returns capture meaningful linkages between financial institutions, this thesis employs equity return volatility as a proxy for the latent risk transfer relationships described above. It must be acknowledged, however, that Billio et al. (2012) study *banks and hedge funds*: in that context, equity co-movement plausibly reflects interbank lending and derivative exposure. For reinsurers, co-movement may instead reflect shared catastrophe-event exposure, common investment portfolio marks, or index membership. V. V. Acharya et al. (2016) further validate market-based systemic risk measures for the banking sector; their applicability to reinsurers rests on the non-traditional activity channels documented by Cummins and Weiss (2014), which create bank-like contagion pathways. While equity volatility remains an imperfect proxy, it constitutes the best available market-observable signal given the proprietary nature of reinsurance treaty data.

Therefore, the following chapters adopt the Generalized Forecast Error Variance Decomposition (GFEVD) framework of Diebold and Yılmaz (2014), embedded within a Bayesian Vector Autoregression (BVAR) model. The GFEVD decomposes the forecast error variance of each reinsurer's volatility into contributors from all others, yielding a directed, weighted adjacency matrix that approximates the network of volatility spillovers. The resulting network is then analyzed using graph-theoretic methods to identify which reinsurers occupy structurally critical positions, and whether the global reinsurance market exhibits the hub-dominated, disassortative topology that would make it vulnerable to targeted shocks (Battiston et al. 2016).

Chapter 2

Methodology and Data

2.1 Methodology

As discussed before, the reinsurance industry is founded on the principle of risk sharing among market participants. While this mechanism provides various benefits, as stated in Chapter 1, it also introduces interdependencies among participants that could spread financial distress, creating substantial systemic risk. Although financial networks are frequently analyzed and subject to research, the existence of systemic risk arising from interconnection and interdependence is still debated, especially for insurance and the reinsurance sector. The goal of this research is to construct a network based on forecast error variance decomposition in order to quantify risk spillovers, assess the stability of the system, and identify systemic risk in the industry.

The research proceeded by first selecting possible internationally influential reinsurance companies and collecting market price data during the analysis period. Yang-Zhang volatility was then calculated on the retrieved data. A statistical BVAR model was chosen, after comparing it with frequentist analogous VAR models. The connectivity measure deriving from the generalized forecast error variance decomposition is calculated, to later construct and analyze the network. The methodology used and the empirical results are reported in this chapter.

2.2 Data Selection

The original dataset of reinsurance companies considered consists of the top 50 global reinsurance groups for the year 2024 according to the rating agency *A.M. Best*[1]. The retained dataset was obtained considering only companies that satisfied two criteria:

- are publicly listed on a stock exchange
- their continuous daily trading data are available on Yahoo Finance for the period October 2015 to March 2025.

[1]The transposition was necessary to comply with the usage of the R package `igraph`, used to perform the presented analysis.

Of the original 50 groups, 27 companies were retained, while 23 were excluded. Table 2.1 shows the reinsurers that were considered for the network construction, the Country in which they are established, the gross written premiums (GWP), and the market capitalization (Mkt Cap) expressed in billion USD. Of the 27 companies composing our dataset, 5 are set in the American continent (4 in the US, 1 in Canada), 6 are set in Bermuda, 8 in Europe, 7 in Asia, and 1 in Australia, ensuring to capture various markets' dynamics.

For the other 23 companies excluded, 15 are not listed or private companies, 5 are listed companies but their trading data are available for a period shorter than the required one, one was delisted as it was acquired by another financial company, and two were removed as their ticker did not purely represent the reinsurance sector. Detailed information of the excluded companies is reported in Table A.2 in Section A.1.

2.2.1 Market Coverage

The 27 retained companies account for an aggregate gross written premium (GWP) of $573.32 billion. According to the International Association of Insurance Supervisors (2025), written global life and non-life reinsurance premiums amounted to approximately $1.75 trillion for the year 2024, making the sample representative for roughly one third of the whole market volume. In addition, A.M. Best estimated GWP was $265 billion. The difference with the values reported in Table 2.1 mainly comes from the fact that these data represent group-level gross written premiums, which include substantial primary insurance operations for several conglomerates. Their inclusion is justified as their reinsurance divisions are ranked within the A.M. Best top 50, and that group-level stock prices - which are the input to our volatility estimation - reflect the market's assessment of the entire group, including its reinsurance exposure. This is consistent with the approach of Billio et al. (2012), who use publicly traded equity prices to measure systemic interconnectedness, and with the Diebold-Yilmaz framework (Diebold and Yılmaz 2014), which relies on market-observable proxies rather than balance-sheet data.

2.2.2 Limitations

Before proceeding further, the data selection poses some limitations that are worth emphasizing.

2.2.2.1 Survivorship bias

A limitation of this sample construction is survivorship bias. The dataset inherently excludes firms that were active during the sample period considered but subsequently exited the market, either through insolvency, merger, or delisting. Such firms may also constitute the most relevant cases for the study of systemic risk, as financial distress and corporate failure are manifestations of the intrinsic systemic risk of the industry. Notable exclusions include XL Group (acquired by AXA in 2018, partially captured through CS.PA post-acquisition) and PartnerRe (delisted in 2015 following acquisition by EXOR). The use of a survivor-conditioned sample may therefore lead to an underestimation of tail-risk spillovers and an overly benign view of network stability during crisis periods.

Table 2.1: Reinsurers considered for the network construction.

Ticker	Name	Country	Mkt Cap	GWP
003690.KS	Korean Reinsurance Co	South Korea	1.34	5.39
0966.HK	Taiping Reinsurance Co. Ltd.	Hong Kong	6.92	1.11
1508.HK	China Reinsurance (Group) Corp	China	8.53	7.85
6809.T	The Toa Reinsurance Co, Limited	Japan	0.25	2.30
8630.T	Sompo International Holdings, Ltd.	Japan	26.68	26.83
8725.T	MS&AD Insurance Group Holdings, Inc.	Japan	32.04	35.34
8766.T	Tokio Marine & Nichido Fire Insurance Co., Ltd.	Japan	76.77	36.30
ACGL	Arch Capital Group Ltd.	Bermuda	33.25	11.11
AXS	AXIS Capital Holdings Limited	Bermuda	7.34	2.39
BRK-A	Berkshire Hathaway Inc.	US	1070.25	88.69
CB	Chubb Limited	Switzerland	112.18	62.00
CS.PA	AXA XL	France	74.46	8.41
EG	Everest Group, Ltd.	Bermuda	14.72	12.94
FFH.TO	Allied World Assurance Co Holdings, Ltd.	Canada	38.91	7.20
G.MI	Assicurazioni Generali SpA	Italy	43.37	81.84
HNR1.DE	Hannover Rück SE	Germany	26.60	22.68
HSX.L	Hiscox Ltd	Bermuda	3.43	1.03
MAP.MC	MAPFRE RE, Compañía de Reaseguros S.A.	Spain	10.77	6.72
MKL	Markel Corp	US	24.30	1.15
MUV2.DE	Munich Reinsurance Company	Germany	62.00	34.42
QBE.AX	QBE Insurance Group Limited	Australia	47.59	22.40
RGA	Reinsurance Group of America Inc.	US	12.81	18.49
RNR	RenaissanceRe Holdings Ltd.	Bermuda	12.04	11.73
SCR.PA	SCOR S.E.	France	4.63	17.25
SPNT	SiriusPoint Ltd.	Bermuda	2.09	1.34
SREN.SW	Swiss Re Ltd.	Switzerland	54.45	44.86
WRB	W.R. Berkley Corp	US	28.89	1.55

2.2.2.2 Data quality consideration

The use of Yahoo Finance as the data source introduces additional considerations. Stock prices across the 27 companies are quoted on 8 different exchanges and spanning multiple time zones. To address non-synchronous trading, only trading days on which all exchanges were open simultaneously are retained, reducing the sample from approximately 2,500 potential trading days to 1,989 common trading days. While this approach eliminates stale-price contamination, it discards information from days when only specific markets experience significant price movements. Furthermore, Yahoo Finance data may not consistently account for corporate actions (stock splits, special dividends) across all exchanges, although the use of adjusted closing prices and log-transformation applied during volatility estimation mitigate the impact of level shifts.

2.2.3 Volatility estimation

Volatility is a crucial element for the estimation of the network and the subsequent findings about spillovers effects in the system. The volatility estimator used in this framework is the one developed by Yang and Zhang (Yang and Zhang 2000), in which an unbiased estimator based on multiple periods of open, close, high and low prices is used. It is assumed that price movements can be modeled as a geometric Brownian motion with undetermined parameters, volatility σ and drift μ. The estimator is shown to be the minimum-variance unbiased variance estimator, also being independent of the drift motion μ, and from the opening price jumps, due to the multi-period consideration. The formula for the Yang-Zhang variance is defined as:

$$V_{YZ} = V_o + kV_c + (1-k)V_{RS} \tag{2.1}$$

where k is a parameter defined as $k = \frac{0.34}{1.34+\frac{n+1}{n-1}}$. Letting O_i, C_i, H_i,L_i be the opening, closing, high, and low of the period considered, the formula can be further decomposed to its core components: the term V_o is the overnight variance, defined as:

$$V_o = \frac{1}{n-1}\sum_{i=1}^{n}(o_i - \bar{o})^2$$

in which $\bar{o} = \frac{1}{n-1}\sum_{i=1}^{n} o_i$ and $o_i = ln(O_i) - ln(C_{i-1})$ is the normalized open price, to account for jumps in the stock quotations, otherwise not considered. The second term of the Yang-Zhang variance is the open-to-close volatility, V_c, used to account for intra-day price variations:

$$V_c = \frac{1}{n-1}\sum_{i=1}^{n}(c_i - \bar{c})^2$$

where $\bar{c} = \frac{1}{n-1}\sum_{i=1}^{n} c_i$ and $c_i = ln(C_i) - ln(O_i)$ is the normalized closing price. The latter element to calculate the variance estimator is the Rogers-Satchell variance V_{RS} that gives independence of the estimator from the drift component. It is computed as:

$$V_{RS} = \frac{1}{n}\sum_{i=1}^{n}[u_i(u_i - c_i) + d_i(d_i - c_i)]$$

where $u_i = ln(H_i) - ln(O_i)$ and $d_i = ln(L_i) - ln(O_i)$ are respectively the normalized high and the normalized low of the trading day. The period used to calculate the estimator was 21 days (approximately one month of data). To normalize the results, the log transformation was applied to the resulting variance (V_{YZ}) to ensure the absence of unit-roots.

The use of this estimator, instead of the more common close-to-close volatility measure, is mainly driven by the fact that the Yang-Zhang estimator preserves the beneficial properties of the original estimator, but at the same time it increases efficiency, allowing for more accurate estimation with less data, and reduces dependence on the drift component which could significantly impact volatility estimations. In addition, V_{YZ} reflects volatility drops that occur when a stock undergoes a steady one direction movement, known as trendy phase. The volatility dataset was then used to estimate a Bayesian Vector autoregressive model, as discussed in the following section. However, while the Yang-Zhang estimator achieves minimum variance under geometric Brownian motion, reinsurance equities frequently exhibit jump-diffusion behavior due to sudden catastrophe loss announcements. Consequently, the Yang-Zhang estimator serves here as a robust, drift-independent baseline proxy for underlying variance, rather than a perfect measure of continuous-path volatility.

2.2.3.1 Stationarity test

Before proceeding with the BVAR estimation, the stationarity of the log-volatility series was assessed using the Augmented Dickey-Fuller (ADF) test (Dickey and Fuller 1979), in which the null hypothesis is the presence of a unit-root, and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test (Kwiatkowski et al. 1992), where the null hypothesis is level-stationarity.

Table A.3 in the Section A.2 reports the results for all the 27 log-volatility series. The ADF test rejects the unit-root null hypothesis at the 5% significance level for all the series, confirming the absence of I(1) behavior, while the KPSS test fails to reject level-stationarity at 5% significance level for only one of the 27 series. This implies that the stationarity of the series is ambiguous, due to the failed level-stationarity test. This pattern is consistent with the long-memory properties of financial volatility (Andersen et al. 2003), where slowly decaying auto correlations cause the KPSS test to over-reject even in absence of unit roots. We point out that strict level-stationarity is not a necessary condition for VAR estimation, as the critical requirement is the absence of unit roots. Furthermore, the BVAR specification already accommodates near-unit-root behavior through the Sum-of-Coefficients (SOC) and Single-unit-root (SUR) dummy priors, which shrink the model towards unit-root compatible dynamics without requiring strict stationarity of the input series (Sims 1993; Giannone, Lenza, and Primiceri 2015).

2.3 BVAR

Vector autoregressive (VAR) models are flexible time series models widely used in macroeconomics and finance to analyze the interrelationships among multiple variables. However, as the number of endogenous variables increases, the number of parameters in an unrestricted VAR increases quadratically, often leading to overfitting, unstable inference, and faulty out-of-sample forecasts. Bayesian VAR (BVAR) models address this issue by introducing a prior distribution that shrinks the unrestricted model toward a more parsimonious specification, improving estimation stability and predictive accuracy (Litterman 1986; Bańbura, Giannone, and Reichlin 2010; Giannone, Lenza, and Primiceri 2015).

A VAR model of finite order p, in its standard form, can be expressed as:

$$y_t = a_0 + A_1 y_{t-1} + ... + A_p y_{t-p} + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \Sigma) \tag{2.2}$$

where y_t is an $M \times 1$ vector of endogenous variables, a_0 is a $M \times 1$ vector of the intercept, $A_j, (j = 1, \dots, p)$ are $M \times M$ coefficient matrices, ε_t is a $M \times 1$ vector of exogenous normally distributed shocks with zero mean and variance-covariance matrix Σ.

Under flat (noninformative) priors, the posterior distribution of the VAR coefficients is centered at the ordinary least squares (OLS) estimator. Nevertheless, since the number of coefficients increases rapidly with M and p, OLS becomes unreliable in high dimensions. Bayesian inference treats the unknown parameters as random variables, and updates beliefs using the observed data. Let $\theta = (A_1, \dots, A_p, a_0, \Sigma)$ denote the VAR parameters, and let γ denote the hyperparameters determining the prior distribution. Bayes' rule yields:

$$p(\theta, \gamma | y) \propto p(y|\theta, \gamma)\, p(\theta|\gamma) p(\gamma) \tag{2.3}$$

$$p(y|\gamma) = \int p(y|\theta) p(\theta|\gamma) d\theta \tag{2.4}$$

where $p(y|\theta, \gamma)$ is the likelihood, used to explore the full posterior hyperparameter space. The introduction of a prior structure allows the BVAR to shrink toward economically plausible values, such as persistence close to a random walk, thus mitigating the overfitting while maintaining flexibility for forecasting (Miranda-Agrippino and Ricco 2018).

2.3.1 Priors Implementation

BVAR models require specifying prior distributions for the VAR coefficients. Priors may incorporate information from economic theory, historical data, or empirical regularities. A widely adopted prior is the Minnesota prior, often supplemented with additional dummy-observation priors to regularize the deterministic components of the VAR.

2.3.1.1 Minnesota Prior

The Minnesota prior is built on the empirical observation that many macro-financial variables are persistent and often well approximated by a univariate random walk. It embodies the economic beliefs that recent observations are more informative than distant ones, own lags are more relevant than lags of other variables, and that many financial time series show near unit-root behavior.

Let β denote the vectorized VAR coefficients, $\beta = vec(A_1, \dots, A_p, a_0)$, with dimension $K = M^2 p + M$. The Minnesota prior specifies

$$\beta \sim \mathcal{N}(\bar{\beta}, \Omega)$$

where $\bar{\beta}$ is the prior mean and Ω is the covariance matrix, and are structured as follows.

- **Mean structure**: for the coefficient on the k-th lag of variable j in the equation of variable i, denoted as $a_{ij}^{(k)}$, the prior mean is:

$$E[a_{ij}^{(k)}] = \begin{cases} \delta_i & \text{if } i = j \text{ and } k = 1 \\ 0 & \text{otherwise} \end{cases} \tag{2.5}$$

  where $\delta_i = 1$ corresponds to a random walk prior and $\delta_i = 0$ to a white noise prior.

- **Variance structure**: the prior variance for the coefficient $a_{ij}^{(k)}$

$$Var(a_{ij}^{(k)}) = \begin{cases} \frac{\lambda^2}{k^\alpha} & \text{if } i = j \\ \frac{\lambda^2 \sigma_i^2}{k^\alpha \psi_j} & \text{if } i \neq j \end{cases} \tag{2.6}$$

  where:

  - $\lambda > 0$ controls the overall tightness of the prior, so smaller values impose a stronger shrinkage toward the prior mean
  - $\alpha > 0$ controls the rate of variance decay across lags, setting how quickly the prior tightens for more distant observation
  - ψ_j controls the prior standard deviation, over lags of variable j
  - σ_i^2 represents the error variance of the variable i

The values for the hyperparameters of Minnesota prior are selected according to the methodology described in Section 2.3.2, and the values used for the final model are reported in Section A.6.

2.3.1.2 Dummy Observation

In addition to the Minnesota prior, to further regularize the BVAR model, dummy observations are added, imposing additional structure on the long-run properties of the VAR, reducing the importance of deterministic components estimated using the initial observations (Giannone, Lenza, and Primiceri 2015). These priors operate by appending artificial dummy observations to the data, effectively modifying the likelihood in a principled manner, as found in Sims (1993).

SOC The Sum-of-coefficients prior addresses the concern that, without additional regularization, traditional VAR models may produce implausible long-run forecasts. This prior imposes the notion that a no-change forecast is optimal at the beginning of a time series, so it shrinks the sum-of-coefficients towards unity, reinforcing persistence while avoiding explosive behavior. The SOC dummy observation is constructed as follows:

$$Y^{SOC} = \frac{1}{\mu} \cdot diag(\bar{y})$$

$$X^{SOC} = [0, \frac{1}{\mu}, \bar{y}, \dots, \bar{y}]$$

where $\bar{y}$ represents an $M \times 1$ vector of means over the first p observations of each variable, and $\mu > 0$ controls the tightness of this prior. As $\mu \to \infty$, the prior becomes uninformative, while as $\mu \to 0$, the model is pulled toward a specification with as many unit roots as variables and no cointegration.

The intuition is that if all lags of variable j sums to 1, then a permanent unit shock to variable j will have a permanent unit effect on itself. The hyperparameter μ determines how strongly this property is applied.

SUR The single-unit-root prior complements the function of the SOC prior, imposing that the VAR system admits at least one unit root. The SUR prior takes the form of:

$$Y^{SUR} = \frac{1}{\delta} \cdot \bar{y}^T$$
$$X^{SUR} = [\frac{1}{\delta}, \frac{1}{\delta}\bar{y}^T, \dots, \frac{1}{\delta}\bar{y}^T]$$

where $\bar{y}^T$ is the row vector of $\bar{y}$ introduced before, and $\delta > 0$ controls the tightness.

This prior shrinks the model towards a specification where the sum of all coefficients across all variables and lags equals the identity matrix plus a constant, preserving long-run co-movement. This is economically motivated in environments where variables are believed to share common stochastic trends.

2.3.1.3 Metropolis-Hastings Algorithm

Given the structure of our priors specification, with hyperparameters $\gamma = (\lambda, \alpha, \mu, \delta)$ treated as random, posterior inference, as in Equation 2.3, requires sampling from the joint posterior distribution $p(\theta, \gamma|y)$. However, the marginal posterior of the hyperparameters $p(\gamma|y)$ generally lacks a closed form solution, requiring the use of the Metropolis-Hastings (MH) algorithm to sample from the marginal posterior (Chib and Greenberg 1995). A deeper explanation about the process used by this algorithm are shown in Section A.4.

Because the posterior distribution is not analytically tractable, the estimation is carried out using a Markov Chain Monte Carlo (MCMC) algorithm, where the combination of Metropolis-Hastings sampling for the hyperparameters and Gibbs sampling for the VAR coefficients allows us to draw from the joint posterior $p(\theta, \gamma|y)$, properly accounting for uncertainty not only in the VAR parameters, but also in the hyperparameters controlling shrinkage intensity, which is crucial for hierarchical BVAR estimation.

2.3.2 Model Selection and Hyperparameter Optimization

Given the reliance of BVAR models on the specifications of prior beliefs, selecting appropriate hyperparameters accurately is a crucial step to avoid overfitting and over-shrinkage of the model. To determine the optimal model specification, a systematic grid search procedure was performed, comparing the predictive performance of various BVAR models, against frequentist VAR models that served as benchmark. The model selection process involved estimating BVAR models across a range of values for the hyperparameters described in the previous section. The grid search focused primarily on the:

- Lag length (p): models were estimated using lags ranging from 1 to 6, in order to obtain an ideal trade-off between model-complexity and information capture.
- Overall tightness (λ): the starting value of the hyperparameter λ were tested at $[0.05, 0.1, 0.2, 0.50]$, to assess various level of informative behavior.
- Variance decay (α): the rate at which distant lags are discounted was tested with $\alpha \in [1, 2]$, representing linear and harmonic decay.

Other settings of the hyperpriors for the Minnesota and dummy priors were left unchanged, as well as the behavior of the Metropolis-Hastings algorithm.

In total, 48 models were estimated, with $15,000$ total draws for each model, of which $5,000$ were used for the burn-in phase, and the remaining $10,000$ draws were saved. Each model was assessed using an in-sample forecast of 10 days horizon. The dataset was therefore split into a training set, excluding the last 10 observations, that served as the test dataset. The same operation was carried out for the VAR models, where the only variable changed was the lag order.

The primary performance metric used to assess model performance was the Root Mean Squared Error (RMSE) of the forecast. For comparative analysis, the difference in RMSE ($\Delta RMSE$) was calculated as:

$$\Delta RMSE = RMSE_{BVAR} - RMSE_{VAR}$$

where a negative value of $\Delta RMSE$ indicates that the Bayesian model has outperformed the frequentist VAR.

2.3.2.1 Model Diagnostics

Statistical diagnostics were evaluated to ensure the stability and reliability of the posterior distribution, calculated using the MCMC method. For every estimated model, the following diagnostics were analyzed:

- Acceptance rate: the proportion of proposed parameter draws accepted by the sampler. The target acceptance range was set between 0.20 and 0.40. Extremely low rates indicate a stuck chain, where no draws are accepted in the posterior distribution, while extremely high rates suggest that the sampler is not adequately exploring the parameter space.

- Geweke Convergence diagnostic: this test compares the mean of the first 10% of the chain against the last 50%. If the samples are drawn from the stationary distribution of the chain, the two means are equal and Geweke's statistic has an asymptotically standard normal distribution

$$Z = \frac{\bar{\theta_1} - \bar{\theta_2}}{\sqrt{Var(\bar{\theta_1}) + Var(\bar{\theta_2})}} \sim N(0,1) \tag{2.7}$$

  A value of $|Z| < 1.96$ indicates that the chain has converged to a stationary distribution.

- Effective Sample Size (ESS): represents a measure of the number of independent samples equivalent to the autocorrelated MCMC chain (Geweke 1992).

$$ESS = \frac{N}{1 + 2\sum_{k=1}^{\infty} \rho_k} \tag{2.8}$$

where ρ_k is the lag-k autocorrelation of the chain. It is used to assess the quality of the convergence of MCMC samples, with higher values indicating better convergence.

2.3.2.2 Hyperparameter Selection

The comparative results revealed distinct patterns in predictive performance relative to lag length, where models with smaller lags (1 to 3) presented underfitting behavior, with RMSE relatively high (> 0.59), although Bayesian models consistently outperformed their frequentist counterpart (negative $\Delta RMSE$). This suggests that these models were too parsimonious to capture the full dynamic relationships between the variables in the system.

The higher lags tested (from 5 to 6), instead, showed overfitting or overshrinkage behavior. Indeed, the BVAR models' performance deteriorated relative to the VAR models of equivalent lags, with a positive $\Delta RMSE$. This is likely due to the Bayesian priors imposing excessive shrinkage on distant lags. Figure 2.1 provides a graphical representation of the $RMSEs$ relative to the number of lags used.

The analysis identified the model with 4 lags, prior belief $\lambda = 0.2$ and linear decay $\alpha = 1$ as the optimal specification (*model 28*). This model achieved the lowest predictive error ($RMSE_{BVAR} = 0.588$) while maintaining a performance advantage over the frequentist VAR model, with a negative $\Delta RMSE = -0.15$.

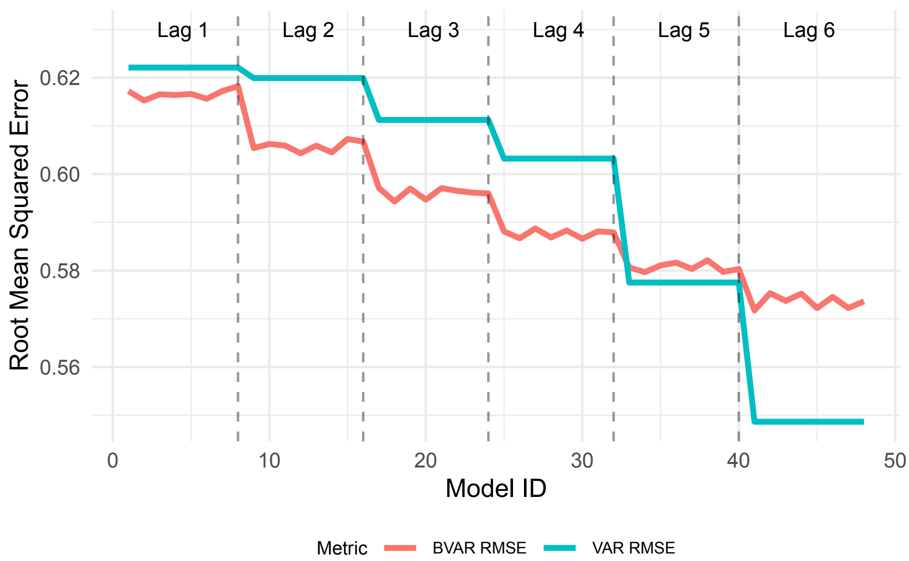

Figure 2.1: Comparison of the Root Mean Squared Error ($RMSE$) for both the BVAR and VAR models, with different lag values: BVAR models perform better (lower $RMSE$) compared to frequentist VARs, up to lag 4, where the difference ($\Delta RMSE$) is greater; for higher lags, VAR models have better predictive performance.

The full prior and MCMC specifications for the final model used in the remaining work are reported in Section A.6

The MCMC process resulted in a posterior distribution with a value of $\lambda = 0.191$, with an acceptance rate of $\alpha_{MH} = 0.259$, falling just inside the acceptance region, indicating a proper exploration of the parameter space. The posterior mode of $\lambda = 0.191$ remains close to the prior mode (0.20). In a highly over-parametrized 27-variable system, this indicates the prior is appropriately regularizing the likelihood. Further diagnostic was performed and the results are reported in Section A.6, along with the source code used to estimate the BVAR model and all the specifications of the priors used.

Having found the model estimation satisfactory, the output was then processed to estimate the forecast error variance decomposition, in order to construct the network. The next section addresses the procedure adopted and the results obtained.

2.4 GFEVD

The adjacency matrix later used to construct and analyze the network was done applying the forecast-error variance decomposition (FEVD), developed in a series of papers by Francis X. Diebold and Kamil Yilmaz. Following Diebold and Yilmaz (2012) findings, the use of generalized forecast-error variance decomposition (GFEVD) allows to generate a spillover measure that at first is independent of variable ordering, and second is able to examine directional spillovers.

Let us define own variance shares as the fractions of the H-step ahead error variance in forecasting y_i that are due to shocks to x_i, for $i = 1, 2, \dots, N$, and spillovers (cross-variance shares) as the fractions of the H-step ahead error variances in forecasting x_i that are due to shocks to x_j, for $i, j = 1, 2, \dots, N$, such that $i \neq j$. Denoting with $\theta^g_{ij}(H)$, for $H = 1, 2, \dots$, the H-step ahead forecast error variance decomposition entries, as developed in the generalized VAR framework of Koop, Pesaran, and Potter (1996) and Pesaran and Shin (1998), we have:

$$d^g_{ij}(H) = \frac{\sigma_{jj}^{-1} \sum_{h=0}^{H-1} (e'_i \theta_h \Sigma e_j)^2}{\sum_{h=0}^{H-1} (e'_i \theta_h \Sigma \theta'_h e_i)} \tag{2.9}$$

where Σ is the covariance matrix for the shock vector in the non-orthogonalized BVAR model, θ_h is the coefficient matrix multiplying the h-lagged shock vector in the infinite moving-average representation of non-orthogonalized shocks, σ_{jj} is the j-th diagonal element of Σ, and e_i and e_j are the selection vectors for the i-th and j-th variables.

However, the sums of the row elements of the variance decomposition are not necessarily equal to 1. In order to use the results in the calculation of the spillover index, we normalize the sum of the elements of each row:

$$\hat{d^g_{ij}}(H) = \frac{d^g_{ij}(H)}{\sum_{j=1}^{N} d_{ij}(H)}$$

By construction $\sum_{j=1}^{N} \hat{d^g_{ij}}(H) = 1$ and $\sum_{i,j=1}^{N} \bar{d^g_{ij}}(H) = N$, therefore, after normalization, the generalized variance decomposition matrix D^{gH} is an $N \times N$ matrix with entries $[d^{gH}_{ij}]$.

Table 2.2: Simplified representation of the connectedness table of the Diebold and Yilmaz framework

[]

	x_1	$\cdots$	x_N	From others	
x_1	d_{11}^{gH}	$\cdots$	d_{1N}^{gH}	$\sum_{j=1}^{N} d_{1j}^{gH}$	$j \neq 1$
$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	
x_N	d_{N1}^{gH}	$\cdots$	d_{NN}^{gH}	$\sum_{j=1}^{N} d_{Nj}^{gH}$	$j \neq N$
To others	$\sum_{i=1}^{N} d_{i1}^{gH}$	$\cdots$	$\sum_{i=1}^{N} d_{iN}^{gH}$	$\frac{1}{N} \sum_{i,j=1}^{N} d_{ij}^{gH}$	
	$i \neq 1$	$\cdots$	$i \neq N$	$i \neq j$	

The off-diagonal elements of D^{gH} are the values of the N forecast error variance decomposition that are relevant from a connectedness perspective, since they represent the pairwise directional connectedness between the variables considered in the system. The pairwise directional connectedness from j to i, is therefore defined as:

$$C_{i \leftarrow j}^{H} = d_{ij}^{H}$$

The sums of off-diagonal rows and columns give two other important measures. The row sum of the off-diagonal elements gives the share of the H-step forecast error variance of variable i coming from shocks arising in other variables ("*from others to i*") , while the column sum of the off-diagonal elements produces the forecast error variance that arises in variable j and is transmitted to the other variables ("*to others from j*"). The two measures are defined as:

$$C_{i \leftarrow \blacksquare}^{H} = \sum_{j=1}^{N} d_{ij}^{gH} \qquad j \neq i$$

$$C_{\blacksquare \leftarrow j}^{H} = \sum_{i=1}^{N} d_{ij}^{gH} \qquad i \neq j$$

These additional measures, in extension to the already calculated GFEVD matrix, compose the connectedness table, which is represented schematically in Table 2.2. The sum of the column labeled "From others" represents the total connectedness measure C^H:

$$C^{H} = \frac{1}{N} \sum_{i,j=1}^{N} d_{ij}^{gH}, \quad \text{with} j \neq i \tag{2.10}$$

The total connectedness index effectively summarizes the degree of interdependence within the system, providing a single measure of the contribution of spillover effects to the total forecast error variance. This metric, combined with the directional pairwise connectedness measures, constitutes the primary output of the GFEVD framework used to assess systemic risk.

To address parameter uncertainty in the GFEVD estimates, the decomposition is computed not only at the posterior median but across all accepted MCMC draws (after thinning by factor 10, yielding 1,000 posterior draws). This yields a full posterior distribution for every element of the connectedness matrix and for the Total Connectedness Index. The plug-in TCI estimate at the posterior mean of the VAR coefficients is $C^H = 74.28$, consistent with the Diebold and Yılmaz (2014) convention. Propagating parameter uncertainty through the MCMC posterior yields a median TCI of 77.09 with a 90% credible interval of [74.03, 80.61], confirming that the high level of interconnectedness is robust to BVAR parameter uncertainty. The upward shift from plug-in estimate reflects Jensen's inequality applied to the nonlinear GFEVD transform. Posterior credible intervals for individual node-level out-strength (risk-spillovers to others) are reported in Table A.5 in Section A.5.

Before turning to the empirical reconstruction of the network, it is worth making explicit what the GFEVD-based connectedness measures do, and do not, represent economically. The pairwise entries of the connectedness matrix quantify the share of the H-step-ahead forecast error variance of one firm's volatility that is statistically associated with shocks to another firm's volatility. By construction, they are a measure of statistical association in volatility innovations within a reduced form VAR, and not a measure of causal or contractual transmission between institutions. A high pairwise value indicates that the two firms' volatilities co-move in a way that is not accounted for by the rest of the system; it does not, on its own, identify the mechanism through which that co-movement arises. This distinction matters for interpretation, as the generalized framework of Koop, Pesaran, and Potter (1996) and Pesaran and Shin (1998) yields order-invariant, directional variance shares, but these directions reflect predictive content in the reduced-form residuals rather than an economically identified flow of risk from one firm to another.

Several economically distinct mechanisms can generate the same observed spillover, and the connectedness matrix cannot, in isolation, extricate them. First, the linkages may reflect common exposure to shared risk factors, such as catastrophe and natural-hazard events, interest-rate and term-structure movements, and broad capital-market shocks, to which reinsurers are jointly exposed and which induce contemporaneous co-movement in their equity volatility without any direct link between the firms. Second, they may capture information transmission and learning, whereby news about one firm's loss experience, reserves, or capital position is rationally extrapolated by investors to peers with comparable business models. Third, they may embody correlated market reactions and sentiment, where common shifts in risk appetite, index membership, or sector-wide repricing move reinsurer valuations in tandem. Fourth, and only partly, they may trace genuine underlying economic and contractual links (retrocession and risk-transfer chains, shared cedents, and overlapping investment portfolios), which are the relationship that this network aims to proxy. Because equity volatility connectedness aggregates all four channels, the estimated network should be read as an indirect, market-implied proxy of interdependence among reinsurers rather than as a direct map of contractual exposure. In line with Billio et al. (2012) and Diebold and Yılmaz (2014) framework, the analysis relies on market-observable proxies precisely because the direct contractual data are not available, a constraint discussed later in the Chapter 6.

Having presented the econometric and mathematical foundations - ranging from the Yang-Zhang volatility estimation to the Bayesian VAR specification and generalized variance decomposition the methodology allows for a robust reconstruction of the reinsurance network. The empirical results derived from this framework, including the estimated connectedness table and the topological analysis of the network, are presented in the following chapter.

Chapter 3

Static Network Analysis

3.1 Introduction

In this chapter, we conduct static analysis of the reinsurance network constructed using the methodology presented previously. We start by analyzing the outcomes of the BVAR model on the reinsurers' volatility and the connectedness table derived from the generalized forecast error variance decomposition (GFEVD), then we move to the network analysis of the reinsurance market, reporting the global topology measures and the connections between market players. We proceed then to a micro-prudential analysis of the network, assessing the key central reinsurers and their systemic importance, to later move on community detection and cluster analysis.

In this research the network is defined as a graph $G = (V, E)$ where

- V is the set of nodes representing the reinsurance companies selected for the study
- E is the set of edges representing pairwise forecast error variance contributions between companies

The weighted directed network is constructed on the adjacency matrix deriving from the BVAR model applied to the Yang-Zhang log-volatility and the GFEVD over a time horizon of 10 days. To construct the network topology, in which the directions of the edges represent the explanation of variance (from source to target), the adjacency matrix was transposed[1]. Consequently, a directed edge from company j to company i reflects the proportion of i's forecast error variance explained by j, ensuring that the out-strength captures a firm's variance contribution. The graph G_{full} consists of $|V| = 27$ nodes and $|E| = 702$ edges. The connectedness table for all nodes is reported in Appendix B. The total connectedness index, which represents the total level of interconnectedness within the network, was $C^H = 74.278$. A value of this magnitude means that roughly three-quarters of the total volatility forecast-error variance in the system is cross-firm spillover and only about a quarter is idiosyncratic: from the market's standpoint the global reinsurance sector acts almost as a single risk pool.

[1] The transposition was necessary to comply with the usage of the R package `igraph`, used to perform the presented analysis.

Figure 3.1 illustrates the distribution of edge weights in the G_{full}. The distribution exhibits positive skewness, with a mean of $\mu = 2.857$ ($\sigma = 2.118$). The majority of pairwise connections are relatively weak, while a small number of edges in the right tail represent the strongest variance contribution relationships.

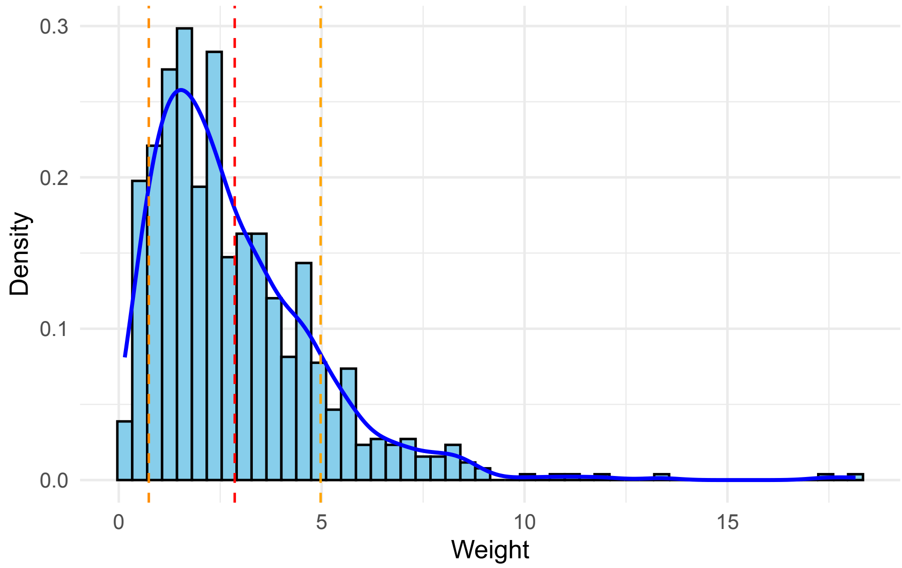

Figure 3.1: Distribution of the pairwise connectedness weights in G_{full}. The dashed red line indicates the mean, dotted orange lines represent the 1 standard deviation. The right-skewed distribution shows that most connections are weak, with few strong linkages driving system connectedness.

Table 3.1 reports the 15 strongest pairwise connections of the network, with the highest value estimated at 18.157. Interestingly, the first 6 records are all between Japanese companies, with values > 10, while all the other strongest connections are between European companies. This indicates a strong statistical connectedness within the Japanese market, where the forecast error variance is predominantly explained by the variance within the market, while American companies show lower pairwise variance contributions. The Japanese links reflect a tightly coupled, largely self-contained regional block whose firms explain each other's volatility without exporting it to the global system. Instead, European links form the global backbone, since these firms set the worldwide reinsurance pricing and capacity at the major renewals and their volatility leads the system.

Table 3.1: Ranked list of the 15 strongest pairwise connections.

From	To	Value
8766.T	8725.T	18.157
8766.T	8630.T	17.422
8630.T	8766.T	13.348
8725.T	8766.T	12.013
8630.T	8725.T	11.334

Table 3.1: Ranked list of the 15 strongest pairwise connections.

From	To	Value
8725.T	8630.T	10.722
HNR1.DE	MUV2.DE	10.012
HNR1.DE	CS.PA	9.023
CS.PA	G.MI	8.922
HNR1.DE	G.MI	8.550
CS.PA	HNR1.DE	8.529
SREN.SW	MUV2.DE	8.434
HNR1.DE	MAP.MC	8.378
G.MI	CS.PA	8.374
MAP.MC	CS.PA	8.296

The Diebold-Yilmaz framework, by construction, produces a complete network, where all companies are connected to each other and contribute to explaining each other's forecast error variance. To consider only weights that provide some impact over other market participants, we applied a threshold to the weights of the connections, where only the edges with weight above the 50^{th} percentile are considered. The 50th percentile threshold serves as a filter to separate the signal (structural transmission channels) from the noise (negligible spillovers). By focusing on the upper half of the distribution, we isolate the effective network where changes in forecast error variance are economically meaningful. This construction and the topology of the network itself allow us to perform deeper and more detailed analysis, otherwise unattainable. The network obtained, G_{sparse}, still comprises the 27 reinsurance companies, but the number of edges is reduced to 351, and from now on it is the base model considered for further examinations. The choice of the 50th percentile as threshold is validated through robustness analysis presented in the following section, where we demonstrate that the key findings remain stable across alternative thresholds (25^{th} and 75^{th} percentile). Accordingly, all node-level measures reported hereafter are computed on G_{sparse}.

3.2 Robustness Analysis

In this section, we present the robustness analysis to assess whether the key findings that are later presented are sensitive to the choice of threshold and to validate the 50th percentile threshold selection. A robust result should exhibit qualitative consistency (same directional findings) and quantitative stability (high rank correlations) across alternative thresholds. We consider Spearman's rank correlations of companies for the 25th, 50th, and 75th percentiles of the weight distribution. We define a stability threshold of $\rho > 0.85$ as indicating robust stability for the different threshold. This criterion is conservative, exceeding the standard statistical benchmark for a "very strong" relationship ($\rho > 0.8$) proposed by Evans (1996). A correlation above 0.85 ensures that the identification of the most systemically important insurers remains consistent regardless of the thresholding choice.

The total connections between nodes for the three graphs reduce to 526 for the

25^{th} , 351 for the 50^{th}, and 176 for the 75^{th} quantile thresholds, corresponding to an edge density of 75%, 50%, and 25% respectively. The metrics analyzed to test the robustness and stability of the method are the topologies and micro-prudential level statistics, to identify the key variance contributors and recipients, and the community detection and creation. The results obtained were then compared by assessing similarity through rank correlation.

3.2.1 Network Topology

Table 3.2 presents the key topological measures across the different thresholded networks. By construction, the average path length and the diameter vary significantly, due to the different densities of the networks. The 25^{th} percentile network shows a dense interconnected structure with a strong core interconnected with the peripheral nodes ($L = 0.419$, $D = 1.169$). Instead, for the 75^{th} percentile network, the metrics indicate a less efficient flow, with an average path length of $L = 0.421$. The maximum geodesic distance ($D = 1.074$) indicates greater isolation for peripheral nodes.

Strength-based assortativity remained consistent across the three networks, showing slightly disassortative behavior. With a value of $r = -0.146$, the 75^{th} percentile network exhibited the strongest disassortative signal, suggesting that the structure best captures the hub-spoke transmission patterns, while the 25th percentile thresholded network showed a value of $r_{25} = -0.097$, very close to the $r_{50} = -0.116$ of the 50^{th} percentile network. This behavior reflects either noise inclusion (25^{th}), or compression of the strength distribution as peripheral nodes completely lose edges, reducing statistical power to detect mixing patterns.

Table 3.2: Key topological metrics across the 25th, 50th, and 75th percentile thresholds.

Network Threshold	Avg. Path Length	Diameter	Assortativity
50th	0.394	1.187	-0.116
25th	0.419	1.169	-0.097
75th	0.421	1.074	-0.146

3.2.2 Node Centrality

The identification of the top variance contributors and variance recipients remained stable across the different threshold levels, with European companies showing high variance contribution properties, and Asian companies maintaining the role of variance recipients. The top contributors and top recipients are reported in Table 3.3.

Table 3.3: Top variance contributors and top variance recipients, based on the net-strength metric.

Ticker	50	25	75	Rank stability
HNR1.DE	80.05	73.05	67.03	Top transmitter
CS.PA	74.31	62.66	81.05	Top transmitter

Table 3.3: Top variance contributors and top variance recipients, based on the net-strength metric.

Ticker	50	25	75	Rank stability
SREN.SW	67.46	53.17	59.56	Top transmitter
1508.HK	-49.49	-62.15	-32.07	Top receiver
003690.KS	-58.67	-65.92	-12.16	Top receiver
0966.HK	-59.91	-62.27	-34.77	Top receiver

The eigenvector centrality rankings remained stable for the three networks, with a strong Spearman correlation of 0.989 between the 25^{th} and 50^{th} quantile networks, and a correlation of 0.97 between the 50^{th} and 75^{th} quantile networks. The classification of hubs and authorities held a consistency value of 0.98 across thresholds, with 25 of the 27 companies maintaining the original designation, while only Tokio Marine & Nichido Fire Insurance Co. (8766.T) and Hiscox Ltd. (HSX.L) saw their denomination change from hub to authority. As shown in Table 3.4, betweenness correlation decreases at higher thresholds due to the sparser network topology, where peripheral nodes lose connectivity and fewer nodes can serve as bridges between the remaining clusters.

Table 3.4: Spearman correlations of the network centrality measures and their respective p-values.

	$\rho_{25,50}$	p-value	$\rho_{50,75}$	p-value
Net	0.993	0	0.829	0.000
Betweenness	0.735	0	0.743	0.000
Eigenvector	0.989	0	0.970	0.000
PageRank	0.813	0	0.650	0.000
Hub	0.989	0	0.971	0.000
Authority	0.980	0	0.855	0.000
Community	0.891	0	-0.314	0.111

3.2.3 Community Structure

Community detection showed the greatest sensitivity to threshold choice. The sparser network produced more fragmented communities, as QBE Insurance Group (QBE.AX) and Hiscox Ltd. (HSX.L) lost all connections and became isolated nodes. This would imply total independence from the system and no variance influence being received or transmitted from or to other nodes, resulting in a connected network of only 25 companies rather than the initial 27. However, as later presented, community identification resulted in a rather low separation, suggesting a tightly connected core. This finding shows that community detection is most subject to thresholding choice, consistent with the relatively low modularity score. The partition reshapes under thinning because no strong natural clusters exist; the instability of community structure and the integration of the market are thus two views of the same fact.

Given the results of the robustness analysis, the 50^{th} percentile provides an effective balance, avoiding noise of denser networks while preserving sufficient connectivity to maintain an interpretable community structure.

3.3 Global Topology Measure

We start by first analyzing the topology of the network, describing the structural stability of the connections, and identifying the core of the structure. To better understand the connections between nodes and the topology of the network, we provided a graphical representation in Figure 3.2. The graph is represented using a Force-Directed algorithm, more specifically the Fruchterman-Reingold algorithm (Fruchterman and Reingold 1991). At first glance, the network presents a tight core of European companies that have higher out-strength values (bigger node size), and American companies. Asian companies occupy the periphery of the network and typically have lower out-strength, indicating a low or no variance transmission role.

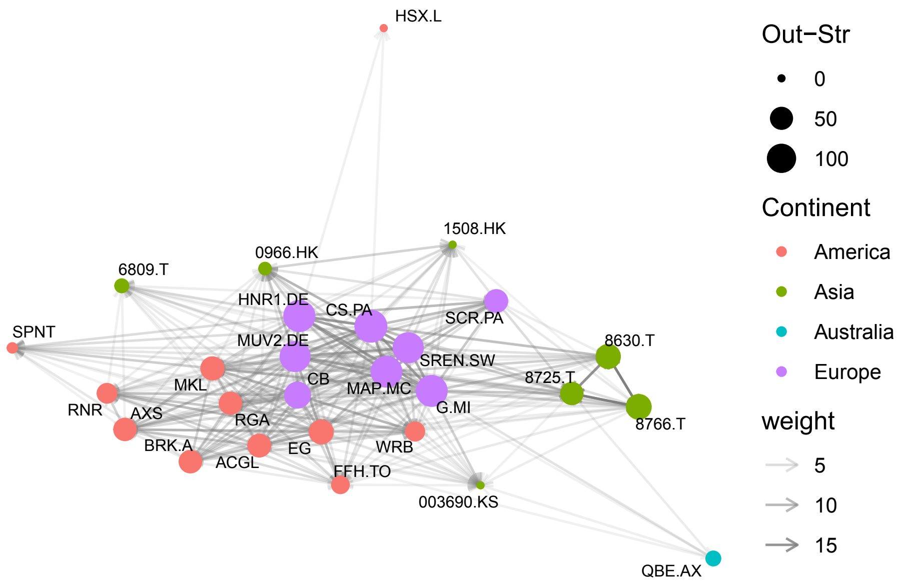

Figure 3.2: Plot of the network using the Fruchterman-Reingold algorithm. Node size is set using the Out-strength value.

3.3.1 Average Path Length

Average path length represents the average number of steps along the shortest paths for all possible pair of nodes in the network.

$$L = \frac{1}{N(N-1)} \sum_{i \neq j} d(i,j)$$

where $d(i,j)$ is the geodesic distance (shortest path) between nodes i and j. In a weighted network, the distance is often defined as the inverse of the weight $1/w(i,j)$, implying that strong links create short distances. This measure captures the efficiency of connectedness pathways: higher values indicate longer pathways, requiring variance contribution to pass through more nodes before reaching the whole system, while smaller values indicate a faster and more direct statistical linkages system-wide.

The average path length of the network is $L = 0.394$, indicating relatively efficient connectedness across the network, where strong variance contribution relationships create short statistical distances across nodes. This relatively short path indicates a tight interconnected network structure, where statistical relationships span the system through few intermediary nodes. Economically, this is a fast-contagion, low-firewall property: a shock that reaches the whole system through few intermediaries, leaving little structural buffering between any firm and the rest of the market. This shows characteristics that are feared by supervisors in crisis scenarios, where transmission is rapid and difficult to break.

3.3.2 Diameter

The network diameter D represents the longest shortest path between any pair of nodes, or in other words, it is the maximum geodesic distance in the network:

$$D = \max_{i,j} d(i, j)$$

In weighted networks, distance is defined as the inverse of edge weight ($d_{ij} = \frac{1}{w_{ij}}$), reflecting the intuition that stronger connections create shorter distances. Consequently, the diameter is measured in units of inverse weight rather than hop count.

While the average path length captures typical transmission efficiency, diameter identifies the most peripheral node pairs and highlights the longest propagation distance. The network exhibits a diameter $D = 1.187$, corresponding to the longest weighted shortest path between any two reinsurers. This metric, combined with the average path length, suggests a structure in which the core companies exhibit strong variance contribution relationships while peripheral nodes, although more statistically distant, remain connected through indirect pathways.

3.3.3 Assortativity

Assortativity, also known as homophily, serves as a macroscopic descriptor of network topology, quantifying the preference for nodes to attach to other nodes that are similar in some way. It measures the degree of correlation between the properties of a pair of connected vertices. It is calculated with respect to node degree or strength, assessing, in this case, whether well-connected hubs tend to associate with other hubs, or whether they instead connect to smaller peripheral nodes. Following Yuan, Yan, and Zhang (2021), for weighted directed networks, the assortativity coefficient is defined as the weighted Pearson correlation coefficient:

$$r_{\alpha,\beta} = \frac{\sum_{(i,j)\in E} w_{ij} \left(s_i^{(\alpha)} - \bar{s}_{\text{sou}}^{(\alpha)}\right) \left(s_j^{(\beta)} - \bar{s}_{\text{tar}}^{(\beta)}\right)}{W \sigma_{\text{sou}}^{(\alpha)} \sigma_{\text{tar}}^{(\beta)}} \tag{3.1}$$

where w_{ij} is the edge weight from node i to node j, $W = \sum w_{ij}$ is the total edge weight, $s_i^{(\alpha)}$ is the out-strength of source node i, $s_j^{(\beta)}$ is the in-strength of target node j, and $\bar{s}$ and σ denote the edge-weight means and standard deviations of the source (sou) and target (tar) edges.

For positive values of r, the network exhibits assortativity mixing, which means that strong nodes tend to connect to other strong nodes, forming a dense cohesive core. For negative values of r, the network exhibits disassortative mixing, in which hubs act as bridges. Using strength-based assortativity to account for edge weights, the network exhibits moderate disassortativity ($r = -0.116$). This indicates that high-strength variance contributors preferentially connect to low-strength peripheral nodes rather than forming a dense interconnected core among themselves. This structure has important implications for systemic risk: variance contributions from major contributors are distributed throughout the network rather than remaining concentrated within the core. This architecture carries two opposing implications: a hub-and-spoke structure absorbs the failure of a random peripheral firm with little systemic effect, yet the distress of a hub such as Munich Re, Swiss Re, or AXA is broadcast efficiently across the periphery rather than bottled up within the core.

3.4 Micro-Prudential Analysis: Node Centrality

These metrics are the core of network analysis and are the keys to assessing the role of each reinsurer's contribution to the system's aggregate forecast error variance. In this section, we try to identify firms that can be considered systemically important firms (Systemically important financial institution - SIFI) and assess the influence that each reinsurer has on the whole system.

3.4.1 Strength

Strength is the sum of the weights of edges attached to a node. Since we are considering a directed network, this measure splits into In-Strength and Out-Strength.

- Out-Strength: $S_{out}(i) = \sum_j \omega_{ji}$ is used to identify the total contribution of node i in explaining other nodes' forecast error variance
- In-Strength: $S_{in}(i) = \sum_j \omega_{ij}$ is used to identify the total forecast error variance of node i explained by other nodes

High Out-strength values indicate that a firm is a major variance contributor, meaning that it explains substantial portion of other reinsurers' forecast error variance, while high In-Strength values indicate that a firm is a major variance recipient, and its forecast error variance is largely explained by other companies' changes. For the network G_{sparse}, the largest variance contributor is the French reinsurer AXA XL (Ticker: CS.PA), with $S_{out} = 137.627$, while the largest variance recipient is the American reinsurer W.R. Berkley Corp (Ticker: WRB) with $S_{in} = 80.372$.

Table 3.5a and Table 3.5b list in descending order, the In-strength and Out-strength, of the nodes composing the network. The first 6 variance contributors are located in Europe, while three companies (one American and two Asian) have 0 out-strength (do not have outgoing edges), implying no contribution in explaining other reinsurers' variance, but acting purely as variance receivers.

Table 3.5: In- and Out-strength values for the 27 reinsurers.

(a) In-Strength: variance recipients.

Ticker	In	Continent
WRB	80.372	America
MKL	79.518	America
RGA	75.611	America
CB	74.801	Europe
MUV2.DE	74.167	Europe
EG	74.055	America
G.MI	71.607	Europe
ACGL	70.844	America
0966.HK	66.518	Asia
MAP.MC	66.356	Europe
8725.T	65.140	Asia
CS.PA	63.318	Europe
FFH.TO	62.436	America
AXS	60.238	America
003690.KS	58.671	Asia
RNR	53.977	America
SCR.PA	52.339	Europe
8630.T	51.667	Asia
SREN.SW	51.112	Europe
1508.HK	49.490	Asia
HNR1.DE	48.792	Europe
SPNT	46.905	America
8766.T	44.445	Asia
BRK.A	44.425	America
6809.T	37.544	Asia
HSX.L	5.046	America
QBE.AX	2.566	Australia

(b) Out-Strength: variance contributors.

Ticker	Out	Continent
CS.PA	137.627	Europe
HNR1.DE	128.843	Europe
G.MI	126.444	Europe
MAP.MC	124.770	Europe
SREN.SW	118.575	Europe
MUV2.DE	117.616	Europe
CB	78.736	Europe
8766.T	69.997	Asia
8630.T	64.799	Asia
EG	62.935	America
MKL	57.597	America
ACGL	55.395	America
SCR.PA	54.745	Europe
BRK.A	53.073	America
AXS	52.189	America
8725.T	52.071	Asia
RGA	51.815	America
RNR	35.534	America
WRB	31.885	America
FFH.TO	24.749	America
QBE.AX	13.329	Australia
6809.T	10.247	Asia
0966.HK	6.612	Asia
SPNT	2.380	America
003690.KS	0.000	Asia
1508.HK	0.000	Asia
HSX.L	0.000	America

To better assess the recipients and contributors of variance of the network's nodes, the difference between the outgoing strength and the receiving strength, the net-strength, was used:

$$S_{net} = S_{out} - S_{in}$$

Companies that have positive S_{net} values are considered net variance distributors, while companies that have negative values are considered net variance recipients. Figure 3.3 provides a graphical representation of the companies and their net weight S_{net}. Again, variance contributors tend to be concentrated in Europe (purple bars), with 66% of companies with a positive net balance being based in Europe, and the remaining being distributed across Asian, America, and Australian companies (respectively green, red, and blue bars). The variance receiving companies ($S_{net} < 0$), are composed predominantly by Asian and American companies. This pattern maps cleanly onto the industrial organization of the sector: global underwriting capacity is concentrated in few European groups whose pricing and capital decisions set the market tone at the major renewals. The market therefore re-prices European firms first, while American companies sit further along the risk transfer chain and therefore receive the signal.

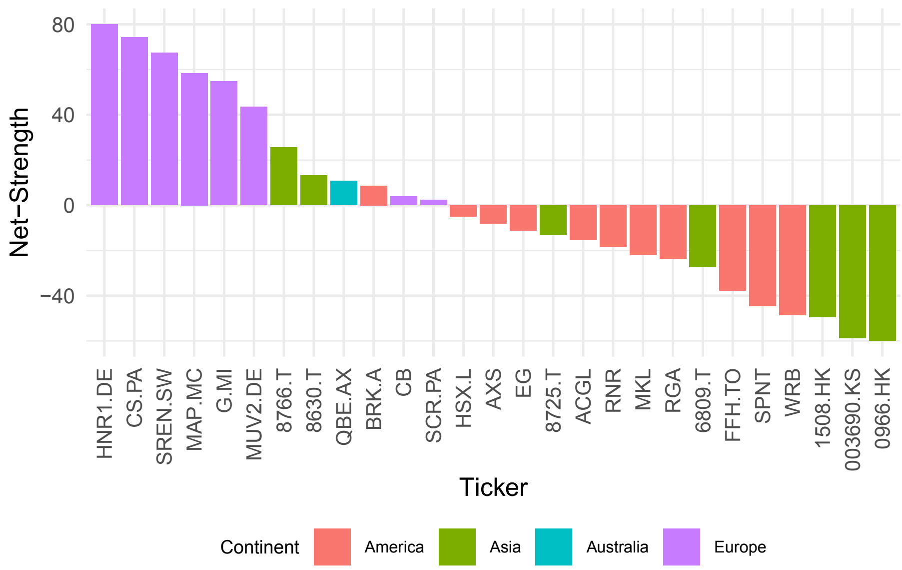

Figure 3.3: Net-strength balance for each ticker. Positive values indicate that the company is a net variance contributor, while negative values indicate that the firm's forecast error variance is more explained by others than it contributes to explaining others.

3.4.2 Betweenness Centrality

Betweenness centrality is the fraction of all shortest paths in the network that pass through a specific node, measuring the extent to which the node lies in the pathways connecting other nodes' forecast error variance relationships.

$$C_B(v) = \sum_{i \neq v \neq j} \frac{\sigma_{ij}(v)}{\sigma_{ij}}$$

where σ_{ij} is the total number of shortest paths from node i to node j, and $\sigma_{ij}(v)$ is the number of these paths that pass through v. For weighted networks, the shortest paths σ_{ij} are computed using the edge distances defined as the inverse of weights ($d_{ij} = 1/w_{ij}$), reflecting the interpretation that stronger variance contributions creates shorter statistical distances.

This metric is used to identify nodes that function as bridges of the network, connecting two or more different clusters. High values of betweenness indicate that the firm occupies a structurally important position in the network's connectedness topology, serving as a bridge between otherwise less-connected regions. Such firms lie on many of the shortest variance-contribution pathways between other node pairs. From a network perspective, this position suggests that these firm's unexpected movements may be statistically relevant to explaining variance patterns across different regions of the network. The firm with the highest value is Assicurazioni Generali SpA (G.MI), with a value of $C_B = 51.043$, followed by Markel Corp (MKL) and Mapfre (MAP.MC), with respective values of 39.338 and 31.804. G.MI serves as the strongest bridge for variance patterns across the system, and the higher values for European reinsurers suggest a strong importance of the European market, serving as variance contributors themselves and as connectors for the variance recipients nodes, with Markel Corp. being the only non European company appearing with high betweenness value. The economic reading is that the maximum betweenness value is small in absolute terms, so there is no structural company whose removal would fragment the network. Contagion has many paths and is diffuse rather than reliant on a single intermediary, in contrast with the banking sector, where a major interbank lender or a central counterparty can be "too-interconnected-to-fail" as a unique gateway.

3.4.3 Eigenvector Centrality

Eigenvector centrality represents a measure of network embeddedness where a node is considered highly embedded if it is connected to other nodes that are themselves considered nested. This approach captures the inherent structure of the network, considering the value of a connection from the prominence of the counterpart involved.

Mathematically, the concept is formalized through the spectral properties of the adjacency matrix A. Let λ_1 denote the largest eigenvalue of A and let $\mathbf{x}$ be its corresponding eigenvector satisfying $A\mathbf{x} = \lambda_1\mathbf{x}$. The eigenvector centrality of node i is then defined as the i-th component of $\mathbf{x}$, denoted x_i. Decomposing this to individual node level, the centrality score satisfies:

$$\lambda_1 x_i = \sum_j A_{ij} x_j$$

or equivalently:

$$x_i = \frac{1}{\lambda} \sum_j A_{ij} x_j$$

In the context of financial stability and systemic risk analysis, eigenvector centrality provides a metric to identify institutions that are "too connected to fail". A high eigenvector centrality score indicates that a firm is enclosed in the network's core, with direct linkages to the most important nodes. Consequently, firms with high eigenvector centrality are nested in pathways connecting the network's most connected nodes, meaning that their idiosyncratic shocks are statistically linked to the movements of other embedded firms. Conversely, a firm that has low eigenvector score interacts mostly with the periphery of the network, interacting mainly with less embedded entities, and with a limited capability of explaining forecast error variance to the core nodes of the network.

Table 3.6: Eigenvector values of the network nodes.

Ticker	Eigenvector	Continent
CS.PA	0.380	Europe
MAP.MC	0.356	Europe
G.MI	0.355	Europe
HNR1.DE	0.352	Europe
MUV2.DE	0.344	Europe
SREN.SW	0.329	Europe
8766.T	0.205	Asia
8630.T	0.202	Asia
SCR.PA	0.183	Europe
8725.T	0.175	Asia

Table 3.6 provides the 10 highest values of the eigenvector centrality. AXA XL (CS.PA) was the company with the highest score, with similar values (> 0.3) for 5 other companies based in Europe, supporting again the idea that European market explains the aggregate volatility movements of the system.

It is interesting to note the high correlation ($\rho = 0.926$, $p < 0.001$) of eigenvector values with the Net-strength: this high positive correlation coefficient highlights the fact that the most influential hubs are also the nodes that contribute the most to the market's aggregate volatility, suggesting a concentrated core of market risk. In markets with clear core-periphery topology, small peripheral firms connected exclusively to central hubs would exhibit high eigenvector centrality despite low strength. The absence of this pattern suggests that the global reinsurance market operates as a relatively homogeneous, densely interconnected system where the largest volatility contributors are also the most systemically embedded institutions. Figure 3.4 provides a graphical representation of the correlation between the Eigenvector values and the net-strength value of the nodes.

In dense neutrally assortative networks ($r \approx 0$), high eigenvector-strength correlation is typical because uniform mixing makes strength the primary determinant of centrality. In this network, the high correlation occurs despite disassortativity ($r = -0.116$), suggesting that the hub-spoke structure amplifies rather than diminishes the importance of variance contribution magnitude. European variance contributors are embedded in the core both because of their high variance contribution magnitude and their connections to peripheral recipients, creating a reinforcing feedback where strength determines structural position.

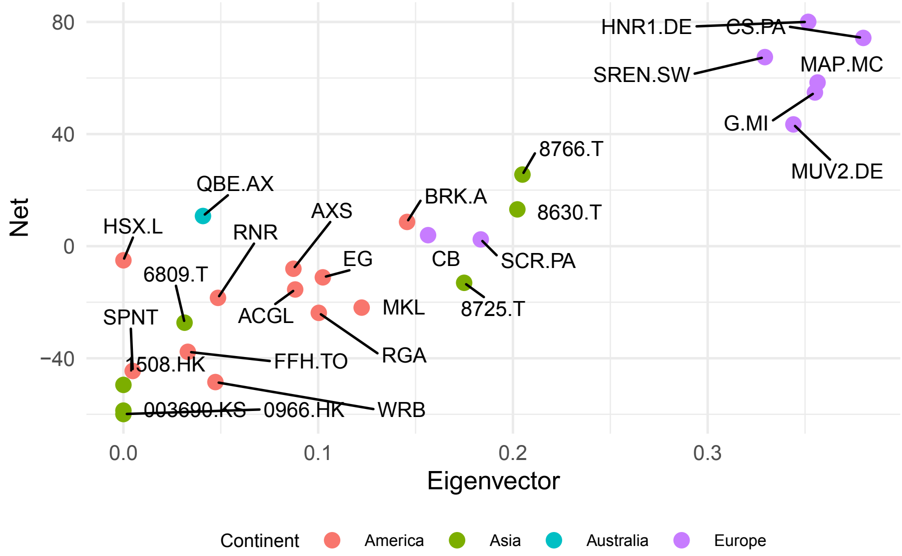

Figure 3.4: Correlation between the net-strength of nodes and the eigenvector measure.

3.4.3.1 Size and Connectedness

One objection is that these centrality rankings might merely re-describe firm size, since the largest reinsurers could mechanically dominate any connectedness measure. The data shows that size and centrality are positively and significantly related, but not equivalent. In fact, the Spearman rank correlation between out-strength and market capitalization is $\rho = 0.606$, a moderate association indicating that larger firms do tend to be more central, yet one that leaves a substantial part of the centrality ordering unexplained by size. Table 3.7 reports the distinctions between the two metrics. For Munich Re (MUV2.DE), Swiss Re (SREN.SW), and AXA XL (CS.PA), size and connectedness coincide, as all rank among the largest firms by market value and among the most central transmitters. Other firms, however, are markedly more central than their size would suggest. Mapfre (MAP.MC), for example, ranks fourth in the out-strength ranking, but only nineteenth for market capitalization. Further, Hannover Re (HNR1.DE) is only a mid-sized firm by market value, but second overall by variance contribution to the system. Conversely, the single largest firm in the sample, Berkshire Hathaway (BRK.A) ranks fourteenth in the transmission rank. Connectedness therefore captures a firm's position in the propagation structure of the market, distinct from its absolute scale or written-premium volume, although being correlated with it. The implication for macro-prudential surveillance is that a screen based on size or market share alone, like the one embedded in the IAIS size indicator, would over-weight large but possibly peripheral firms and under-weight smaller but potentially highly embedded transmitters.

Table 3.7: Network centrality vs firm size. First 10 firms according to the out-strength rank.

	Ticker	OutRank	EigenRank	MktCap	MktRank	GWP
12	CS.PA	1	1	74.46	4	8.41
16	HNR1.DE	2	4	26.60	14	22.68
15	G.MI	3	3	43.37	8	81.84
18	MAP.MC	4	2	10.77	19	6.72
26	SREN.SW	5	6	54.45	6	44.86
20	MUV2.DE	6	5	62.00	5	34.42
11	CB	7	11	112.18	2	62.00
7	8766.T	8	7	76.77	3	36.30
5	8630.T	9	8	26.68	13	26.83
13	EG	10	14	14.72	16	12.94

3.4.4 PageRank

PageRank is used to identify node importance based on each node's connections. It was originally designed for web search (the algorithm that laid the foundations of Google's web search success), and is best interpreted as the "flow" of information that travels in the network (Brin and Page 1998). The PageRank algorithm differently accounts for the weighting of out-degree nodes. A node's influence is divided equally among all its outbound links, therefore, a prestigious node that links to only a few other nodes confers higher importance to those neighbors than a prestigious node that connects to hundreds of nodes. Another distinction in the PageRank algorithm is the addition of a damping factor d, usually set to 0.85, that transforms the adjacency matrix A into a stochastic, irreducible, and aperiodic Google Matrix. These characteristics allow to effectively model a "random shock" moving through the graph, ensuring that the entire network is taken into account.
Mathematically it is represented as

$$PR(i) = \frac{1-d}{N} + d \sum_{j \in M(i)} \frac{PR(j)}{L(j)}$$

where $M(i)$ are nodes pointing to i, and $L(j)$ is the out-degree of j.

While the eigenvector centrality identifies which nodes are connected to the core, the PageRank measure identifies which nodes receive the largest aggregate contributions from the system. The highest PageRank was observed for Markel Corp. (MKL), with a value of $PR = 0.072$, followed by China Reinsurance Corp (1508.HK) and the US Everest Group (EG).

These companies are reported in negative net values (recall Figure 3.3), and unsurprisingly also rank among the highest PageRank values, highlighting their tendency to receive variance contribution without proportionally contributing to others' variance explanation. These firms act as "terminal nodes" in the variance contribution topology. Eigenvector-central firms are embedded in the network's core structure, while PageRank-central firms represent primary recipients of aggregate variance contribution.

3.4.5 Hubs and Authorities

Hubs and Authorities concept was introduced by Jon Kleinberg through the Hyperlink-Induced Topic Search (HITS) algorithm and serves as another measure to identify central nodes and the pathways of variance explanation relationships. Hubs and Authorities distinguishes between two roles within the variance relationship:

- Hubs: nodes that point to good authorities, they act as the main variance contributors
- Authorities: nodes that are pointed to by good hubs. They are considered main receivers from the hubs.

This creates a cycle of mutual reinforcement: a hub is valuable if it points to valuable authorities, while an authority is valuable if it is pointed to by valuable hubs (Kleinberg 1999). Mathematically, this involves calculating the principal eigenvectors of both the adjacency matrix A for the authorities and its transpose A^T for the hubs.

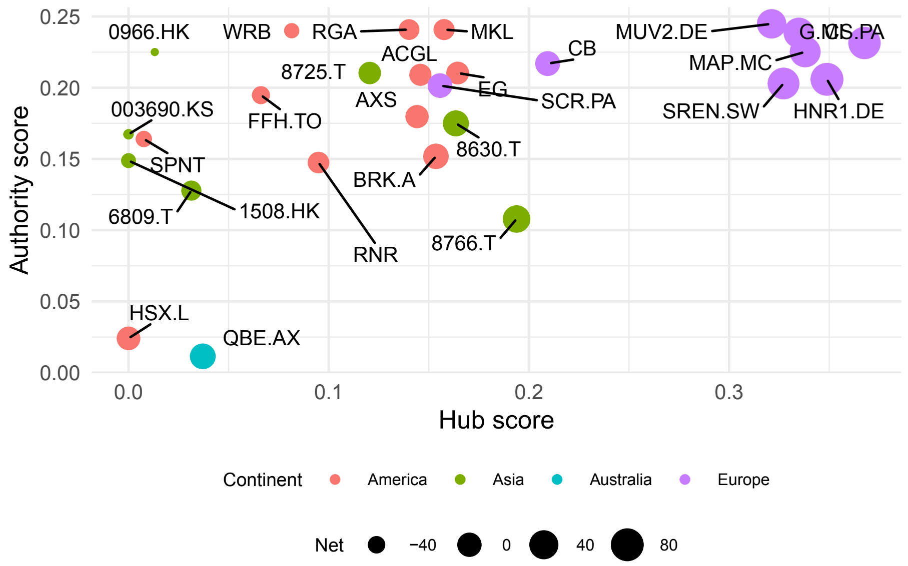

Figure 3.5: Scatterplot of the Hub scores (x-axis) and the Authority score (y-axis).

The elements in quadrant 1 of Figure 3.5 are nodes that present high hub scores and high authority score, while the second quadrant is categorized by nodes that have high authority scores but low hub scores. As is possible to see, the upper right corner is constituted by only European companies, once again highlighting the major variance explanation that these companies have on the whole system. Figure 3.6 provides a representation of the network displaying the nodes according to their hubs or authorities categorization. Such distinction was made by subtracting the two scores obtained: for positive values, the node was classified as *Hubs*, meaning that these nodes act as intermediaries in variance contribution pathways, while negative values were classified as *Authorities* nodes, meaning that they are usually the primary recipients of explained variance.

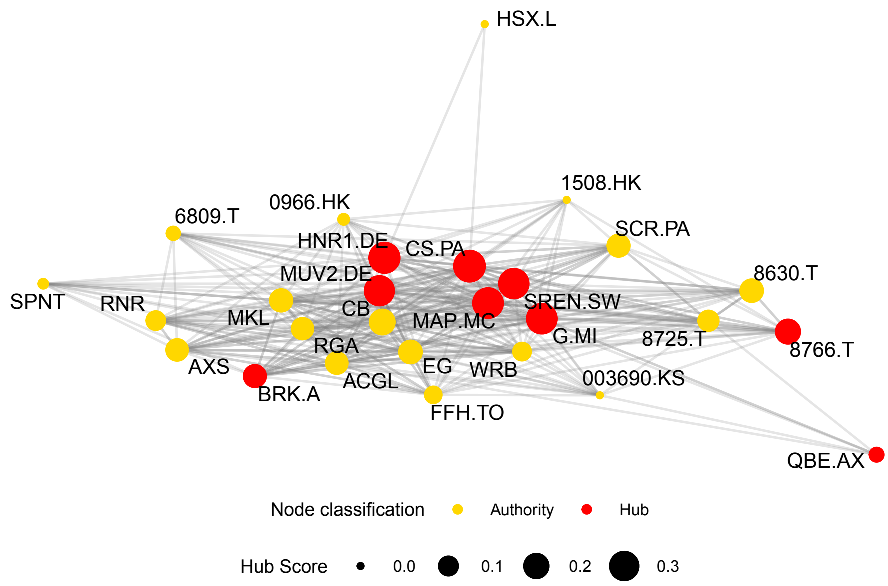

Figure 3.6: Network plot classifying hubs and authorities.

3.5 Clustering and Community Detection

3.5.1 Clustering Coefficient - Transitivity

Transitivity is a structural metric that quantifies the degree to which nodes in a network tend to cluster together, measuring the prevalence of triadic closures. It specifically measures the probability that two neighbors of a given node are themselves connected, assessing the local cohesiveness of the network comparing the number of closed triangles to the number of connected triplets. Mathematically it is defined as

$$C = \frac{\text{number of triangles connected to } i}{\text{number of triples centered on } i}$$

or more precisely

$$C_i^{\omega} = \frac{1}{S_i(k_i - 1)} \sum_{j,h} \frac{w_{ij} + w_{ih}}{2} a_{ij} a_{ih} a_{jh}$$

where S_i is the strength of vertex i, a are elements of the adjacency matrix A, k_i is the vertex degree, and w_{ij} are the edge weights (Csárdi et al. 2025). The transitivity value for the network of reinsurers is relatively high at 0.78, implying that the variance contributions flow through the dense local connection patterns, with multiple pathways linking neighboring nodes.

3.5.2 Community Detection

Community detection offers a broader view of the concept that transitivity just introduced. It seeks to aggregate the local clusters identified into large, coherent clusters, known as communities. A community is typically defined as a subdivision of the network where nodes are more densely connected internally than the rest of the network.

3.5.2.1 Modularity

Modularity score (Newman 2006), serves as the measure to identify the strength of the division of a network into modules, comparing the actual links within a community to the expected density if links were distributed randomly, but maintaining the same degree distribution.

$$Q = \frac{1}{2m} \sum_{ij} \left(A_{ij} - \frac{k_i k_j}{2m} \right) \delta(c_i, c_j) \tag{3.2}$$

where $\delta(c_i, c_j) = 1$ if nodes i and j are in the same community.

Communities are detected by maximizing the modularity score Q, and multiple algorithms can be used. In this analysis, we adopted the Louvain algorithm (Blondel et al. 2008), due to its high speed in identifying communities, even for large networks, and its efficiency and accuracy in achieving high modularity scores and correctly determining communities.

To maximize modularity, Louvain's algorithm functions through two phases:

1. In the first phase, the algorithm positions each node in its own community. It then examines a node and calculates the gain in modularity (ΔQ) that result from moving that node to the community of its neighbors. The node is therefore moved to the community that yields the highest modularity increase. This process is repeated for all nodes
2. The second phase consists of the aggregation of communities. Once a local maxima is reached, the network is collapsed, and each community found in the first phase of the iteration is considered as single nodes. The edges between the new single nodes are the sum of the weights of the edges between the constituent nodes in the original graph, while internal links within the community become self-loops in the new network

These phases are iterated for all nodes composing the collapsed network until modularity no longer increases, obtaining the optimized value of Q.

Prior to maximizing modularity with the Louvain algorithm, the directed adjacency matrix was transformed into an undirected graph. Specifically, a single undirected edge was created for each pair of vertices connected by at least one directed link. This symmetrization step ensures compatibility with Equation 3.2, which is formulated for undirected networks.

The Louvain algorithm therefore identified 3 different communities (reported in Table 3.8). Algorithmic stability was verified by repeating the procedure across 10 independent runs with distinct random seeds; all runs returned the same community identification, confirming that the identified partition is unique and not randomly initialized. The modularity score reached a maximum value of $Q = 0.162$, which indicates relatively weak community structure. For reference, modularity scores above 0.3 are typically considered meaningful indicators of community division (Newman 2006), while values between 0.1 and 0.3 suggest only modest clustering tendencies.

The primary finding is therefore that the global reinsurance market exhibits weak community structure. This result is consistent with the low network disassortativity ($r = -0.116$) and short average path length ($L = 0.394$) identified before, suggesting that the system behaves as a tightly integrated rather than isolated regional markets. A secondary descriptive observation shows that the three detected communities show *partial* alignment with broad geographic groupings: community *1* includes mainly American companies, community *2* is composed predominantly of European based companies, and community *3*, includes three Japanese companies and the only Australian reinsurer. Given the low $Q = 0.162$, this geographic pattern reflects a weak tendency for firms in similar time zones and regulatory environments to exhibit marginally higher co-movement.

Table 3.8: Membership of the community for each ticker.

	Community	Continent
003690.KS	1	Asia
ACGL	1	America
AXS	1	America
BRK.A	1	America
CB	1	Europe
EG	1	America
FFH.TO	1	America
MKL	1	America
RGA	1	America
RNR	1	America
WRB	1	America
0966.HK	2	Asia
1508.HK	2	Asia
6809.T	2	Asia
CS.PA	2	Europe
G.MI	2	Europe
HNR1.DE	2	Europe
HSX.L	2	America
MAP.MC	2	Europe
MUV2.DE	2	Europe
SCR.PA	2	Europe
SPNT	2	America
SREN.SW	2	Europe
8630.T	3	Asia
8725.T	3	Asia
8766.T	3	Asia
QBE.AX	3	Australia

To better assess this distinction, however, we introduce the measure of conductance that allows to quantify how well-separated a community is from the rest of the network. Lower level (< 0.2) of conductance indicates that the

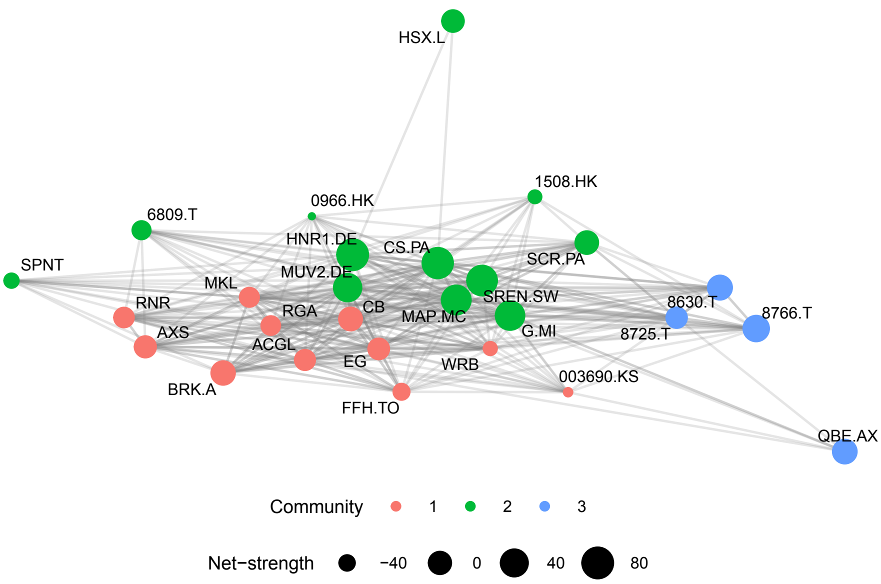

Figure 3.7: Network plot highlighting the different communities detected using the Louvain algorithm.

community is better-defined with fewer external connections relative to internal ones, suggesting stronger cohesion within the community. Mathematically, for a weighted graph, it is defined as:

$$\Phi(C) = \frac{cut(C, \bar{C})}{\min(vol(C), vol(\bar{C}))} \tag{3.3}$$

where $cut(C, \bar{C})$ is the sum of the weights of the edges that have one endpoint in the community C and the other endpoint outside the community, $vol(C)$ is the volume of the community C, corresponding to the sum of the weights of the edges incident to nodes in C, and $vol(\bar{C})$ is the volume of the rest of the network, excluding community C.

From Table 3.9 we can see the conductance level for each community. Since communities overlap (Community 1 includes several European firms and a South Korean and Community 2 includes a Japanese firm), the conductance scores are particularly low. Community 2 exhibits the lowest conductance, suggesting marginally stronger internal cohesion. However, the values are similar across all three communities and consistent with the globally integrated market structure identified by the low modularity score.

Table 3.9: Conductance level for each community.

Community	Conductance
Community 1	0.432
Community 2	0.411
Community 3	0.527

3.6 Summary and Key Findings

This chapter applied network analysis to the BVAR-GFEVD framework to identify potential systemically important reinsurers and define the network topology of the global reinsurance market. Some key findings emerged:

- **Network Structure**: The global reinsurance network exhibits a short average path length ($L = 0.394$) and slight disassortative behavior ($r = -0.116$), with a high total connectedness value ($TC = 74.278$), indicating a tightly interconnected system where European companies tend to form hubs but still connect to Asian and American nodes, rather than forming isolated cores. Table 3.10 reports the first 5 companies with the highest metrics analyzed;
- **Systemic Importance**: European companies appeared to be the most significant variance contributors of the system, with AXA XL, Swiss Re, and Munich Re ranking highest in the centrality measures, while Asian companies emerged as variance recipients;
- **Community Detection**: The Louvain algorithm yields $Q = 0.162$, indicating weak community structure and confirming that the global reinsurance market operates as an integrated system rather than a collection of regionally segmented markets. Three communities with partial geographic overlap are identified, but given the low modularity and the similar conductance values, this geographic pattern is a secondary descriptive observation rather than a structural finding;
- **Limitations**: Some limitations of this workflow need to be addressed. First, the use of stock volatility poses a strong bias over short-term sentiment and does not serve as an insightful indicator for financial distress. A balance-sheet based analysis could have offered deeper and more accurate identification of potentially systemic reinsurers, at the expense of the number of observations available, due to the limited availability of balance sheet data. Chapter 5 tries to address this issue by running market-based systemic risk measures. Second, the static analysis captures a stationary system which does not reflect the moments of crisis and the dynamic evolution of the system itself. In the following chapter we are going to address this issue, performing a dynamic analysis over the dataset and review the network topology and systemically important reinsurers during specific events. A single full-sample estimate yields only an average network and cannot reveal that connectedness spikes in crises as idiosyncratic risk gives way to common factors; the static picture therefore understates tail-state systemic risk and cannot identify which firms become systemic specifically during a crisis. The next chapter, through the dynamic, rolling-window analysis, tries to address this limitation.

Table 3.10: Top-ranked firm for each centrality metric.

Rank	Out-Strength	In-Strength	Net	Eigenvector	Betweenness
1	CS.PA	WRB	HNR1.DE	CS.PA	G.MI
2	HNR1.DE	MKL	CS.PA	MAP.MC	MKL
3	G.MI	RGA	SREN.SW	G.MI	MAP.MC
4	MAP.MC	CB	MAP.MC	HNR1.DE	MUV2.DE
5	SREN.SW	MUV2.DE	G.MI	MUV2.DE	CS.PA

Chapter 4

Dynamic Network Analysis

4.1 Introduction

In Chapter 3 we provided a static analysis of the network obtained applying a BVAR model over the entire dataset of the log-volatility of reinsurance companies. Although these data produce a solid base for statistical validity, static network analysis has limitations with regard to the dynamic developments of the relations and components of the network. To better assess such elements, we performed a dynamic network analysis (DNA), that allows us to observe changes over the entire time period, making the identification of systemically important reinsurers and the dynamics of the whole system more insightful and robust. This chapter displays the methodology used to perform the DNA, the evolution of the network topology and the firm level dynamics over the 10 years period, with a later focus on the effect of significant events over the system's behavior.

The structural stability of the dynamic analysis was tested performing structural break test on the total connectedness index time series. The tests performed are the Bai-Perron and the ordinary least square cumulative sum (OLS-CUSUM), and their results are reported in Section C.2.

4.2 Methodology

The data used to perform the subsequent analysis is the same as presented in Chapter 2, consisting of 27 reinsurance companies with log volatility data for the period starting at 2015-11-27 and ending at 2025-05-30, for a total of 1993 observations per company. This data set was divided into 349 windows, each containing 250 observations. Each window advances by a step of 5 observation between the previous one. The window size corresponds roughly to a full trading year (in our case, it amounts to more than one year, due to the constraint of gathering companies' prices from different markets), while the step size corresponds to a weekly delay. The choice of this window length fixes the economic memory of each estimate, long enough to yield a stable estimate yet short enough to register a regime change within a smaller period. This design captures sustained shifts in interdependence rather than transient single-week contagion.

The fragmented data set was subjected to the same BVAR model shown in Chapter 2 to obtain the GFEVD matrix to estimate the network. However, due to technical limitations and computational requirements, the burn-in phase and the total draws used to estimate the BVAR model for each window have been reduced, respectively, to $3,000$ and $9,000$ draws, compared to the $5,000$ burn-in and $15,000$ total draws used to estimate the static network. The reduced number of draws shortens the estimation time for the whole process, but remain sufficient to ensure reliable point estimates. In order to ensure the stability of the model, the convergence of the MCMC chains has been tested every 20 windows using the Gelman-Rubin diagnostic. Across the 349 windows, the Potential Scale Reduction Factor (PSRF) for both hyperparameters and the marginal likelihood remained below 1.02, significantly better than the standard 1.10 threshold, indicating robust convergence of the MCMC chains, letting the regime switching dynamics documented below to be attributed to genuine features of the market, rather than to instability of the sampler.

As a result, each matrix represents the edges between reinsurers and the fully weighted connected graph for each window. The total connectedness value C^H was computed along with the "From" and "To" values. After analyzing the distribution of the fully connected graphs created using the connectedness matrix C, the 50^{th} percentile threshold was applied to allow comparison with the analysis presented in Chapter 3. In the next section, we present the analysis of the topology evolution and the firm level dynamics during the period analyzed.

4.3 Network Evolution

4.3.1 Total Connectedness Evolution

Figure 4.1 shows the evolution of Total Connectedness Index (TCI) over time. We use this metric as a primary identifier for significant changes and to assess the overall connectedness of the system. Compared to the static value obtained in the previous chapter (74.28), the average value over the whole 349 windows analyzed is 68.2, with a maximum value of 88.75 and a minimum of 56. To better assess the distribution of the connectedness values obtained, the distribution was divided into three regimes:

- Low regime for values below the 25^{th} percentile of the distribution
- Normal regime, for values between the 25^{th} and 75^{th} percentiles
- High regime, for values greater than the 75^{th} percentile

The higher standard deviation of the high regime values, as reported in Table 4.1, suggests that the presence of outliers typically originates for values higher than the normal mean. The TCI distribution is shown in Figure 4.2. Clearly, the distribution is bimodal, showing a marked regime change and two peak values in the normal and high regimes. The connectedness value does not revert towards its mean value, but bimodality suggests that the system alternates between two distinct regimes: one standard equilibrium and one with higher connectivity. Following this classification, the result of the static analysis would be categorized as a high connected network, as its TC value is above the threshold quantile. This indicates that the static analysis might have been inflated by periods with higher connectivity, and that the reinsurance market

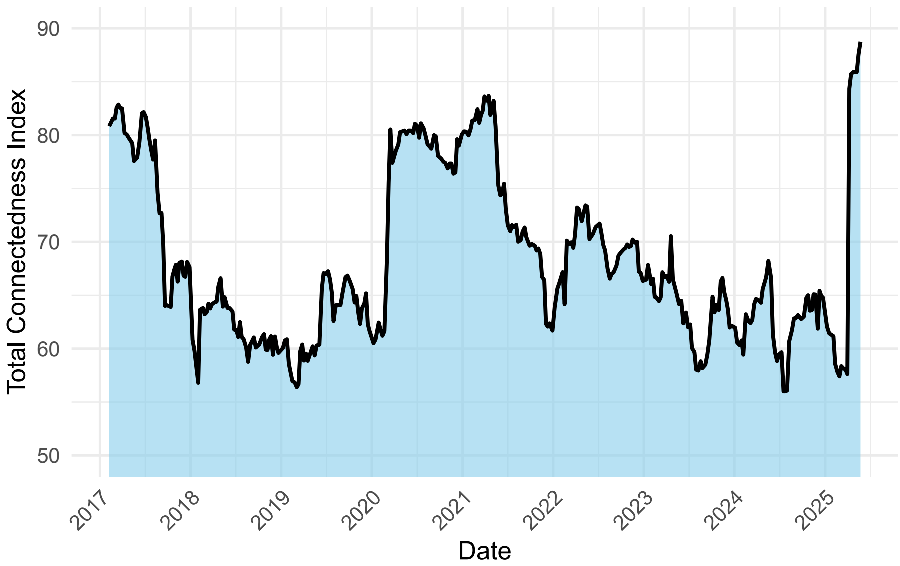

Figure 4.1: Total connectedness value for every rolling window computed over the whole period.

Table 4.1: Regime classification of the TCI values recorded for the 349 windows estimated.

Regime	Windows	Mean TCI	SD TCI	Pct
Normal	175	66.50	2.91	50.14
High	87	80.04	2.91	24.93
Low	87	59.81	1.59	24.93

might not possess a single equilibrium level of integration, but rather two states between which it switches. This is interesting for a supervisory standpoint, reframing the question from how connected the system is at a given moment to which regime the system is currently in and whether it is transitioning from one state to another.

The pairwise edge distribution of the static network has been compared with the distributions obtained from each computed network. The results showed similar values for the mean of the distributions, while the standard deviation of the dynamically computed networks appeared to be slightly higher. The evolution of the mean and the standard deviation for the networks are shown in Figure 4.3. The highest standard deviation are in conjunction with lower TCI values, indicating that a less connected network results in a wider range of pairwise connection, as shown by the negative Pearson's correlation ($\rho = -0.41$) between the two metrics.

Regarding the strongest pairwise connections between reinsurers, a similar pattern between static and dynamic analysis was found. The Japanese reinsurers tend to maintain stronger effects and influence within their domestic market, and above the average value of the system. The persistence of this block in more than half of the windows identifies it as a structural rather than episodic

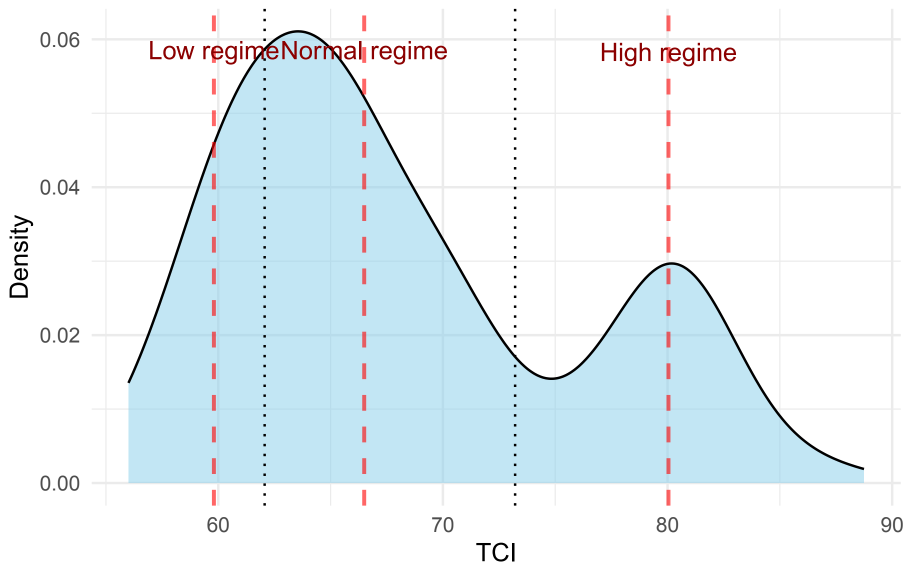

Figure 4.2: Density plot of the TCI values, separated by regime. The red dashed line represents the mean of the regime class.

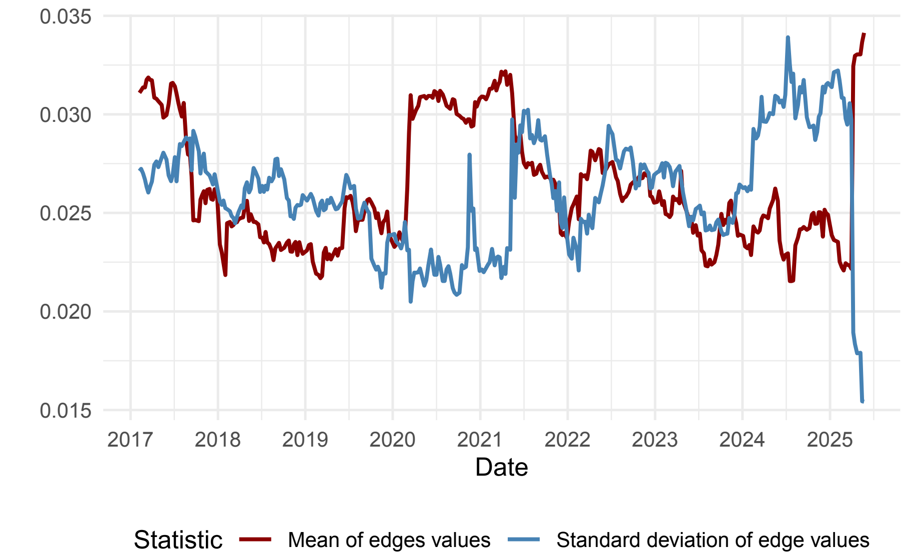

Figure 4.3: Values of the mean and standard deviation for each of the 349 windows computed.

Table 4.2: Frequency and average values of the strongest 15 pairwise connections between reinsurers collected from all the 349 windows.

From	To	Frequency	Freq. pct	Avg. Strength	Max Strength
8725.T	8766.T	264	75.64	16.09	25.51
8766.T	8630.T	262	75.07	12.91	20.98
8725.T	8630.T	235	67.34	15.67	21.79
8766.T	8725.T	199	57.02	15.20	22.00
8630.T	8725.T	179	51.29	13.03	19.11
8630.T	8766.T	173	49.57	15.58	25.52
MUV2.DE	HNR1.DE	140	40.11	10.72	15.98
WRB	RNR	114	32.66	12.38	18.84
WRB	CB	86	24.64	10.23	14.52
MUV2.DE	SREN.SW	85	24.36	9.38	11.66
SREN.SW	CS.PA	85	24.36	9.67	11.62
RNR	WRB	81	23.21	17.32	26.58
EG	RNR	73	20.92	12.13	14.68
HNR1.DE	G.MI	73	20.92	10.55	13.40
SREN.SW	MUV2.DE	71	20.34	9.96	12.61

feature, showing a tightly interconnected cluster that co-moves internally without exporting substantial variance to the global system. Table 4.2 reports the pairwise connections that appeared more frequently in the top 15 connections of the system.

4.3.2 Network Topology Evolution

The static analysis in Chapter 3 revealed a disassortative, globally integrated network. However, one might ask whether this behavior is maintained throughout stable and crisis periods, or if different scenarios reshape the network's architecture.

The 50^{th} percentile threshold is applied independently within each rolling window, producing a constant network density of 0.50 by construction. Changes in topology metrics thus do not reflect network densification, but rather the redistribution of edge weights among the retained connections, and the changes in which specific edges exceed the threshold between windows.

The results of our analysis show that during periods of elevated volatility and high connectedness regime, the reinsurance network contracts significantly. The average path length drops from its time series average of 0.39 to a minimum of 0.29 during mid 2021 high TCI period. This means that during crisis periods, the network tends to have shorter and closer connections, so variance spreads more rapidly, accelerating contagion when its speed is most dangerous.

The diameter instead shows weaker correlation with the TCI ($\rho = -0.368$), meaning that TCI changes primarily affect the average network distance, without proportionally shrinking the longest path. During high regime periods, firms get closer, while the most peripheral ones remain relatively distant, being less affected by the higher variance explanation of core firms. As Figure 4.4 shows, the diameter also exhibits higher volatility compared to the APL, reflecting less independence from the TCI regime status.

Notably, the topology response to TCI changes are not exactly symmetrical. An asymmetric regression of APL changes on positive and negative TCI increments yields coefficients of -0.003 ($p < 0.001$) for TCI increases and -0.007 ($p < 0.001$) for TCI decreases. The ratio of 2.33 between the two coefficients shows that the average path length reacts roughly $2.2\times$ more strongly to declines in connectedness that to increases, meaning that the network dissolves more rapidly when stress decreases than it compacts when stress increases. The relationship between connectedness and topological compactness is therefore nonlinear and asymmetric, with the network loosening abruptly once a stress episodes passes.

The rolling assortativity provides further insights. While the static analysis yielded a disassortative coefficient ($r = -0.116$), the rolling window average is slightly positive ($\hat{r} = 0.078$). The positive correlation with the TCI ($\rho = 0.326$) suggests that during high connectedness periods, the network shifts toward assortativity, with hub firms being more connected to other hubs rather than to peripheral nodes. This structural shift means that the hub-spoke topology is not permanent but tends to emerge during lower connectedness regimes. As connectedness rises, the large hubs couple directly to one another rather than feeding peripheral recipients, so the systemically important firms cease to diversify against each other and begin to move as a single block. The hub-and-spoke resilience documented in Chapter 3 is thus a normal times property that dissolves under stress, when the network adopts its most dangerous systemic configuration.

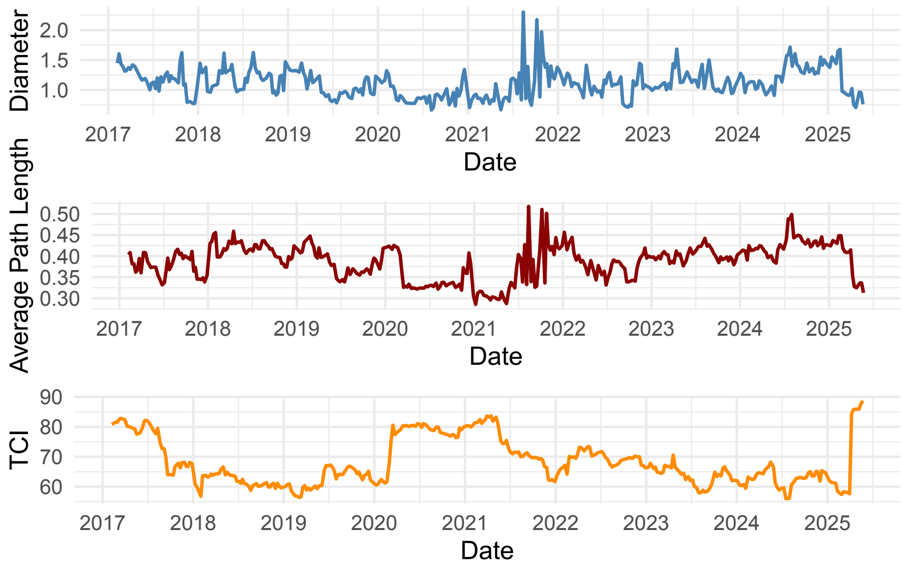

Figure 4.4: Values of the Diameter and of the Average Path Length (APL) in comparison with the TCI. The Diameter shows greater standard deviation, while the APL clearly show opposite responses to the TCI changes.

4.3.3 Network Stability and Community Evolution

We tested network stability and community evolution performing the Jaccard stability analysis for consecutive TC values, and community evolution analyzing the modularity measure for the various rolling windows.

4.3.3.1 Jaccard Stability

To quantify the similarity of the connections between consecutive networks we used the Jaccard similarity measure. Mathematically, the index is defined as the ratio between the intersection size and the union size, with values ranging from 0 for sets with no elements in common, and 1 for identical sets.

$$\text{Jaccard} = \frac{A \cap B}{A \cup B}$$

The observed mean Jaccard similarity score is 0.88 ($\sigma = 0.06$), with a right skewed distribution. While this high baseline suggests strong continuity, it must be interpreted in the context of the rolling window methodology applied: with a step size of 5 and a window of 250 observations, consecutive windows share 98% of the underlying data.

Consequently, the various "dips" in the Jaccard similarity observed in Figure 4.5 are highly significant. These sharp declines often coincide with rapid spikes or falls in the TCI, signaling periods of topological changes. The weak positive correlation between the Jaccard stability and the TCI ($\rho = 0.178$, $p < 0.001$) suggests that the level of connectedness does not intrinsically determine stability, but rather that stability is compromised during the transition between regimes. This supports the hypothesis that while a "stable core" of reinsurer relations exists during the normal market condition, systemic shocks force a rapid realignment of reinsurers' roles. The transition windows are therefore the moments of maximum structural fragility, since it is during these episodes that the identity of transmitters and receivers is in flux.

4.3.3.2 Modularity

To assess the presence of cluster structures, we computed the modularity evolution and the number of communities detected for each rolling window. The clustering method used to estimate the modularity and the communities was the Louvain algorithm (Blondel et al. 2008). The average modularity score was 0.179 ($\sigma = 0.039$), with values ranging from a minimum of 0.055 to a maximum of 0.265. The communities detected fluctuated between 2 and 5, with 96.5% of the observations having 3 or 4 clusters. The networks where 5 communities were observed result in lower interconnectedness values, as a less clustered core structure identifies more nodes that could be treated as communities.

However, the low absolute values of modularity (remaining below the 0.3 threshold typically associated with strong community structure as described by Newman (2006)) indicate that the division in communities remains relatively minor, with the communities identified remaining tightly connected to the core, showing that the reinsurance market maintains a high degree of global interaction throughout the entire period. Furthermore, modularity presents a strong negative correlation with TCI ($\rho = -0.698$, $p < 0.001$), as visualized in Figure 4.6. This relationship shows that during periods of high systemic stress, an increase in total

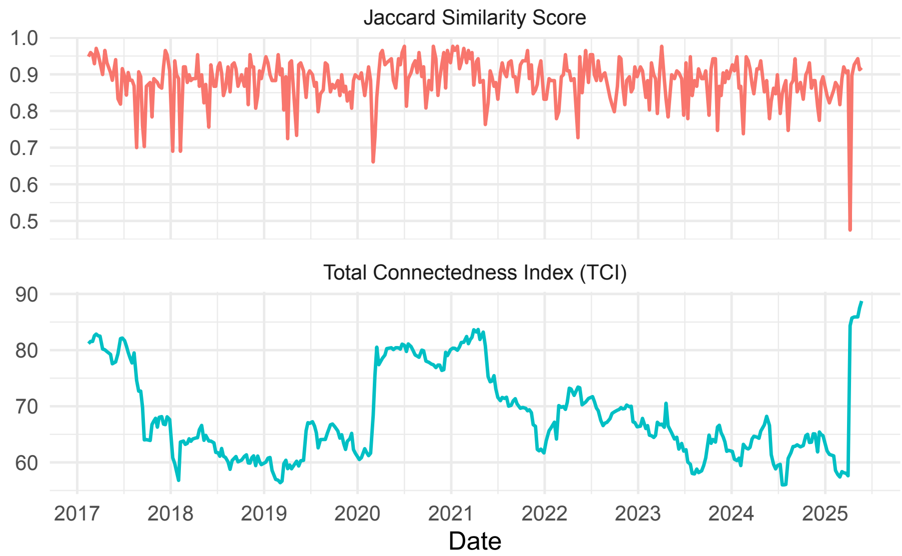

Figure 4.5: Jaccard similarity score (top plot) of the network's edge set between consecutive windows and TCI (bottom plot) for comparison of their relationship.

connectedness values reduces the distinction of regional or functional clusters, with variance influence propagating across the entire network without being contained in specific community boundaries.

4.4 Firm Level Dynamics

Firm level analysis is performed to assess changing dynamics at node level, identifying reinsurers that consistently act as variance distributors and those who are variance receivers, and address the evolution of the main components of the network during different time periods. To grasp the changing environment, the net-strength was used as proxy measure. We calculated the net values for each rolling window, to spot the net variance contributors and the variance receiver nodes. We aggregated each region's reinsurer to identify patterns related to specific areas to later move to the recognition of systemically important institutions that act as the main source of variance contribution over time.

4.4.1 Directional Connectedness

Figure 4.7 represents the values of net-strength for each company across the 349 rolling windows. Firstly, some companies maintained persistently a negative net balance throughout the whole period, like Chinese reinsurers China Re (1508.HK) and Taiping Re (0966.HK), remaining net variance receiver. This is consistent with their peripheral role identified in the static analysis, suggesting a limited spillover capacity due to their market position rather than an ephemeral phenomenon.

Secondly, one sharp positive spike is visible for the London-based company

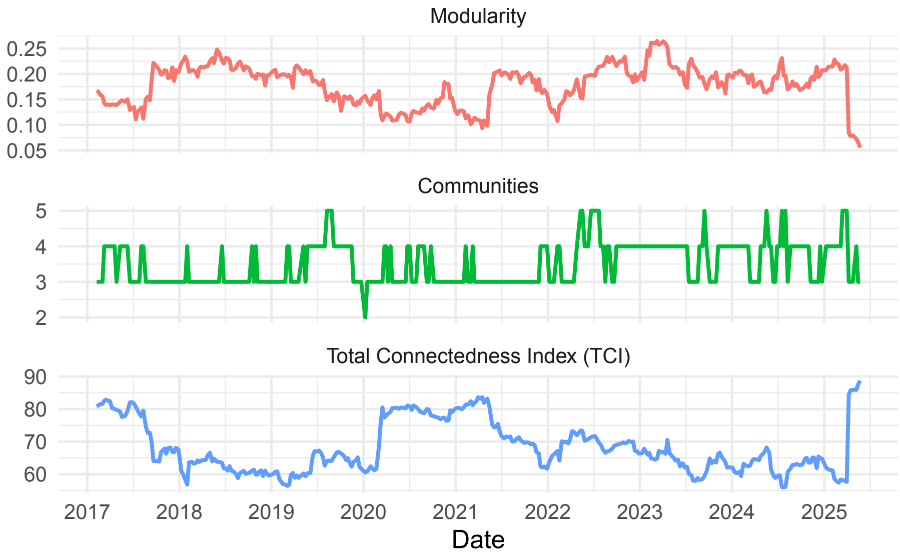

Figure 4.6: Evolution of the modularity (upper plot) and of the number of communities detected (middle plot) during the period analyzed. The TCI (bottom plot) is reported as well to better assess the community creation in relation with the system's connectedness.

Hiscox Ltd (HSX.L) in the late 2021 and early 2022. The value can be associated with high spikes in the volatility value likely due to a high-profile dispute regarding "Business Interruption" claims with the UK Financial Conduct Authority. The spike in volatility elevated the British reinsurer to a strong variance contributor, as its spillover effect strongly affected the network.

Third, the COVID-19 pandemic and the subsequent market crash in March 2020, and throughout the following year, reveals a clear separation of net roles. As the TCI entered the high connectedness regime, the contrast between net explainers (red) and net receivers (blue) sharpened, suggesting that lasting high connectedness periods could amplify the existing directional structure rather than raising all firms' role uniformly, as strong contributors become stronger, while receivers absorb more variance. Finally, the most recent windows (2024-2025) exhibit a notable shift: Japanese companies show increase net values and a switch to net transmitter roles. On the contrary, European firms show a decline in their net transmitter role, as their net-strength change appear opposite to the ones of the Japanese firms.

Regarding the aggregated regional net values, the trend remained constant for almost the entire sample. Asian companies remained in negative territory until the second half of 2023, reaching minimum peaks in July 2017 and April 2021. European companies on the other end remained net variance contributors until the second half of 2023, where the trend reverted in conjunction with the Asian role takeover. The European peak strength was achieved in July 2017, and presented high positive values after the pandemic burst, in 2020, for nearly 2 years, where a progressive decline began. American based companies presented a changing trend, with less extreme values compared to both the European

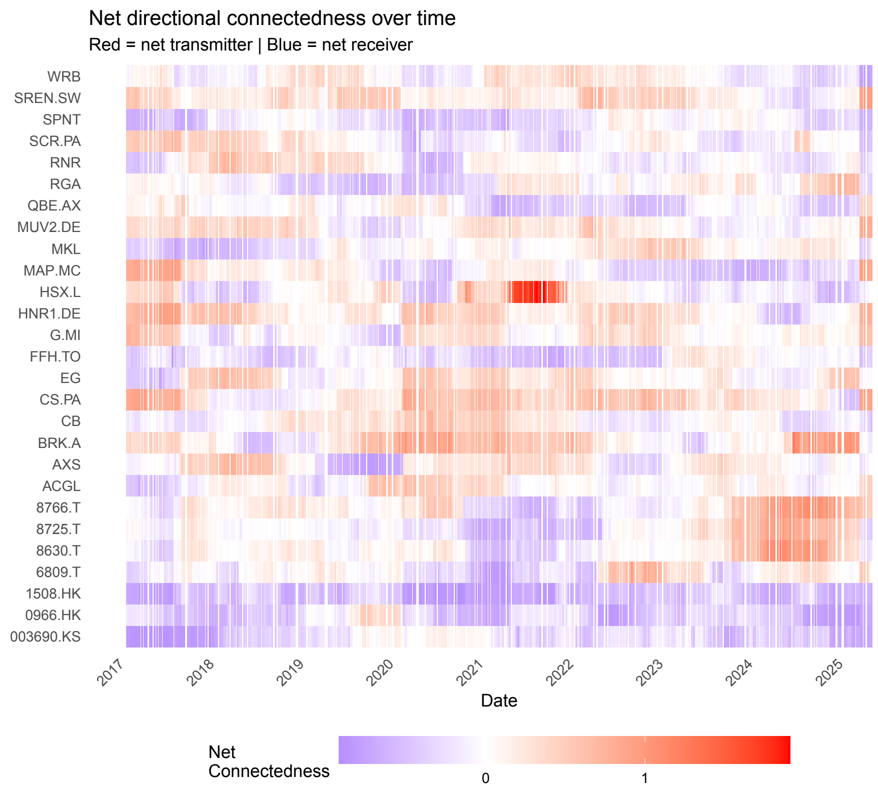

Figure 4.7: Heatmap representation of reinsurers' net-strength evolution over the windows analyzed.

and Asian markets, with movements typically ranging between the Asian and European values. The highest strength was reached in July 2021, where the market showed the highest similarity to the European one.

From this analysis, represented graphically in Figure 4.8, a strong negative correlation between the Asian and the European markets ($\rho = -0.74$, $p < 0.001$) can be noted. These findings remain consistent with the results of Chapter 3, where the European market explained variance movements, especially for the Asian market. However, the magnitude of the strength varied significantly during the sample period, showing a reverting trend in the last years of the sample.

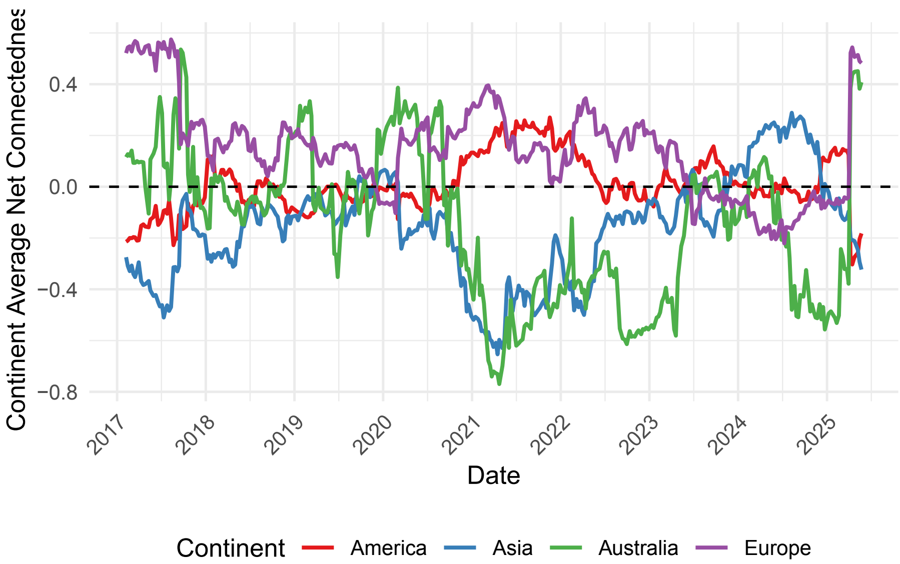

Figure 4.8: Net-strength aggregated by region in which the reinsurer has registered office.

4.4.2 Centrality

Eigenvector centrality, PageRank, and betweenness centrality were computed for each of the 349 rolling windows to assess the core-periphery structure across the various market regimes. A notable finding emerges: Chubb (CB) ranks among the top eigenvector and betweenness nodes, despite not being identified as central nodes in the static analysis. CB shows consistently elevated eigenvector values throughout the sample (mean eigenvector of 0.218), with relatively low standard deviation (0.035), showing a persistent structural position that was not captured by the full sample BVAR model.

To better grasp the dynamics of centrality, firms were classified into two groups based on their eigenvector centrality distribution: the five firms with the highest mean eigenvector centrality in the *persistent hubs*, while the five firms with the highest standard deviation as the *reactive firms*. Table 4.3 summarizes both groups.

It is worth noting the geographic composition of the two groups: persistent hubs are predominantly European (Chubb, Munich Re, Mapfre Reasseguros,

Table 4.3: Firms classified by eigenvector centrality dynamics. Persistent hubs show high mean centrality with low standard deviation; reactive firms exhibit high variability, indicating centrality shifts during the sample period.

Category	Firm	Mean Eigen	SD eigen	Continent
Persistent	CB	21.82	3.46	Europe
Persistent	MUV2.DE	21.31	4.54	Europe
Persistent	MAP.MC	21.08	6.01	Europe
Persistent	MKL	20.95	4.81	America
Persistent	G.MI	20.73	4.16	Europe
Reactive	SPNT	16.69	7.67	America
Reactive	BRK.A	15.04	7.43	America
Reactive	003690.KS	12.83	7.86	Asia
Reactive	QBE.AX	11.67	8.04	Australia
Reactive	HSX.L	10.14	7.36	America

Generali), with Markel Corp as the only non-European member, consistent with the European core documented in Chapter 3. Reactive firms, by contrast, are geographically dispersed. These firms occupy peripheral network positions during normal conditions but experience significant centrality shifts during stress episodes, transitioning from the periphery to the temporary core. This distinction carries the prudential message of the dynamic analysis. The rolling analysis reveals that hidden systemic risk does exist in a time-varying form, since reactive firms that appear peripheral in the full sample of Chapter 3 become transmitters precisely during high connectedness regimes. A static SIFI screen would miss these conditionally systemic institutions.

This conclusion reinforces the size-connectedness distinction highlighted in Section 3.4.3.1. Markel Corp. and Mapfre are not the firms a size-based screen would flag (Markel Corp. ranks fifteenth for market size, while Mapfre ranks nineteenth), yet both occupy persistently central positions across all rolling windows. The dynamic analysis thus shows that network embeddedness is not only a different dimension from size, but a stable one. The firms whose systemic importance is invisible to size and market share remain central through time, while the reactive firms reveal a further, conditional layer of systemic importance that no static balance sheet size-based indicator could anticipate.

The two groups also differ in how their centrality metric co-move. For reactive firms, eigenvector and PageRank are highly correlated ($\rho = 0.631$), meaning that when these firms become influential they also become variance receivers. For persistent hubs instead, the correlation is lower ($\rho = 0.217$), indicating a difference between the influence and absorption phases, maintaining a more specialized structural role. Persistent hubs tend to behave as disciplined intermediaries that separate their phase of influence from their phase of absorption, whereas reactive firms are swept up as crisis absorbers at the same moment they transmit.

Finally, comparing this classification with the later SIFI stability in the next section shows that central firms and net transmitters capture different dimensions of systemic importance. No company appears among both the persistent eigenvector companies and the chronic top-5 transmitters. This divergence suggests that the spillover magnitude may overlook firms that are structurally central to the network architecture, without necessarily dominating

Table 4.4: Number of times that the reinsurer appeared in the top 5 net-strength.

Company	Top 5 appearances	Top 5 pct.
CS.PA	215	61.605
BRK.A	153	43.840
HNR1.DE	153	43.840
SREN.SW	116	33.238
EG	110	31.519

directional variance flows.

4.4.3 SIFI Stability

To identify a systemically important institution for the reinsurance market, we classified the companies that consistently appeared as the strongest net values.

The companies that appeared consistently as the variance explainer of others' movements are reported in Table 4.4, and their net-strength evolution is represented graphically in Figure 4.9. Berkshire Hathaway (red line) is the most volatile firm in the group, as it fluctuates between being a net variance explainer and a variance receiver. AXA XL and Swiss Re exhibit the most consistent behavior, remaining in the positive territory for almost the whole sample. Their lines remain relatively stable, suggesting a consistent role in the industry. The German Hannover Re shows a significant regime shift, as it starts as a net variance contributor in 2017-2018, but experiences a downward trend in its net role, at the onset of COVID-19, a notable reversal for a firm that dominated variance contribution in the initial years. Everest Group instead showed the inverse trend, as it started as a variance receiver and transitioned to a variance contributor in late 2018 and early 2019.

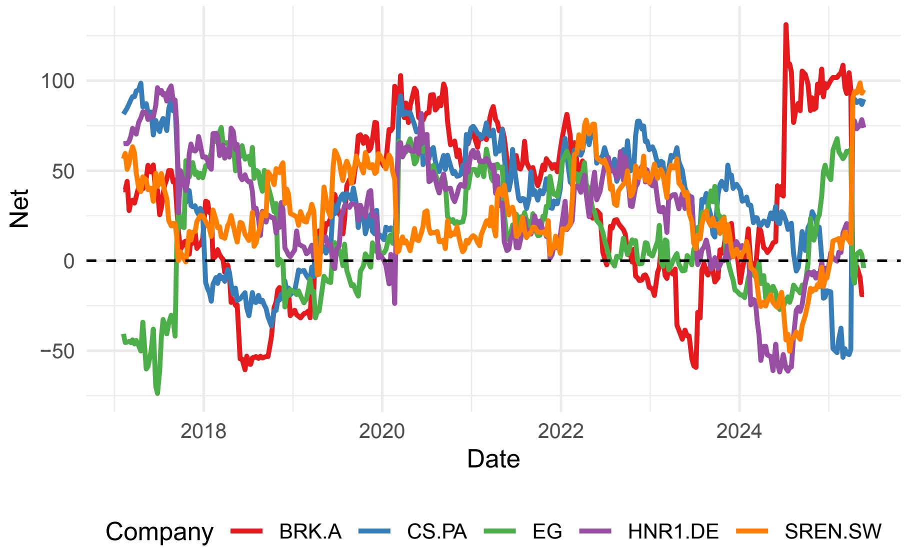

Figure 4.9: Net Evolution of the five most persistent net transmitters.

Table 4.5: Events considered in the event analysis. The estimated insured losses are expressed in billion $.

Date	Event	Type	Estimated Insured Loss	Loss Type
2017-09-01	HIM Hurricanes	NatCat	92	Liability-side
2018-11-15	CA Wildfires	NatCat	18	Liability-side
2018-12-24	Equity Crash	Financial	NA	Asset-side
2020-03-16	COVID-19	Pandemic	40	Both
2021-02-15	Winter Storm Uri	NatCat	15	Liability-side
2022-02-24	Ukraine War	Geopolitical	10	Asset-side
2022-09-28	Hurricane Ian	NatCat	55	Liability-side
2023-03-10	SVB Collapse	Financial	NA	Asset-side
2024-07-19	CrowdStrike	Cyber	NA	Liability-side
2024-10-05	Helene Milton Hurricane	NatCat	44	Liability-side
2024-10-29	Spanish Floods	NatCat	4	Liability-side
2025-01-07	LA Wildfires	NatCat	45	Liability-side
2025-04-02	Trump Tariffs	Financial	NA	Asset-side

4.5 Event Analysis

This section focuses on the analysis of events that had an impact both on the reinsurance market directly, like natural disasters, and other major geopolitical and financial events that might have affected the reinsurance industry indirectly, to assess the impact over the estimated connectedness of the system and the response of the network dynamics to these events.

4.5.1 Event Selection

The events reported are gathered through web searches and annual reports over the main events sustained by the reinsurance companies analyzed. The cases considered are the major occurrences of the last decade, considering both the events related directly to the reinsurance market and those affecting the financial and economic world as a whole in terms of magnitude of social and economic impact. Thirteen events are analyzed, spanning five categories. Seven are natural catastrophes ranging from the HIM hurricanes trio of 2017 ($92B insured losses) to the 2025 Los Angeles wildfires ($45B), testing liability-side propagation through underwriting claims. Three financial events - the 2018 Q4 equity sell-off, the 2023 SVB collapse, and the 2025 US trading tariffs - shock-test the asset-side propagation through investment portfolio revaluation. The 2022 Ukraine conflict represents a geopolitical event with mainly asset-side contagion channels, while the 2024 CrowdStrike outage represents an emerging cyber risk. The COVID-19 pandemic uniquely combined both channels, with simultaneous equity market collapse and widespread insurance claims. Table 4.5 reports the full event list with estimated insured losses and shock channel specification, while Figure 4.10 overlays the event dates onto the TCI time series, allowing visual identification of the system's connectedness response to each shock.

4.5.2 Event Impact Ranking

To assess the impact that each event yields on the network dynamics, we observed the Total connectedness evolution around the time of the event. The TC mean of the previous 20 windows and the peak total connectedness measure for the next

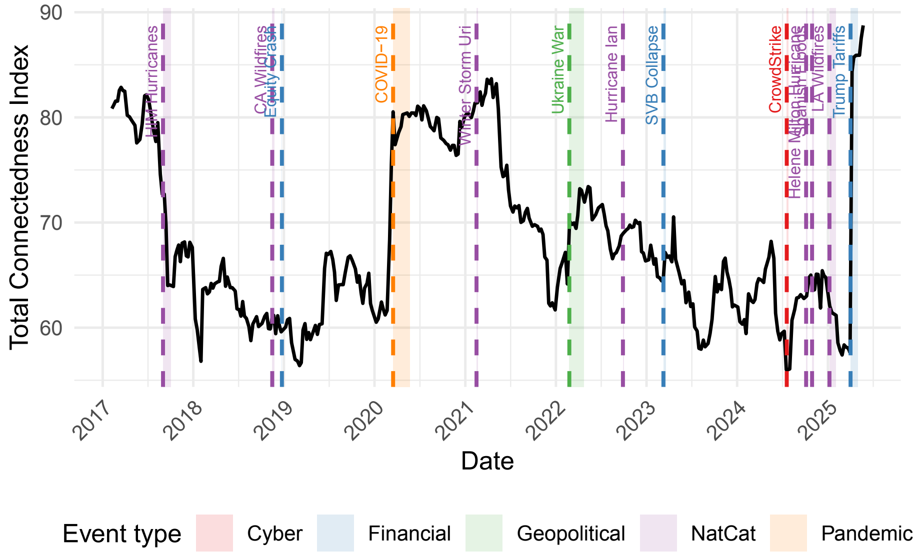

Figure 4.10: Total connectedness index and major events considered.

10 windows were computed (a 50-day time period), and compared to identify the events that created the highest change in the TCI. These findings are reported in Table 4.6, and are classified by their impact on the TC response.

To formally assess whether the observed event response are distinguishable from random fluctuations, we conducted three statistical tests. First, a Monte Carlo permutation test drawing 10,000 random 11 event clusters from the 349 window pool yields $p = 0.108$ for the mean ΔTCI across the pooled event set. Second, a Kruskal-Wallis test across the three loss-type groups returns $H = 2.13$ ($df = 2$, $p = 0.344$). Third, a one-sided Wilcoxon rank-sum test comparing asset-side versus liability-side ΔTCI values yielded $W = 22$, $p = 0.184$. None of these tests is significant, reflecting the fundamental constraint of having only a 13-event sample. The event analysis is therefore exploratory: the patterns described below constitute descriptive evidence for hypothesis generation rather than confirmed causal claims. Nevertheless, descriptive evidence reveals a consistent pattern: the two events producing the highest TCI increases - the Trump tariff shock of April 2025, and the COVID-19 market crash of March 2020 - both involved asset-side transmission through simultaneous investment portfolio revaluation. By contrast, the largest pure liability-side event, the HIM hurricane trio, with estimated insured losses of $92 billion, produced a negative TCI change, consistent with the idiosyncratic rather than systemic risk propagation.

Together, these patterns point to the chapter's central economic conclusion: the reinsurance sector's systemic risk originates on the asset side, through its role as a large holder of correlated financial assets, rather than on the liability side of catastrophe underwriting. The reduction of interconnectedness value in the most severe catastrophe loss is not a paradox, but rather an evidence of the risk-pooling model working as designed, since large insured losses are absorbed idiosyncratically by the exposed firms and do not become contagious. This aligns with the FSB and IAIS view that traditional underwriting activities do not pose

Table 4.6: Impact of each event on the TCI level and volatility (pre- vs post-event).

Event	Loss Type	TCI pre	TCI post peak	Change	Pct. Change
Trump Tariffs	Asset-side	62.01	88.75	26.74	43.12
LA Wildfires	Liability-side	62.46	84.35	21.89	35.04
COVID-19	Both	64.16	80.53	16.37	25.51
Ukraine War	Asset-side	66.89	73.42	6.53	9.77
Winter Storm Uri	Liability-side	78.75	83.68	4.92	6.25
Spanish Floods	Liability-side	61.77	65.42	3.65	5.91
Helene Milton Hurricane	Liability-side	61.85	65.42	3.57	5.77
SVB Collapse	Asset-side	67.85	70.54	2.68	3.95
CrowdStrike	Liability-side	62.75	64.78	2.02	3.23
Equity Crash	Asset-side	60.60	60.88	0.28	0.46
CA Wildfires	Liability-side	61.19	61.17	-0.02	-0.04
Hurricane Ian	Liability-side	70.28	70.22	-0.07	-0.10
HIM Hurricanes	Liability-side	80.06	72.70	-7.36	-9.19

systemic risk, but this is rather created from financial-market activities.

One notable exception is the Los Angeles wildfires in January 2025 ($\Delta TCI = 21.89$), classified as a liability-side event. However, the post-event window for this event coincides with the turmoil related to the market uncertainty following the January 20 American presidential inauguration, suggesting that the measured TCI increase reflects the asset-side pressure rather than a pure underwriting loss response. It is also worth noting that the LA wildfires and the Trump tariffs events do not have a full subsequent window, given their recent nature, not allowing to verify the complete subsequent response of the system to these events.

4.5.3 Deep-Dive Analysis

Five events are selected to perform a deeper analysis on the system's connectedness dynamics, each representing a distinct shock archetype:

- HIM hurricanes (2017): the largest liability-side shock in the sample ($92B), testing whether catastrophe claims propagate systemically or remain idiosyncratic;
- COVID-19 (2020): a simultaneous asset-side and liability-side shock, combining a global equity market crash with widespread business interruption and event cancellation claims;
- Ukraine conflict (2022): a geopolitical shock transmitting primarily through the asset-side via energy price volatility, sanctions-related asset write-downs, and sovereign risk repricing;
- SVB collapse (2023): a financial sector shock testing contagion from banking sector credit risk to the reinsurer investment portfolios;
- Trump "Liberation Day" tariffs (2025): a pure asset-side shock, producing the highest TCI in the sample, included despite the limited post-event window.

These five events were selected to represent distinct shock archetypes, enabling a comparative analysis of how different transmission channels affect connectedness dynamics. For each event, the short-run TCI reaction is analyzed by identifying the peak value within the subsequent 40 rolling windows. The pre-crisis mean

Table 4.7: Comparison of pre and post event TCI values.

	HIM Hurricanes	COVID-19	Ukraine War	SVB Collapse	Trump Tariffs
Pre crisis mean	80.06	64.16	66.89	67.85	62.01
Peak TCI value	82.86	82.43	73.42	70.54	88.75
Peak TCI window	-19.00	40.00	8.00	5.00	6.00
Recovery time	NA	NA	1.00	1.00	NA
Post crisis TCI mean	63.59	79.06	68.09	62.58	NA
Change	-16.47	14.90	1.20	-5.28	NA
Pct. change	-20.57	23.22	1.79	-7.78	NA

is computed over the 20 windows preceding the event; the post-crisis mean is computed over windows 20 through 40 following the event, excluding the acute phase to assess whether the shock induced a persistent regime change.

The HIM hurricanes generated the highest negative response to the TC ($\Delta TCI = -7.36$), as the event occurred during an already-declining trend. The connectedness level before the event was on a high regime state, with an already mean reverting trend, as the peak was registered 8 windows prior to the event. However the massive NatCat event and the estimated loss of 92$B did not stop the declining trend. This supports the idea that liability-side losses are absorbed asymmetrically by the system, as only exposed firms are hit by the shock, which reduces the co-movement rather than increasing it.

The COVID-19 pandemic, the second highest ΔTCI change among the selected events, showed a steady increase after the initial surge following the event, with a post-crisis mean 23% above the pre-crisis level, computed for windows 20 to 40. This event caused a regime change, moving a low connected system to a high connected one, with the peak registered at the 40^{th} window, meaning that the increase persisted during the entire period, with only brief temporary dips in the TCI within the overall upward trend, resulting in a more connected network. A similar reaction is visible from another asset-side event, the Trump Tariffs of April 2025. Despite the right-censoring of the post-event window, the acute response was fully captured, resulting in the highest TCI change between the peak and the pre-event crisis, with a $\Delta TCI = 26.74$ increase. Although the post-crisis mean cannot be computed, due to the missing windows, the volatility shocks caused by the market crash and equity market uncertainty induced co-movement among reinsurance companies, increasing the connectedness of the system.

The outbreak of the Russian-Ukrainian conflict instead caused a delayed response that reached its peak value after 8 rolling windows, while just 5 were needed for the SVB collapse. Moreover, the two events had an equal response subsequent to the peak TCI value, as both the financial crisis of 2023 and the Ukraine war of 2022 recovered just after one rolling window, with the bank sector turmoil causing a reduction from the pre crisis level by $\Delta TCI = 2.68$, while the geopolitical event left the TCI almost unchanged ($\Delta TCI = 7$), not producing significant regime changes. This restricted effect on TCI shows that the contagion to the reinsurance market was contained. Instead, for the Ukrainian case, the system responded more slowly, stabilizing the post crisis mean after 16 windows. Despite the initial similar reaction, the geopolitical importance of the conflict and the relative instability it introduced meant that its effect was more persistent over time.

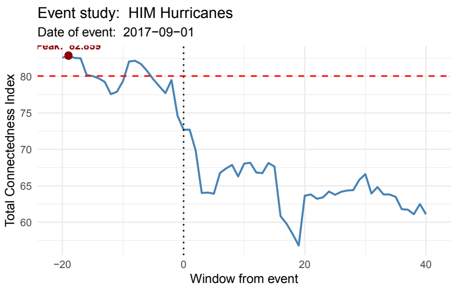

(a) Harvey-Irma-Maria hurricanes

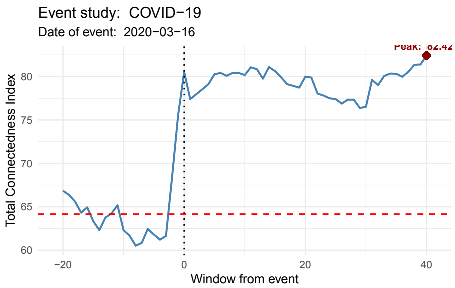

(b) COVID-19 pandemic

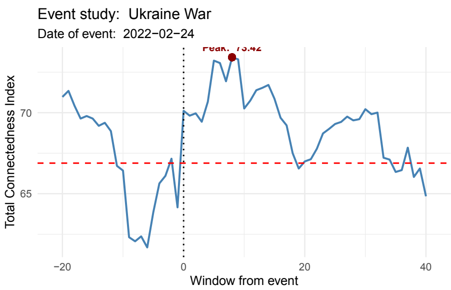

(c) Russian-Ukraine conflict

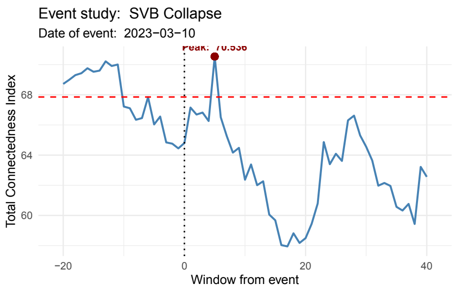

(d) Silicon Valley Bank and subsequent Credit Swiss collapse

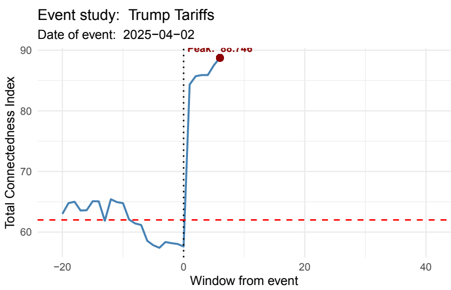

(e) Trump's "Liberation day" tariffs

Figure 4.11: TCI plot for the 20 windows prior to the event and the 40 windows subsequent.

4.5.3.1 HIM Hurricanes

The liability-side shock produces a clear shift in roles among specialty reinsurers. RenaissanceRe, whose business model is concentrated in US property catastrophe reinsurance, transitions from the dominant net transmitter to a net receiver, consistent with its disproportionate exposure to HIM losses. Conversely, AXA XL's diversified European platform increased its net transmitter role, suggesting that the shock redistributes influence from exposed specialists toward less-exposed diversified generalists. This asymmetric firm-level response explains the negative aggregate TCI change: rather than synchronizing the network, the NatCat shock fragmented it along geographic exposure lines.

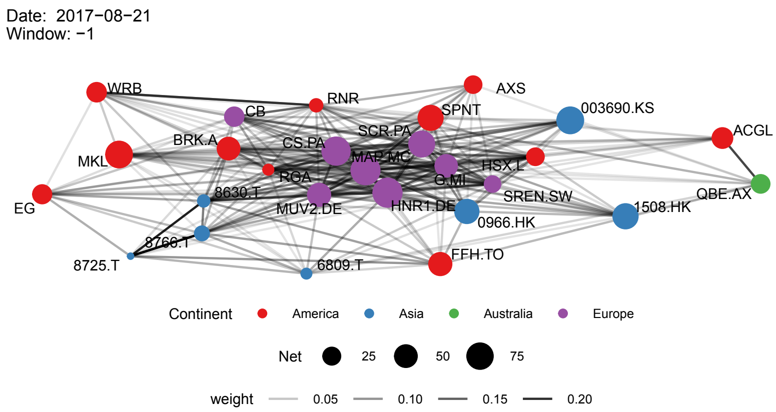

Figure 4.12: Network at the window before the HIM (Harvey, Irma, Maria) hurricanes.

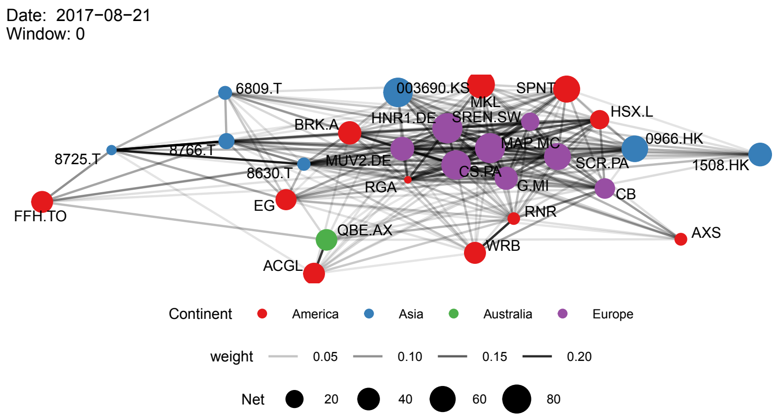

Figure 4.13: Network at the window for the HIM (Harvey, Irma, Maria) hurricanes.

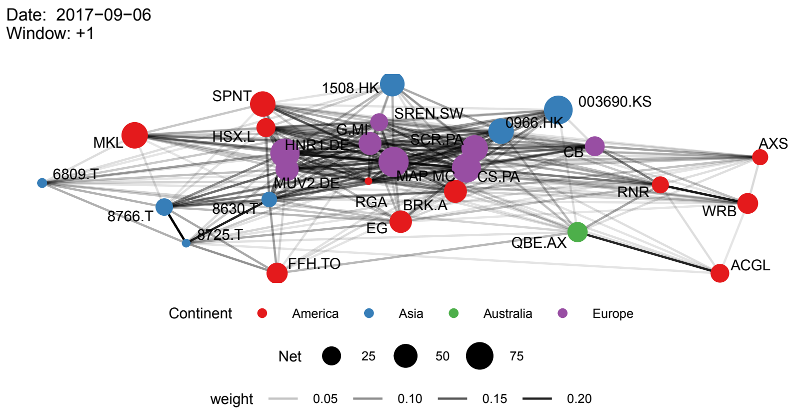

Figure 4.14: Network at the window after the HIM (Harvey, Irma, Maria) hurricanes.

As Figure 4.12 through Figure 4.14 shows, the networks display a trend of declining TCI value, with subsequent windows showing a more dispersed periphery, with net variance contribution roles becoming less clear and more uniform across the network. This trend is consistent with what is shown in Figure 4.11, and provide a visual representation of the ability of the sector to sustain even large catastrophic events, without significantly increasing the interconnectedness of the system, raising the volatility spillovers between nodes.

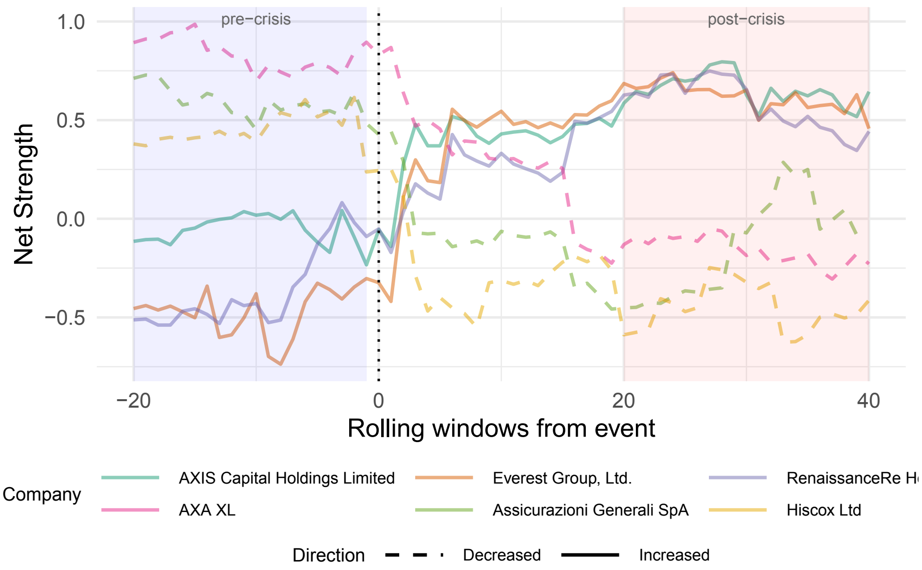

Figure 4.15: Three largest positive changes in net-strength (solid lines) and three largest decreases in net-strength (dashed lines) before and after the HIM event.

Figure 4.15 shows instead the three companies that underwent the highest decrease and increase in net-strength values in the pre-crisis and post-crisis period. The windows used to analyze the pre-crisis are the 20 preceding the event, while the post-crisis values are computed from windows 20 to 40, to consider a rather stabilized phase. The highest increase or decrease are calculated on the average values of the pre- and post- event periods. In this case, the companies that changed their role in the network are the ones more exposed to risk: American companies exposed to hurricanes risk increased their transmitter role, but the overall market was not affected by this increase, and it rather decreased the overall transmission role of the other major reinsurers.

4.5.3.2 COVID-19

The simultaneous asset-side and liability-side shock rearranges roles along different dimensions. European reinsurers Generali and Munich Re, among the largest by total assets, emerge with increased net variance explainer roles. Their dominant investment portfolios, heavily allocated to European sovereign and corporate bonds, experience simultaneous mark-to-market losses during the March 2020 liquidity shock, consistent with the asset-side channel amplifying their influence on the network. Contrariwise, Asian firms (Taiping Re and The Toa Re) and specialty player (QBE) declines. The pattern is consistent with European equity markets, heavily weighted in reinsurance portfolios, driving the initial contagion signal, while Asian and Australian firms respond with delay.

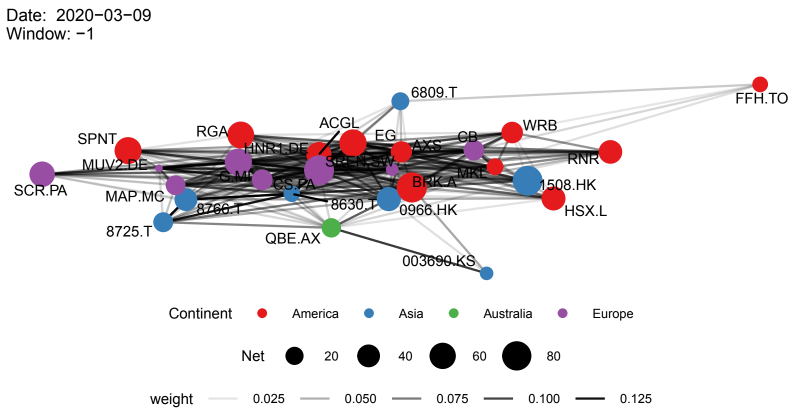

Figure 4.16: Network at the window before the COVID-19 market decline.

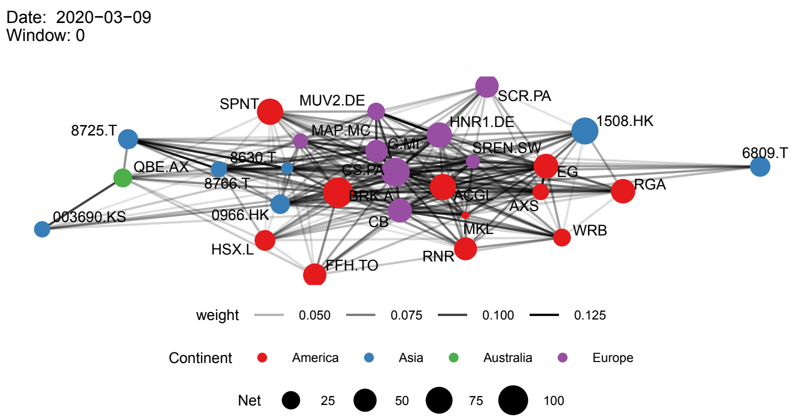

Figure 4.17: Network for the window at the COVID-19 market decline.

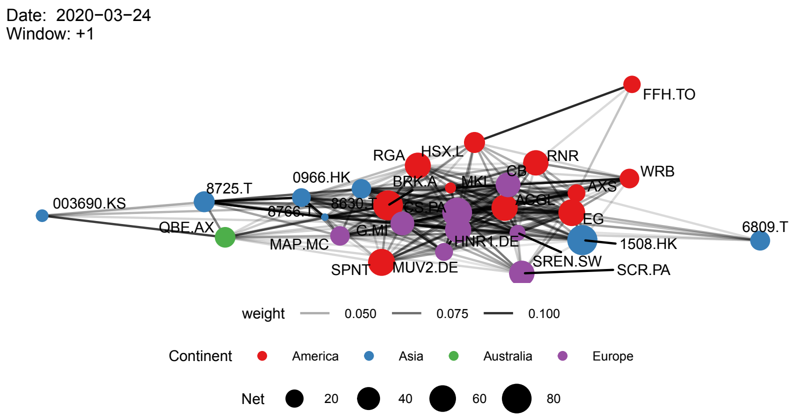

Figure 4.18: Network at the window after the COVID-19 market decline.

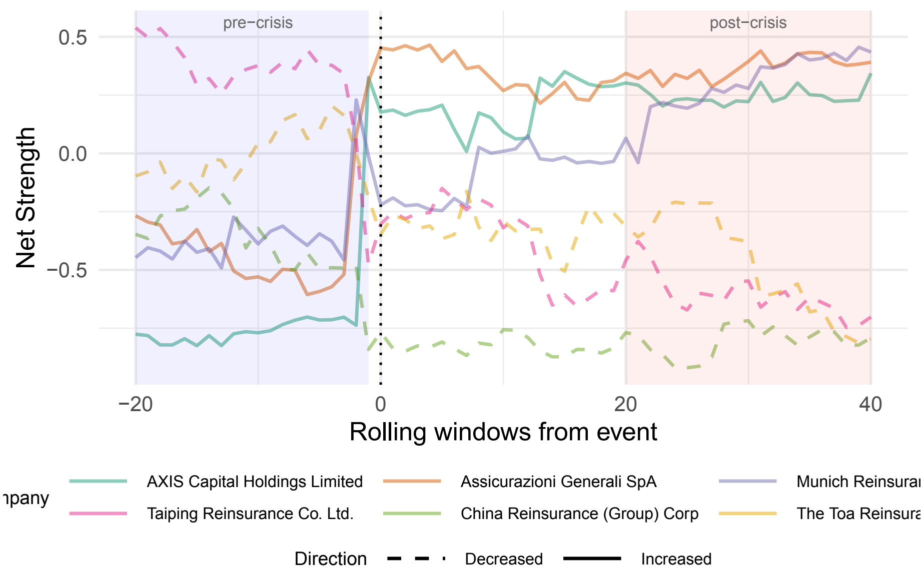

Figure 4.19: Three largest positive changes in net-strength (solid lines) and three largest decreases in net-strength (dashed lines) before and after the COVID-19 event.

Visual representation of the structure dynamics and main cluster evolution can be seen in Figure 4.16, Figure 4.17, and Figure 4.18: all windows show European reinsurers forming a dense central cluster connected by thicker edges, while Asian firms remain at the periphery of the network with smaller sizes. The high connectedness values for the pre-event (77), event (80.5) and post-event (77.50) are a consequence of the spike in the previous windows, as shown in Figure 4.11.

The tight connectedness value (all in the high regime section), and the consequent close disposition of the nodes in the plots show the systemic vulnerability of crisis of this type: the volatility explanation spread more rapidly due to the conformation given by the stress period, and the high interconnectedness of the global reinsurance market.

Figure 4.19 confirms the dual-channel rearrangement: Generali and Munich Re climb in net-strength across the post-event window while Taiping Re and China Reinsurance Corp decline, as market losses on European bond portfolios lift the large European contributors over Asian and specialty firms.

4.5.3.3 Ukraine War

Although the Ukrainian war affected the European continent more closely, the highest changes in net variance explanation roles did not include dominant European reinsurers. The highest positive changes were in fact seen for the Japanese (6809.T and 8725.T), while Hiscox Ltd saw its net variance role decline. However, the absence of major European spikes in net-strength values can be explained looking at the recent past: COVID-19 already created a strong position of variance contribution to the whole market from major European and American reinsurers, as Figure 4.7 showed. The geopolitical crisis of early 2022 consolidated this market condition rather than radically disrupting it.

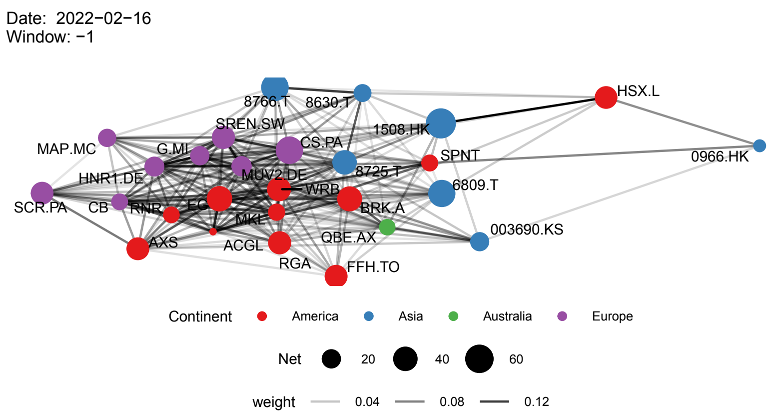

Figure 4.20: Network at the window before the onset of the Russia-Ukraine war.

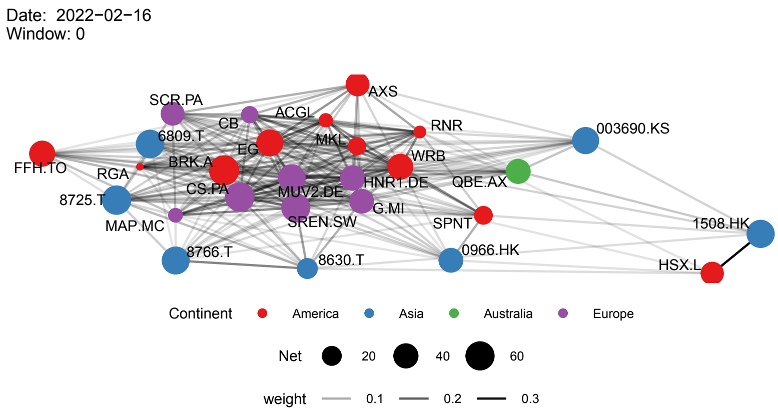

Figure 4.21: Network for the window of the onset of the Russia-Ukraine war.

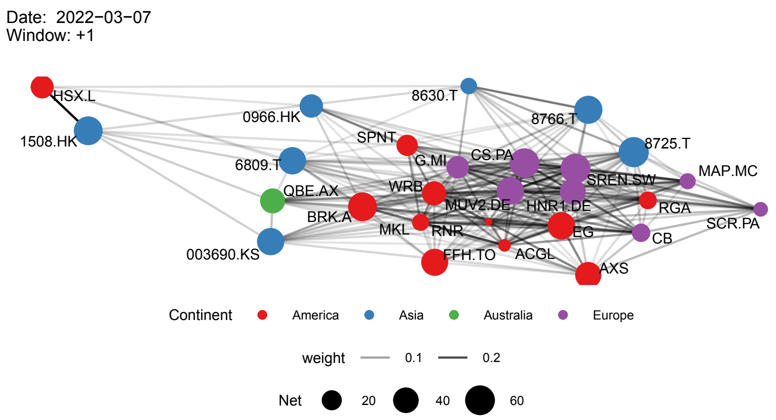

Figure 4.22: Network at the window after the onset of the Russia-Ukraine war.

Figure 4.20 shows that the pre-event network exhibits elevated node sizes, continuing the high connectedness regime established during COVID-19. The relative connectedness measure remains stable throughout the event and the post-event windows (Figure 4.21, Figure 4.22), without any relevant detectable structural change.

Figure 4.23 shows only modest, mean-reverting shifts, with magnitudes small enough to confirm the conflict consolidated the existing COVID period hierarchy rather than reshaping it.

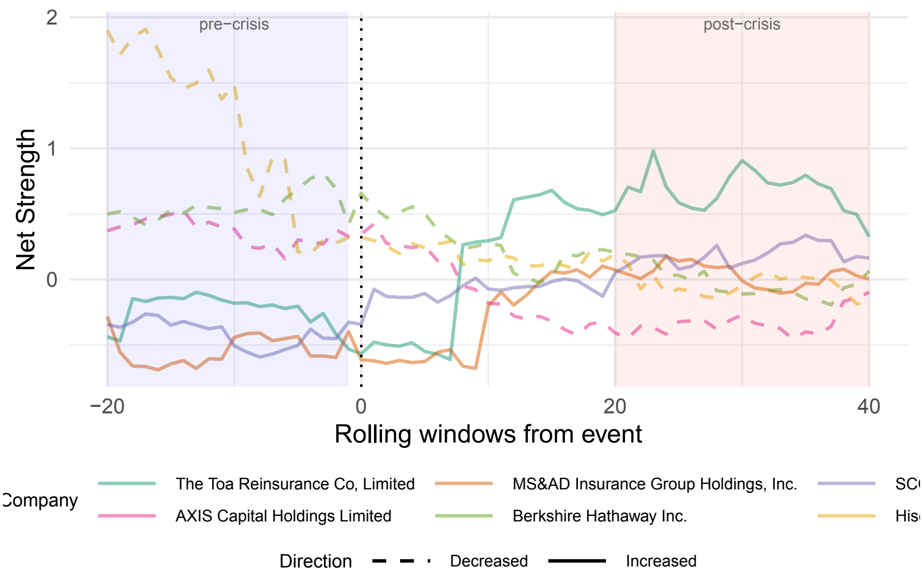

Figure 4.23: Three largest positive changes in net-strength (solid lines) and three largest decreases in net-strength (dashed lines) before and after the Ukraine-Russian conflict."

4.5.3.4 SVB Collapse

The banking crisis of March 2023 saw the Japanese reinsurer Tokio Marine, changing its position from a net variance receiver to a net variance explainer, as shown in Figure 4.27. Even more marked was the movement of AXIS Capital, which moved from a net-strength value of ≈ -0.5 to a peak value of ≈ 0.8. This is consistent with the firm's exposure to US fixed-income assets and the interest rate shock that precipitated SVB's failure and simultaneously created mark-to-market volatility across Japanese and dollar denominated bond portfolios, increasing the variance contribution of these companies in the network. Swiss Re, Toa Re and Hannover Re instead moved in the opposite direction, transitioning to net receiver positions. The SVB crisis activated peripheral firms with asset or liability exposure in the US banking sector.

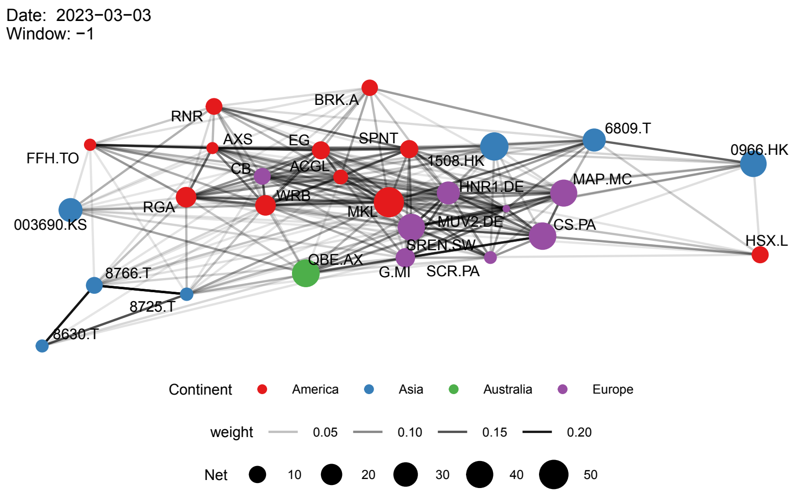

Figure 4.24: Network at the window before the SVB collapse.

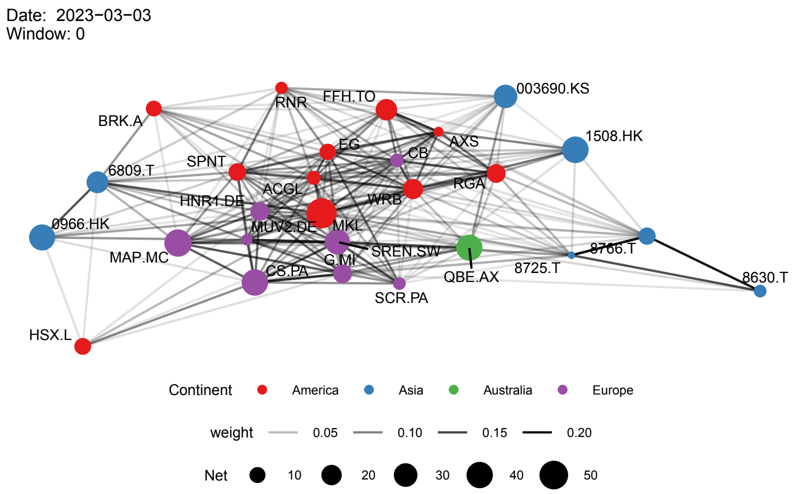

Figure 4.25: Network at the window of the SVB collapse.

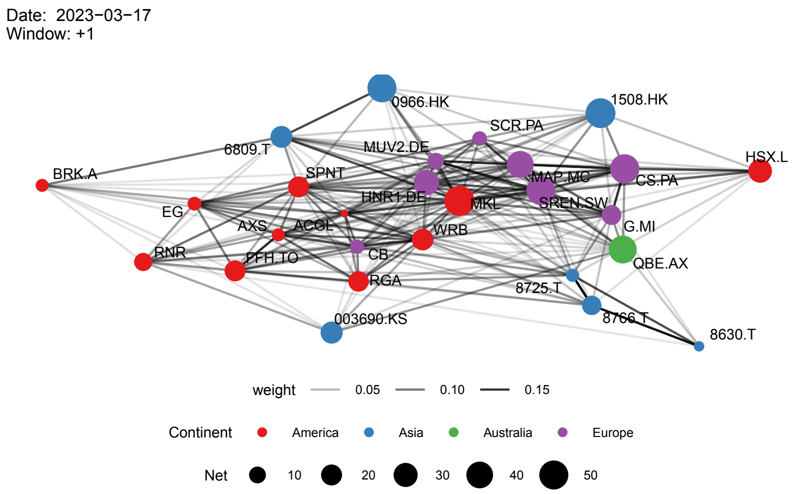

Figure 4.26: Network at the window after the SVB collapse.

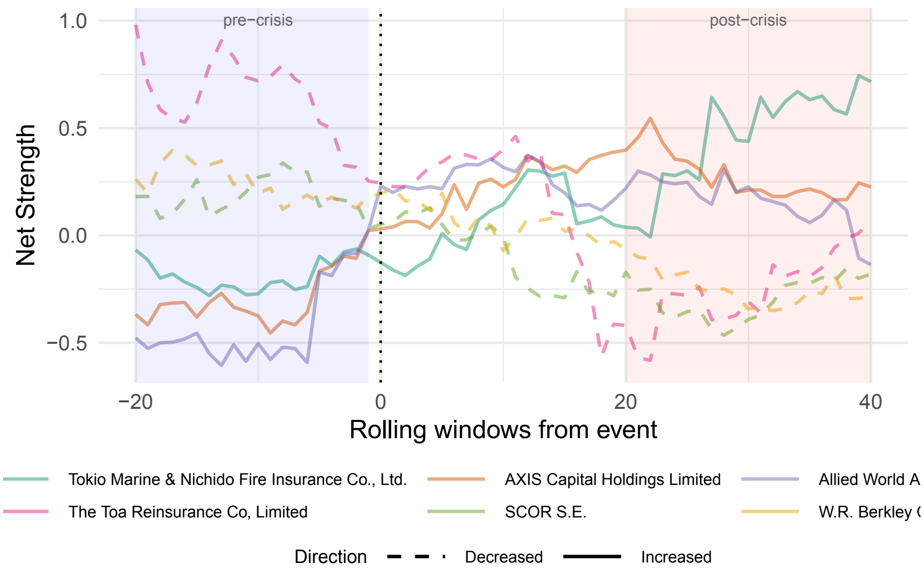

Figure 4.27: Three largest positive changes in net-strength (solid lines) and three largest decreases in net-strength (dashed lines) before and after the SVB collapse.

Figure 4.26 clearly shows an already increased node size for Asian companies, compared to Figure 4.25 and Figure 4.24, consistent with the finding in Section 4.4.1 where the trend of major variance contributors started to shift towards Japanese and Asian reinsurers. The increased connectedness is coherent with what is shown in Figure 4.11, as the post-event window saw a TCI increase of 2.64 indicating a relatively quick market response to the event, but a later tendency to return to the pre-event TCI measure.

Figure 4.27 illustrates the role reversal of US-exposed firms: AXIS Capital and Tokio Marine rise sharply into net transmitter role, while Swiss Re and Toa Re drop to net receivers, isolating the rate shock to firms with US fixed income exposure.

4.5.3.5 Trump Tariffs

The pure asset-side shock compresses firm-level net-strength toward the mean, as the indiscriminate equity-bond sell-off affects all reinsurers simultaneously. Swiss Re, the pre-crisis dominant transmitter (Net-Strength ≈ 1.0), declines sharply, while previously net receiver firms, like Berkshire Hathaway, AXA XL, and Mapfre, increase rapidly becoming the net variance contributor companies of the system. The post-event window is truncated due to sample end, limiting interpretation of the post-crisis regime.

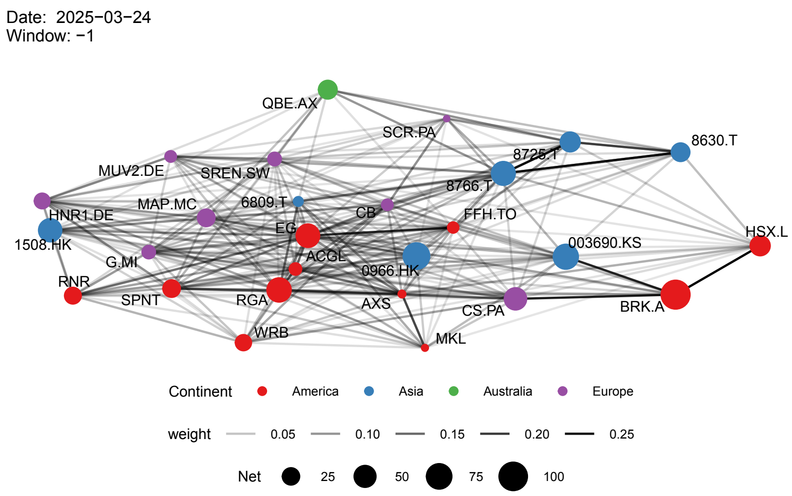

Figure 4.28: Network at the window before the Trump “Liberation Day” Tariffs.

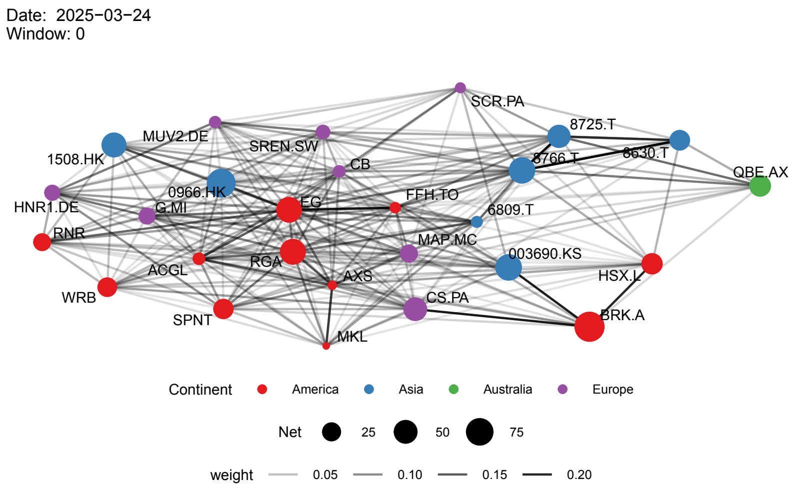

Figure 4.29: Network at the window of the Trump “Liberation Day” Tariffs.

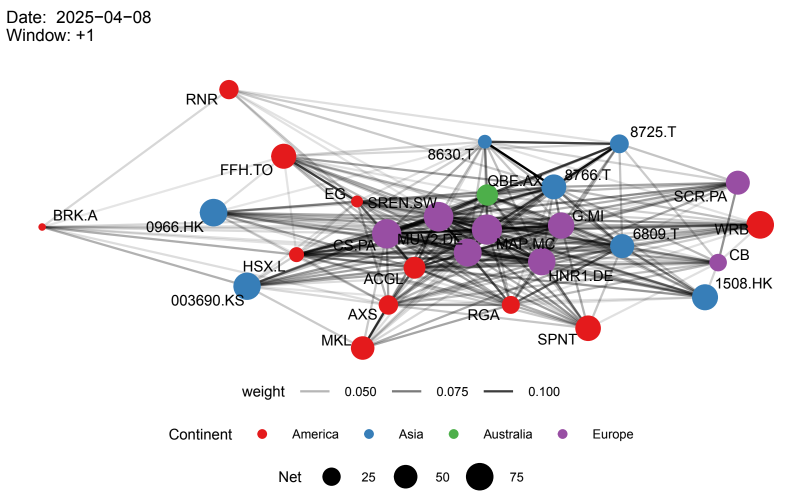

Figure 4.30: Network at the window after the Trump “Liberation Day” Tariffs.

The Trump tariffs event provides the clearest visual contrast across all five events. The network for window -1 (Figure 4.28) is dispersed with major variance contribution roles in American and Asian companies, as the TCI value confirms (57.62). By window +1 (Figure 4.30), all nodes collapse into a tightly compressed center, increasing the connectedness index by 26.74 TCI units value and showing

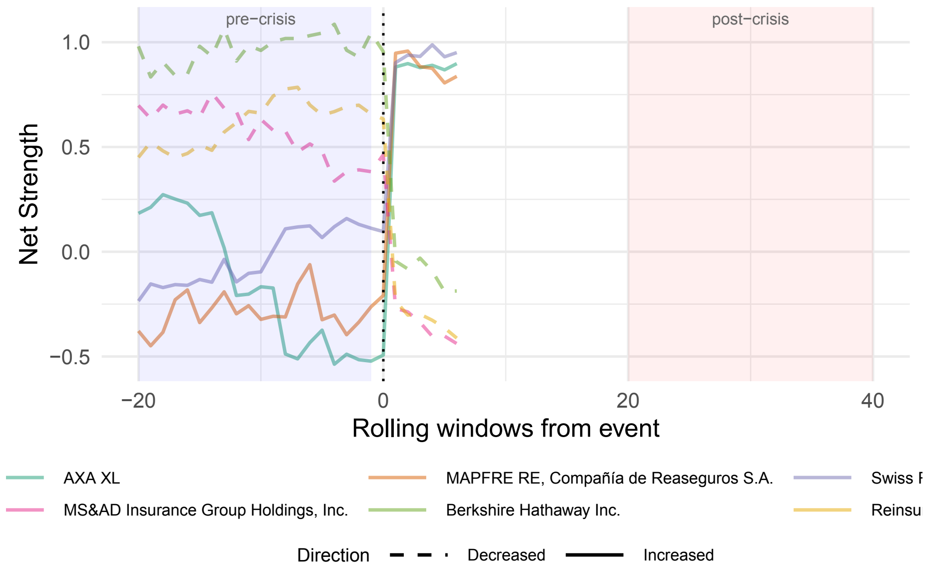

Figure 4.31: Three largest positive changes in net-strength (solid lines) and three largest decreases in net-strength (dashed lines) before and after the Trump tariffs.

the return as variance explainer of European reinsurers, a behavior possibly explained by the historical commercial partnership between European countries and the US, that faced increasing tumbles with US tariffs threats. The rapid response and structural change of the network show that systemic events external to the industry can have an impact on the dynamics of variance spillovers.

4.5.4 Cross-Event Stability of Firm's Rankings

To test whether the ranking of net transmitters is preserved across shock types, the pairwise Spearman rank correlation of company net-strength ranking across the five crisis networks is computed (Table 4.8). Two pairs of events show statistical significance at the 5% level: COVID-19 and Ukraine conflict ($\rho = 0.466$, $p = 0.012$), and the Ukraine conflict and the SVB collapse ($\rho = 0.462$, $p = 0.012$). This suggests that for moderate shocks with an asset-side component, the structural hierarchy of net transmitter, dominated by large European reinsurers, is partially preserved. However, neither the HIM hurricanes nor the Trump tariff shock exhibit statistical significance with any other.

4.5.5 Cross-Event Net Delta Correlation

To test whether the same firms consistently change their systemic role in the same direction across events, we computed the spearman rank correlation of net-strength variations (Δ Net-Strength) for the 27 companies in each event pair (Table 4.9). No pairwise correlation achieved statistical significance following adjustments for multiple comparisons. The strongest observed relationship is between the SVB collapse and Trump tariff shock, indicating that firms

Table 4.8: Pairwise Spearman correlation of companies' net-strength ranking across the five events.

Event 1	Event 2	ρ	p-value
COVID-19	Ukraine War	0.53	0.00
COVID-19	SVB Collapse	0.12	0.56
COVID-19	HIM Hurricanes	0.30	0.12
COVID-19	Trump Tariffs	0.21	0.28
Ukraine War	SVB Collapse	0.40	0.04
Ukraine War	HIM Hurricanes	0.40	0.04
Ukraine War	Trump Tariffs	0.23	0.24
SVB Collapse	HIM Hurricanes	0.12	0.56
SVB Collapse	Trump Tariffs	0.34	0.09
HIM Hurricanes	Trump Tariffs	0.30	0.13

Table 4.9: Spearman correlation of net-strength variation for the five events considered.

Event 1	Event 2	ρ	p-value
HIM Hurricanes	COVID-19	-0.22	0.28
HIM Hurricanes	Ukraine War	0.04	0.86
HIM Hurricanes	SVB Collapse	0.40	0.04
HIM Hurricanes	Trump Tariffs	-0.35	0.07
COVID-19	Ukraine War	-0.35	0.07
COVID-19	SVB Collapse	-0.21	0.30
COVID-19	Trump Tariffs	0.22	0.26
Ukraine War	SVB Collapse	-0.35	0.08
Ukraine War	Trump Tariffs	-0.06	0.78
SVB Collapse	Trump Tariffs	-0.29	0.14

exhibiting heightened net transmitter functionality during the banking crisis subsequently demonstrated reduced functionality during the trade policy shock, and vice versa. The line plots of the net-strength change confirm this result: the company composition does not follow a specific pattern, and varies across events, with no single firm consistently maintaining a dominant position. This absence of stable, cross-event roles implies that firm-level systemic importance within the reinsurance sector is contingent upon the nature of the shock, rather than reflecting a fixed structural property of the network.

Taken together, these results suggest a partial stable core hierarchy of net transmitters that persists across moderate asset-side shocks, co-existing with event-specific perturbations that prevent any single firm from consistently amplifying or absorbing shocks across all crisis types.

4.6 Early Warning Signal - VIX Analysis

To test the predictive capacity of TCI relative to the whole market movements, we tested the reinsurance network's volatility against the Chicago Board Options Exchange Volatility Index (VIX), which is usually used as a proxy for market volatility, as in Diebold and Yılmaz (2014). The data for the VIX were collected from Yahoo Finance. To ensure alignment with the rolling window methodology - which used a 5-day step size - the VIX was aggregated into a 5-day trailing average, mapping the macroeconomic sentiment directly to the time-step of the connectedness index. Furthermore, to avoid spurious regression problem, as highlighted by Granger and Newbold (1974), related to the high persistent TCI resulting from the rolling window construction - attributable to 98% of shared data created by the 5-day shift - the cross-correlation was evaluated on the first difference of the two series (ΔTCI and ΔVIX). The innovation series were then tested applying a cross correlation function and a multi-lead lag correlation, and the latter results are graphically shown in Figure 4.32.

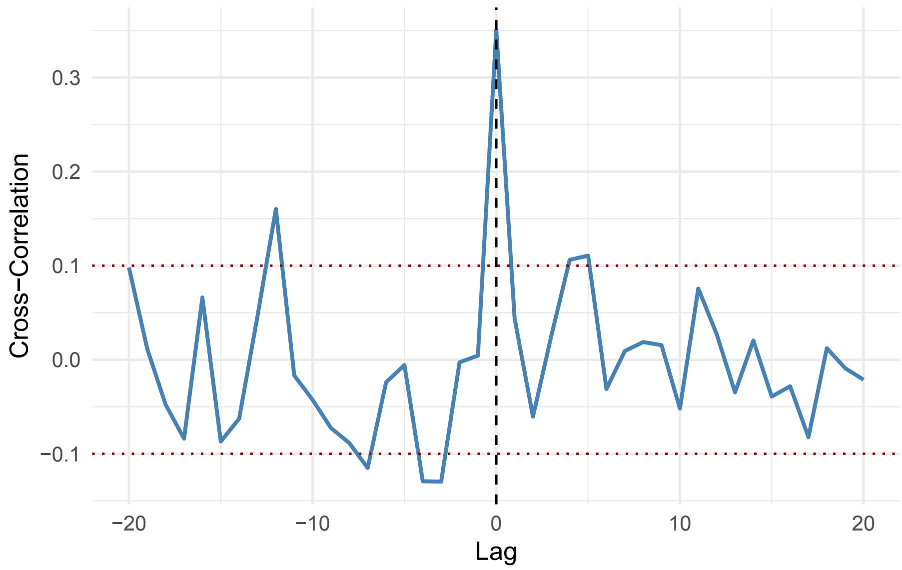

Figure 4.32: Lead lag correlation of the first difference of the TCI and VIX time series. Each lag corresponds to a 5-day period.

The correlation functions peaks at lag 0, showing a contemporaneous transmission of reinsurance connectedness. The network reacts within the same 5-day trading window as the board market, suggesting that the reinsurance sector is integrated into the global financial system, making the TCI a “Coincident Indicator” rather than a “lagging indicator”, capturing market shocks in real time. However, the highest correlation achieved ($\rho = 0.350$) implies that approximately only 12% of the weekly variation in reinsurance market connectedness is explained by equity market volatility. This highlights two different behaviors: while the reinsurance market reacts rapidly and almost simultaneously to market movements, the intensity of the reaction remains largely driven by idiosyncratic factors. The remaining variance is likely attributable to liability shocks, internal network rewiring, or lag effects. Thus, while the sector is well connected in terms of information transmission, it remains insulated in terms of magnitude, preventing a one-to-one coupling with the broad market sentiment.

Regarding other major fluctuations in the correlation structure, observed spikes at non-zero lags (like at lag -12 or at lag +5) fall within the bounds of statistical noise and lack significance. The absence of significant lead or lag effects confirms that the VIX cannot be used to predict future movement of TCI.

4.7 Conclusion

This chapter extended the static connectedness framework of Chapter 3 to a full dynamic analysis, estimating 349 rolling BVAR windows over the 2015-2025 sample and tracing the evolution of network topology, firm-level roles, and structural stability.

At the system level, the Total Connectedness Index reveals a structurally non-stationary market. The OLS-CUSUM test rejects mean stationarity, and the Bai-Perron algorithm identifies four structural breaks corresponding to five regimes: an initial high-connectivity phase, a stabilization period, the COVID-19 peak, a prolonged elevated phase, and a gradual normalization. The bimodal TCI distribution confirms that the reinsurance system does not revert to a single equilibrium but oscillates between two distinct states. This finding directly qualifies the static results of Chapter 3: a Total Connectedness value of 74.278 exceeds the rolling window mean of 68.2 by approximately 6 TCI units. Since the static BVAR was computed over the full 1,993 observation sample, this gap reflects the upward pull of the high-connectivity regime on the full sample average.

The topology analysis revealed an asymmetric network that contracts faster than it expands. The 2.2× ratio between the APL response coefficients to TCI increases(-0.003) and decreases (-0.007) shows that the network's structural compactness reacts far more sharply when connectedness recedes than when it builds: the architecture tightens gradually under accumulating stress but loosens abruptly once that stress passes. The modularity evidence reinforces this: community boundaries erode sharply during high-connectivity periods ($\rho(\text{Modularity}, \text{TCI}) = -0.78$), meaning that geographical or functional segmentation provides less containment precisely when it is most needed.

At the firm level, the persistent versus reactive centrality classification uncovers systemic institutions not detected through static BVAR. Chubb and Mapfre rank among the top eigenvector hubs across all rolling windows despite their absence from the static centrality ranking, while reactive firms shift between the network core and periphery during stress. The event analysis further shows that the stable core hierarchy of net transmitters is only partially preserved across shocks: COVID-19 and Ukraine conflict showed the highest significant rank correlation ($\rho = 0.47$), but the HIM hurricanes and the Trump tariffs produce no significant rank correlation with any other event. This shows that firm-level systemic importance within the reinsurance network is not a fixed structural property, but it is rather contingent on the transmission channel. Asset-side shocks synchronize the network and amplify the European transmitter hierarchy, while liability-side shocks fragment it along idiosyncratic exposure lines.

The VIX cross-correlation analysis supports this aspect. The TCI behaves as an indicator of broad-market stress, with correlation peaking at lag 0 ($\rho = 0.35$), but only $\approx 12\%$ of weekly TCI variation is explained by equity market sentiment, leaving the bulk attributable to network-internal dynamics and idiosyncratic reinsurance shocks. This seclusion limits the scope for using broad-market instruments as early-warning proxies for reinsurance connectedness.

These findings establish two components of systemic risk that the topological and connectedness analysis alone cannot resolve: the *magnitude* of the economic damage that a distress firm inflicts on the broader system and whether firms identified as persistent network hubs constitute genuine Systemically Important Financial Institutions in the regulatory sense. Chapter 5 addresses both questions by combining tail-risk measures with cross-validation against the centrality and net-strength rankings.

Chapter 5

Systemic Importance

5.1 Introduction and Motivation

In previous chapters we focused primarily on connectedness measures and topological analysis of how volatility propagates across reinsurers, both in a static and dynamic state. However, systemic risk emerges not only from interconnectedness of the system, but also from correlated leverage exposure. To better assess these two components of risk, we use measures widely used by financial regulators to assess the vulnerability of financial institutions to market downturns and observe the response of the market and systemically important institutions to shock scenarios that simulate extreme tail losses.

5.2 Systemic Risk Measures: a Framework

The network and connectedness analysis of the preceding chapters characterize how volatility disturbances propagate across the reinsurance system: which firms transmit, which absorb, and how the topology of those linkages evolves over time. What they cannot reveal is the magnitude of the economic damage a firm's distress inflicts on the broader financial system, nor the degree to which individual reinsurers would themselves be undercapitalized under a severe and prolonged market decline. The analysis can be divided into two complementary dimensions of systemic risk.

The first dimension addresses *systemic contribution*: the marginal increase in system-wide tail risk attributable to a specific institution entering distress. The Conditional Value at Risk (ΔCoVaR), introduced by Adrian and Brunnermeier (2016), captures this by measuring the change in the system's Value at Risk when a firm transitions from its median state to its distress state. A firm with a large $|\Delta\text{CoVaR}|$ amplifies system-wide tail risk when it fails; one with a small value remains relatively contained, even if it is highly connected. This outward perspective, from the firm towards the system, is complemented by an inward one: the Marginal Expected Shortfall (MES) of V. Acharya et al. (2010), which quantifies the expected loss a firm sustains on the days the market itself is in its worst tail. MES identifies institutions that are most vulnerable when the system is already under stress, and thus most likely to require recapitalization at precisely the moment it is most costly.

The second dimension concerns *capital adequacy under stress.* Building on the MES framework and the Capital Shortfall Approach of V. Acharya, Engle, and Richardson (2012), Brownlees and Engle (2016) developed the SRISK measure, which integrates the firm's market-based loss sensitivity with its balance-sheet leverage and regulatory capital requirements. SRISK reports the expected capital shortfall conditional on a severe market decline sustained over several months. Unlike MES, which is a return-based measure, SRISK is expressed in currency units and is therefore directly comparable across firms of different size and geographical location.

Together, ΔCoVaR, MES, and SRISK provide three complementary lenses on systemic importance: contribution to tail risk, vulnerability under market stress, and expected capital shortfall. To these market-based measures we add a fourth, structural perspective through the Generalized Impulse Response Functions (GIRF) derived from the BVAR model estimated in Chapter 2. Rather than conditioning on market returns, the GIRF framework traces how a one-standard-deviation shock to a specific firm's volatility propagates through the network over a defined horizon, yielding the Systemic Impact Index (SII) as a summary of each firm's total contagion capacity.

These four measures map different economic directions: the ΔCoVaR runs from the firm to the system, capturing the systemic contribution of each firm; MES and LRMES run from the system to the firm, capturing the vulnerability of each firm compared to the whole system; SRISK reads the balance sheet data and express capital adequacy in currency units; the SII runs from the firm into the network, summarizing contagion capacity.

5.3 Conditional Value at Risk

5.3.1 Methodology

CoVaR measures the Value at Risk of the financial system conditional on an institution being in distress, while the ΔCoVaR isolates each institution's marginal contribution to systemic risk, allowing identification of each reinsurer's marginal contribution to the entire system's risk when it enters distress.

Mathematically, the CoVaR measure of firm q can be expressed as:

$$CoVaR_q^{sys|i} = VaR_q(r^{sys}|r^i = VaR_q(r^i)) \tag{5.1}$$

where r^{sys} represents the system returns conditional on r^i, the returns of the firm in distress, and the market state variables M_t. The ΔCoVaR is then computed as

$$\Delta CoVaR_q^{sys|i} = CoVaR_q^{sys|dis} - CoVaR_q^{sys|med} \tag{5.2}$$

Firms with higher ΔCoVaR are those whose distress destabilizes the system, while firms with low ΔCoVaR cause less material distress to the system, even if their net-strength and connectedness measures are high.

For the empirical analysis, firms' returns were tested against the MSCI World Index ETF (Ticker: IWDA.AS) to test the distress effects of reinsurance companies on the broader market. The choice of the system proxy is consequential for ΔCoVaR. The MSCI World Index represents the aggregate global equity market, much larger than the actual reinsurance sector. Measuring each reinsurer's contribution to tail risk therefore produces conservative $|\Delta\text{CoVaR}|$ estimates: a firm central to the reinsurance network but small relative to global equity markets will register low values against the MSCI even if its distress is effective in the reinsurance sector.

The macro state vector comprised 6 economic variables:

- The volatility index (VIX) retrieved from Federal Reserve Economic Data (FRED) with daily frequencies over the sample period considered, to account for log-volatility daily changes
- The 10 years and 3 months Treasury Bond rates spread, to account for long-term and short-term differences in the risk free rates during the sample period, with data retrieved from the FRED database with daily frequency
- The credit spread between bonds classified at "baa" level against the 10 years treasury bond risk free rate, to account for variations in the cost of risky credit. The data was gathered from the FRED database with daily frequency
- The short term 3 month risk free rate, retrieved using daily data from the FRED database
- The daily market returns using the MSCI World Index as the market proxy
- The real estate returns, calculated against the Vanguard Real Estate ETF (VNQ), to account for illiquid assets

The returns were fitted on the 0.5 and 0.05 quantiles, to capture the median state and the tail distress scenario of the return distribution.

5.3.2 Results

The results of the ΔCoVaR analysis show high absolute values for European companies, as it can be seen in Figure 5.1. The significance of the estimation was tested using Huber sandwich (NID) standard errors, with 22 of the 27 firms retaining statistically significant $\hat{\beta}_i$ at 5% significance level. The full ΔCoVaR values with their significance level is reported in Table D.4 in the Section D.2. European companies appear to contribute more to the destabilization of the system, as their distress scenario would worsen the tail risk of the entire reinsurance system.

The ΔCoVaR analysis reveals a scenario that resembles the one found in Chapter 3 static analysis for the first companies, but with a divergent picture for the remaining ones, with bottom positions occupied primarily by American companies. The relatively low absolute values can point to stable and contained effects of the reinsurance companies over both the adjacent and related market, and more remote and disconnected environments.

Figure 5.2 reports the evolution of the aggregated average ΔCoVaR values of all the reinsurance companies during the whole sample, with bands representing the upper and lower 5th quantile values for comparison with the MSCI World Index. The vertical dashed lines contextualize the measure against the events considered in Chapter 4. The values remain similar through the whole sample, with later measures having a smaller range (quantile values are closer) compared to the measures at the start of the period. The most notable decline is in correspondence with the COVID pandemic.

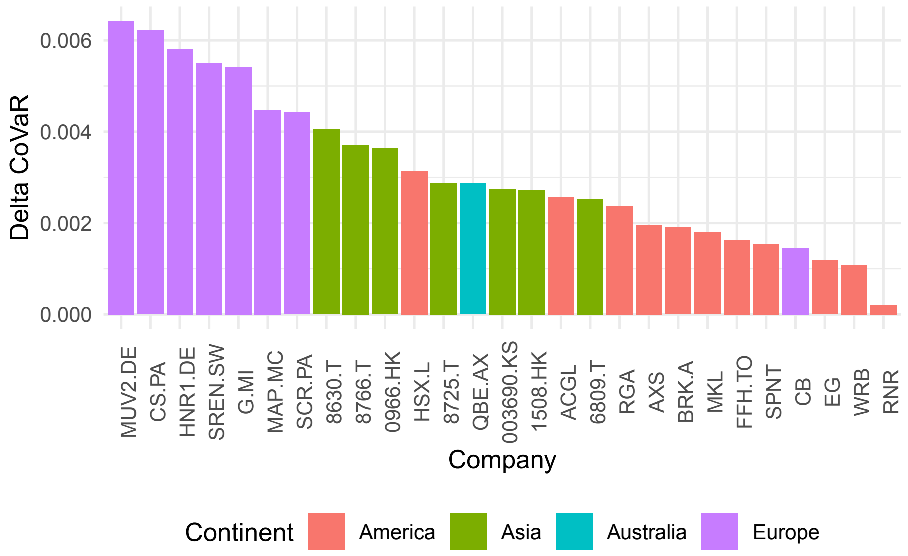

Figure 5.1: Delta CoVaR absolute values for each company in the sample.

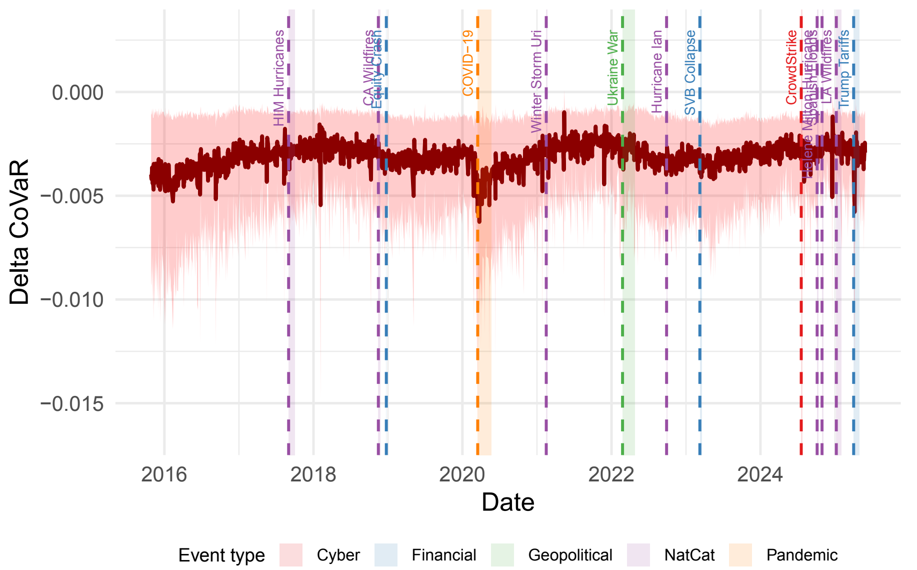

Figure 5.2: Delta CoVaR evolution over time with MSCI benchmark bands at 5th and 95th quantiles.

As a robustness check, the ΔCoVaR analysis was replicated using the iShares U.S. insurance Sector ETF (IAK) as an alternative system proxy, measuring each reinsurer's contribution to US insurance sector tail risk rather than the global equity market. The two results yield a Spearman rank correlation of $\rho = -0.433$, confirming that they capture almost inverse dynamics of systemic risk. Under the IAK specification, American reinsurers occupy the top position, displacing the European firms that dominate the MSCI-based ranking, with American companies' distress amplified within the US insurance sector through domestic co-movement. However, it is important to note that Chubb, Arch Capital Group, and W.R. Berkley Corp. (the top 3 companies in the replicated CoVaR analysis) are also contained in the IAK ETF, amplifying their effects over the considered sector. Also given this limitation, the MSCI-based proxy remains the primary analysis, as it measures spillover potential into the broader financial system. This near-inverse ranking offers a key result. Against the global market, the European reinsurers act as cross-sector contagion vectors, whereas against the US insurance sector, American firms dominate within-sector co-vulnerability. The two specifications answer different questions, and the divergence underscores that the label systemic is incomplete without specifying the reference system against which it is measured.

5.4 Marginal Expected Shortfall

5.4.1 Methodology

Marginal Expected Shortfall quantifies the expected firm return conditional on extreme negative market returns that exceed the left-tail 5% threshold. It is driven by three main factors: the market beta β, the correlation between the market returns and the firm returns, and the volatility of the firm returns.

$$MES_i = E[r_t^i | r_t^{mkt} < VaR_{0.05}(r^{mkt})] \tag{5.3}$$

As market conditions vary significantly and rapidly, static MES provides limited informational value in a dynamic setting. To grasp the dynamic behavior of the market and its risk, we estimated the dynamic MES through a GJR-GARCH DCC model (Engle 2002) to compute the dynamic conditional correlation of the market and the returns of the firm in a distress scenario. We assume that the standardized firm return $z_{i,t}$ can be decomposed into a component driven by the market innovation $z_{m,t}$ and an idiosyncratic innovation $\xi_{i,t}$, such that $z_{i,t} = \rho_{i,t} z_{m,t} + \sqrt{1-\rho_{i,t}^2}\xi_{i,t}$.

$$\widehat{\mathrm{MES}}_i = \sigma_{i,t}\left(\rho_{i,mkt,t} \cdot E[Z_{mkt} \mid Z_{mkt} < C] + \sqrt{(1-\rho_{i,mkt,t}^2)} \cdot E[\xi_i \mid Z_{mkt} < C]\right)$$

where C is the threshold corresponding to the market's 5% quantile. In this framework, the first expectation captures the systematic tail component, while the second expectation captures any residual tail dependence not accounted for by the DCC correlation.

Since the static MES captures only single-day dynamics, to include long-lasting crisis periods the Long Run Marginal Expected Shortfall (LRMES) was introduced, extending the MES concept to multi-month horizon without the limitations of the single-day sensitivity. It can be expressed mathematically as

$$LRMES \approx 1 - \exp\big(\log(1-d) \times \beta_{i,t}\big) \tag{5.4}$$

where d is the fraction of market decline and β is the time-varying market beta of firm i. It is estimated from the GJR-GARCH DCC model as $\beta_{i,t} = \rho_{i,t} \times \frac{\sigma_{i,t}}{\sigma_{m,t}}$, where $\sigma_{i,t}$ and $\sigma_{m,t}$ are the GJR-GARCH conditional volatilities and $\rho_{i,t}$ is the DCC conditional correlation. The time-varying beta is evaluated at each balance-sheet date using the DCC-estimated $\beta_{i,t}$, but it is treated as constant over the 126-day horizon within the closed-form approximation shown in Equation 5.4.

The combination of the GJR-GARCH with the DCC model allows to capture the evolution of volatility and correlations over time, encoding the typical behavior of financial markets, such as volatility clusters, asymmetrical impact of bad news, and correlation spikes during crises.

The GJR-GARCH(1,1) and DCC models and their relative coefficients are presented in more detail in Section D.3.

However, some important approximations are worth noting. The closed-form approximation assumes log-normally distributed returns and that daily shocks are independent across time. It also implies constant proportional leverage over the horizon, as the formula does not account for endogenous balance-sheet adjustments. The use of the same time-varying beta formula for all reinsurers introduces, moreover, a time-zone bias: Asian companies reflect the previous day's US closure, creating potential spurious intraday correlations. These limitations might result in an underestimation of true crisis losses and for a more detailed and regulatory framework they should be addressed in careful and precise manner.

5.4.2 Results

Figure 5.3 provides a graphical representation of the average MES value of the reinsurance companies, with dispersion around the 5% and 95% quantiles. In particular, the highest losses are in correspondence with the COVID-19 pandemic, with average expected losses greater than 10%. The widening dispersion accompanying the MES declines suggests that it is not uniform across the firms, indicating that losses may be concentrated in a subset of companies, not spreading further. This concentration is consistent with the asset-side idiosyncratic split identified in Chapter 4, where the sector as a whole is more resilient than its most-exposed members, whose losses deepen disproportionately as the market enters its tail. Other events, even related to the reinsurance sector, do not seem to cause spikes in the system average MES, showing a greater resilience of the sector to market declines. Notable spikes are at the end of the sample period, coinciding with the "Liberation day" tariffs, and one in the second half of 2016, not coinciding with any detected event.

Figure 5.4 shows the 5 companies with the highest average MES values. Reinsurance Group of America, represented by the purple line, is consistently above other companies, with sharper and more frequent declines, possibly indicating greater sensitivity to market downturns.

High correlation was noted especially between Generali and AXA ($\rho > 0.8$), indicating greater co-movements and similar loss responses by these reinsurers. It is worth noting that all three companies operate also in other financial sectors, as

Table 5.1: MES values averaged over the whole observations. The values are reported in ascending order: first values (greater absolute values) mean higher expected losses in bad market days.

Ticker	average MES
SCR.PA	-0.024
CS.PA	-0.021
G.MI	-0.021
MUV2.DE	-0.020
RGA	-0.018
SREN.SW	-0.018
HNR1.DE	-0.018
MAP.MC	-0.017
0966.HK	-0.016
HSX.L	-0.014
SPNT	-0.013
8630.T	-0.013
8766.T	-0.013
8725.T	-0.012
ACGL	-0.012
MKL	-0.011
AXS	-0.010
EG	-0.010
BRK.A	-0.010
1508.HK	-0.009
CB	-0.009
6809.T	-0.009
FFH.TO	-0.009
RNR	-0.008
WRB	-0.008
QBE.AX	-0.006
003690.KS	-0.005

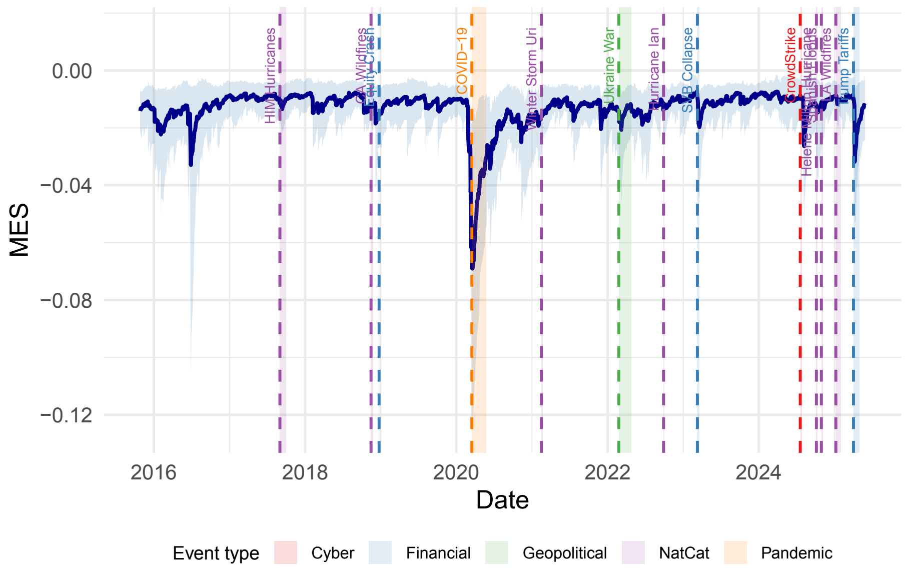

Figure 5.3: Aggregate average Marginal Expected Shortfall evolution over the sample period.

Generali and AXA undertake also primary insurance services, possibly explaining the higher correlation in the expected loss. This is a reminder that reinsurance companies span heterogeneous business models, and that the most MES exposed firms tend to be the most multi-line and asset-heavy companies, reinforcing the reading of the systemic risk driven by asset-side channels.

Figure 5.5 instead reports the average values of the aggregated MES by continent. Clearly, European companies account for higher expected losses, confirming the pattern already seen in Chapter 3 and 4. American companies highly correlate with European ones ($\rho \approx 0.90$), while Asian reinsurers tend to correlate less with European and American markets ($\rho \approx 0.56$ and $\rho \approx 0.46$ respectively). From the line chart, it is possible to see the tendency of Asian firms to recover more rapidly from market losses: the declines are more contained (smaller absolute value) and rapidly return closer to zero, indicating that the expected losses of these firms are smaller and they last for a more limited period. Economically, Asian reinsurers appear more insulated from global market turmoil, partly due to the domestic orientation and a distinct underwriting cycle, and partly through the time-zone artifact that may inflate their measured correlation. In either way, the pattern reinforces their peripheral role in the system.

As done in Section 5.3, the MES specification was tested against the US insurance sector proxy (IAK), and here as well the geographic pattern inverts: American firms record the largest expected losses on days when the US insurance sector is in its left tail, while European and Asian firms rank lower. The cross-specification Spearman rank correlation is $\rho = 0.156$, reflecting the different answer of the MES in the MSCI and IAK based systems. The MSCI results are consistent with the greater integration into international capital markets and larger equity float. The IAK result, despite the high global interconnectedness and the low community separation detected (recall the low modularity score,

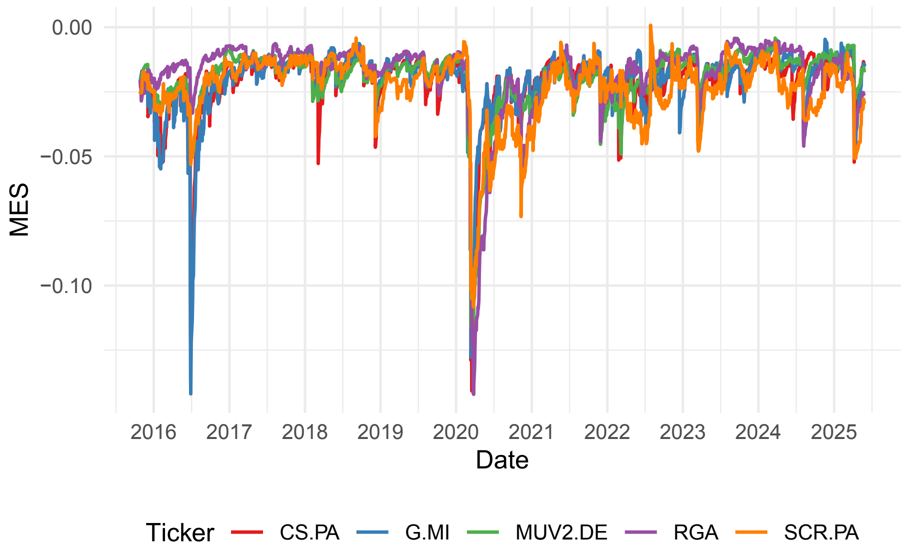

Figure 5.4: Aggregate Marginal Expected Shortfall evolution over the sample period for the 5 companies with the highest values.

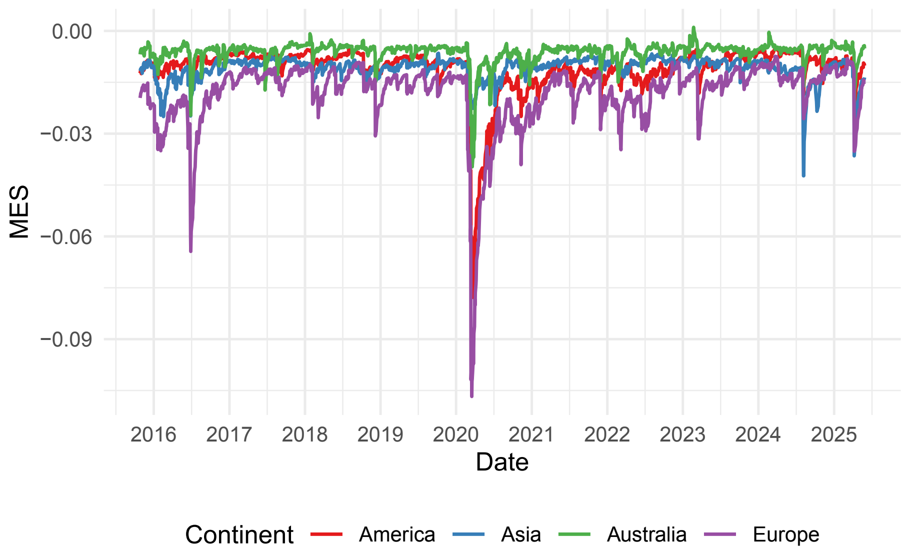

Figure 5.5: Aggregate MES based on the continent of the reinsurer.

$Q = 0.162$), still shows domestic-market influence.

5.5 SRISK

SRISK extends the MES frameworks by incorporating leverage considerations and capital requirements theory. This measure combines market-based risk measures with balance sheet data, based on the core idea of Capital Shortfall Approach introduced by V. Acharya, Engle, and Richardson (2012), in which a new method was proposed to estimate the capital required by a financial institution to survive another financial crisis.

The SRISK assesses the capital structure of reinsurance companies in stress scenarios, combining leverage, market sensitivity (LRMES) and regulatory capital requirements. It accounts for the fact that, in a crisis scenario, a firm's ability to raise funds is typically hampered by adverse market conditions. Mathematically, it is defined as follows:

$$SRISK_{i,t} = k \times \text{DEBT}_{i,t} - (1-k) \times \text{EQUITY}_{i,t} \times (1 - \text{LRMES}_{i,t})$$

where k is the capital prudential ratio that determines the minimum equity that a firm must hold as a fraction of total assets. Following the benchmark prudential ratios adopted in the systemic risk literature (Brownlees and Engle 2016), a capital ratio of $k = 8\%$ is applied uniformly for all reinsurers.

The DEBT values represent the financial leverage and were extracted from the balance sheet data available on Yahoo Finance for the financial years 2024, 2023, and 2022, under the line item 'Total Debt'. The $EQUITY$ was then calculated by extracting the 'Ordinary Shares Number' for each fiscal year and multiplying it by the closing price at the date at which the balance sheet was published. The values are then reported in USD. As some companies have different fiscal year-end dates, the price and exchange rate were calculated at the indicated closing date on the balance sheet, but the overall calculation was carried out considering the end of the year as the standard closing date for the fiscal year. In addition, different accounting standards were not considered. However, these limitations are acceptable for the purpose of assessing the framework and are left for more detailed future studies.

Under the data we extracted, no firm records positive SRISK under any of the severity scenarios tested (small, medium, and large severity, with $d = 0.20, 0.40, 0.60$ respectively) or fiscal year. Aggregate capital shortfall equals zero across all specifications. This result reflects a structural feature of reinsurance balance sheets: firms fund operations primarily through float, not financial debt, making financial leverage modest relative to market capitalization. SRISK was designed for bank-like leverage structures where financial debt is much greater than equity. Applied to reinsurance companies, the measure confirms capital adequacy rather than identifying shortfall. This result complements the ones reported earlier: European reinsurers are important contagion vectors, with high ΔCoVaR and MES values, but not as leveraged counterparties facing capital shortfall. This implies that being a transmitter of systemic risk and being a fragile, under-capitalized counterparty are distinct, and that the European core identified in the previous chapters exhibits the first without the second.

Two qualifications sharpen rather than overturn this reading. The balance-sheet sample spans only three fiscal years (2022–24), a moderate-connectedness window, so capital-adequacy risk during the high-TCI regimes of Chapter Chapter 4 is plausibly understated: zero shortfall now does not guarantee zero shortfall in a high-regime crisis. Moreover, the uniform $k = 8\%$ prudential ratio is a banking benchmark, whereas insurance capital regimes such as Solvency II operate differently; part of the zero-shortfall result thus reflects a model–instrument mismatch, which strengthens rather than weakens the conclusion that SRISK is not the natural capital tool for float-funded insurers.

5.6 Descriptive Shock Propagation Scenario

The second part of systemic risk assessment focused on the stress testing through the use of counter factual shock scenarios. The BVAR model estimated in Chapter 2 is used to estimate a Generalized Impulse Response Function (GIRF) in which one standard deviation shock is applied to firm k's volatility equation, tracing propagation through the network via IRF over a h-step horizon.

5.6.1 Methodology

Building on the BVAR model estimated in Chapter 2 and applied in the static and dynamic analyses of Chapters 3 and 4, we extend the same framework to a stress testing exercise through Generalized Impulse Response Functions (GIRF). Following Pesaran and Shin (1998), we adopt the GIRF framework rather than orthogonalized IRFs to preserve invariance to the ordering of variables in the system and provide consistency with the Generalized Forecast Error Variance Decomposition approach used throughout this work, as GIRFs are not structural, but rather ordering-invariant, reduced-form conditional forecasts with no causal identification strategy. Under this framework, a shock to firm i is not assumed to leave contemporaneous innovations in other equations at zero but instead, the expected response of the full system is conditioned on the observed covariance structure. Formally, denote the h-step moving-average coefficient matrices Φ_h, derived from the posterior-median lag matrices A_k via the recursion:

$$\Phi_0 = I_N, \qquad \Phi_h = \sum_{k=1}^{\min(h,p)} \Phi_{h-k} A_k, \quad h = 1, 2, \dots, H$$

The generalized impulse response to a one-standard-deviation shock originating in firm i is then:

$$\mathrm{GIRF}_i(h) = \Phi_h \cdot \frac{\Sigma \mathbf{e}_i}{\sqrt{\sigma_{ii}}} \tag{5.5}$$

where Σ is the posterior-median residual covariance matrix, $\mathbf{e}_i$ is the i-th unit selection vector, and σ_{ii} is the i-th diagonal element of Σ. The term $\Sigma \mathbf{e}_i / \sqrt{\sigma_{ii}}$ constitutes the generalized impulse vector: it scales the shock to one standard deviation of firm i's innovation while propagating the contemporaneous cross-equation correlations implied by Σ, rather than suppressing them as orthogonalized methods do. Both A_k and Σ are taken at their posterior median, providing a tractable point estimate consistent with the plug-in GFEVD computed in Chapter 2.

A forecast horizon of $H = 50$ trading days is selected, capturing the short-term propagation window over which volatility shocks are most persistent. To aggregate the full cross-firm propagation capacity of each firm into a single measure, we define the Systemic Impact Index (SII):

$$\mathrm{SII}_i = \sum_{j \neq i} \left| \sum_{h=1}^{H} \mathrm{GIRF}_{i \to j}(h) \right| \tag{5.6}$$

where $\mathrm{GIRF}_{i \to j}(h)$ denotes the j-th element of $\mathrm{GIRF}_i(h)$. The inner sum accumulates the response of firm j throughout the horizon; the absolute value prevents positive and negative responses from canceling; the outer sum aggregates across all $N - 1$ receiving firms. A firm with high SII transmits large cumulative volatility disturbances to the rest of the network following an idiosyncratic shock, regardless of direction.

5.6.2 Results

The results obtained from the application of the BVAR model impulse response highlighted the predominant contribution by European reinsurers, as Munich Re, Generali, AXA XL, and Hannover Rück, with Reinsurance Group of America being the only non-European to appear in the top 5 SII ranked companies. Table 5.2 reports the SII values for each reinsurer. The average SII values for European companies is 22.01, while for American companies this decreases to 16.94, and decreases even further for the Asian companies (9.83). This reflects the findings that have been highlighted throughout the entire work, with European firms contributing to volatility disturbance, affecting, with less magnitude, also peripheral nodes. Figure 5.6 and Figure 5.7 show the different propagation effects to the log-volatility for Munich Re and Taiping Reinsurance Co, respectively the first and last companies in the SII ranking. The color scale applied for the two heatmaps is the same, making the comparison between the two visually honest. As the dark red bar shows, the propagation of Munich Re lasts for longer period, while the effects of Chinese company tend to become null or even negative for further observations. Interestingly, the effects of Munich Re are more prominent across other European companies, like AXA, Mapfre, or Swiss Re, while the effects of the Chinese companies, although shorter in time, tend to span across multiple reinsurers with the similar magnitude. This structure highlights asymmetric risk propagation across the system, with risk flowing primarily from Europe to the other closer European companies, and to the other reinsurers with lower effects, but with more persistent effects, while Asian firms tend to act primarily as shock absorbers, as their effect is limited in terms of duration.

Geographical similarity was observed especially between European reinsurers and between Japanese reinsurers, as visible in Figure 5.8, Figure 5.9, and Figure 5.10 showing a stronger volatility disturbance, both in terms of duration and absolute value, while this behavior was not observed for American companies. This result recalls the observation of the pairwise connections in Table 3.1, where the explanation of Japanese reinsurers' volatility was derived primarily from other domestic companies.

Table 5.2: Systemic Impact Index rank.

Ticker	SII
MUV2.DE	25.382
G.MI	23.888
RGA	23.763
CS.PA	23.394
CB	22.839
WRB	22.754
HNR1.DE	22.723
MKL	22.435
MAP.MC	21.287
SREN.SW	20.949
BRK.A	18.427
EG	18.121
AXS	16.758
ACGL	16.551
RNR	16.253
SCR.PA	15.607
FFH.TO	13.653
8630.T	12.713
1508.HK	12.280
8725.T	11.684
8766.T	11.182
QBE.AX	10.099
HSX.L	10.058
003690.KS	9.331
SPNT	7.619
6809.T	6.462
0966.HK	5.171

5.7 Cross-Validation

To ensure the robustness of our findings, we conduct a cross validation between the metrics of Chapter 3 and the systemic risk indicators of this chapter. Recalling the methodology used in the static analysis, the Net-strength and Out-strength values were used as primary indicators of volatility disturbance of the system. The correlation used to compare the different metrics was the Spearman's rank correlation. The Net-strength exhibits a negative rank-correlation with ΔCoVaR ($\rho = -0.653$) and with the average MES values ($\rho = -0.486$), both significant at the 5% confidence level. Instead, positive correlation was observed between the Net-strength and the SII, at $\rho = 0.358$. The presence of significant correlation confirms that the reinsurers identified as primary explainer of log-volatility drift are also those whose distress generates the most significant marginal increase in the system's tail risk and are typically predicting the dynamic contagion potential during stress events. The negative sign of the ΔCoVaR and MES correlations is directionally consistent: high positive Net-strength (net variance

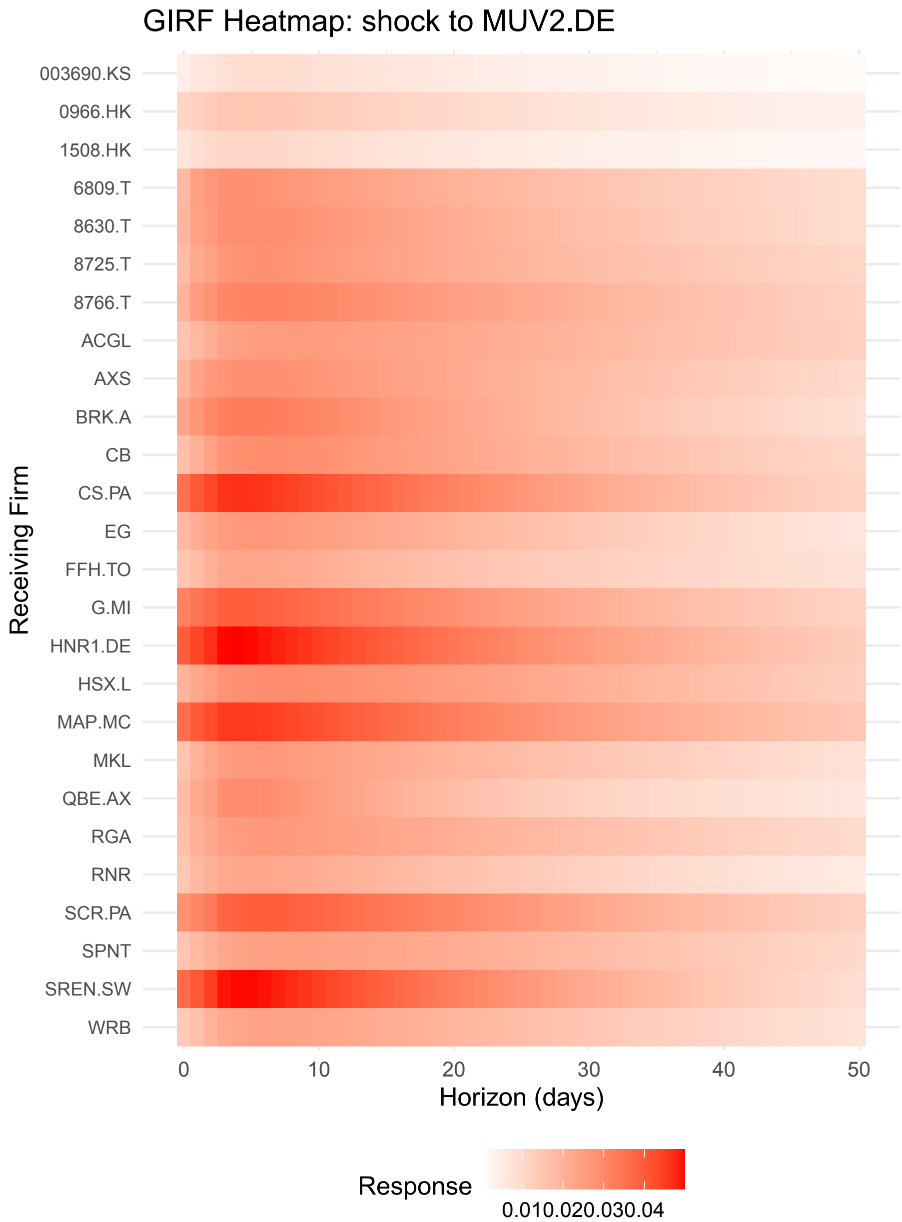

Figure 5.6: GIRF heatmap: volatility response of all firms to a shock to Munich Re (MUV2.DE, highest SII rank).

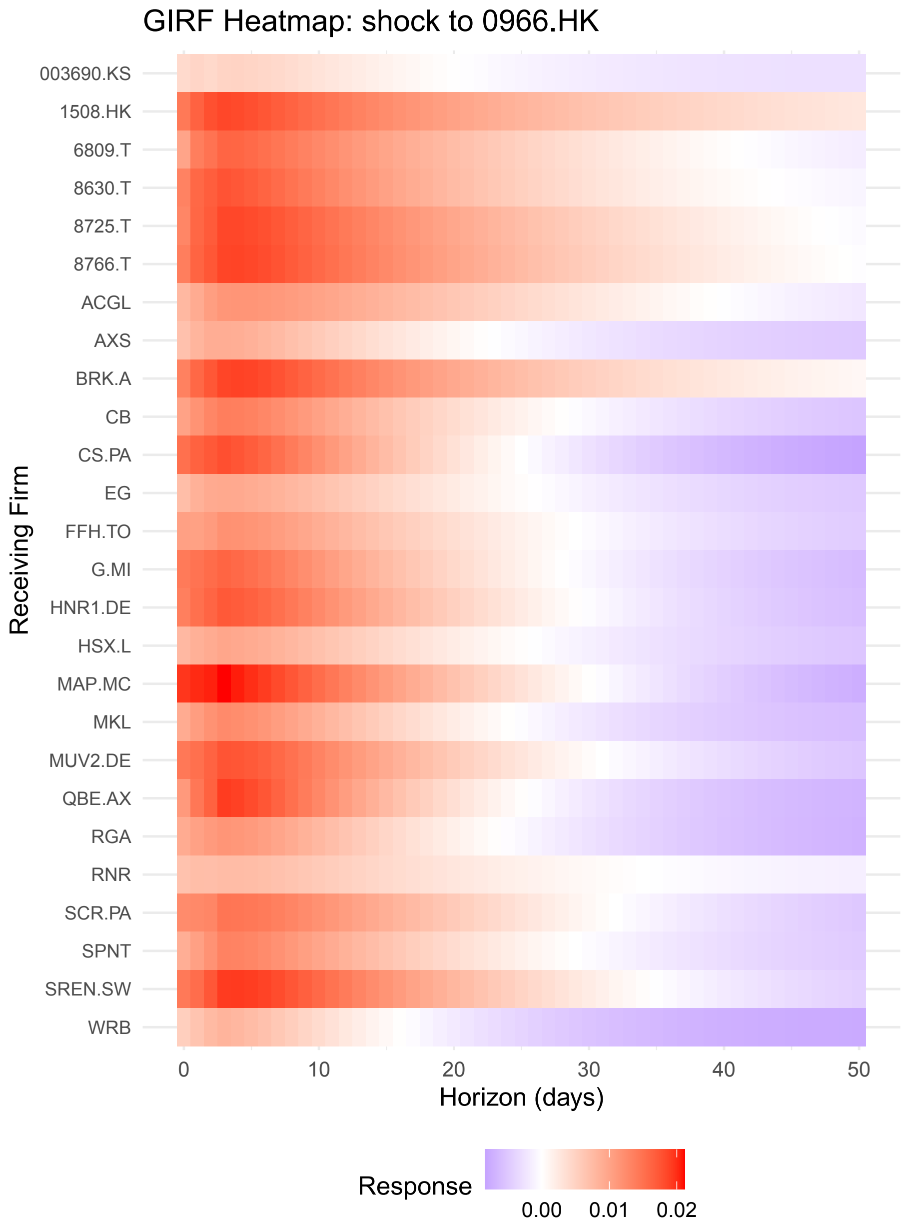

Figure 5.7: GIRF heatmap: volatility response of all firms to a shock to Taiping Re (0966.HK, lowest SII rank).

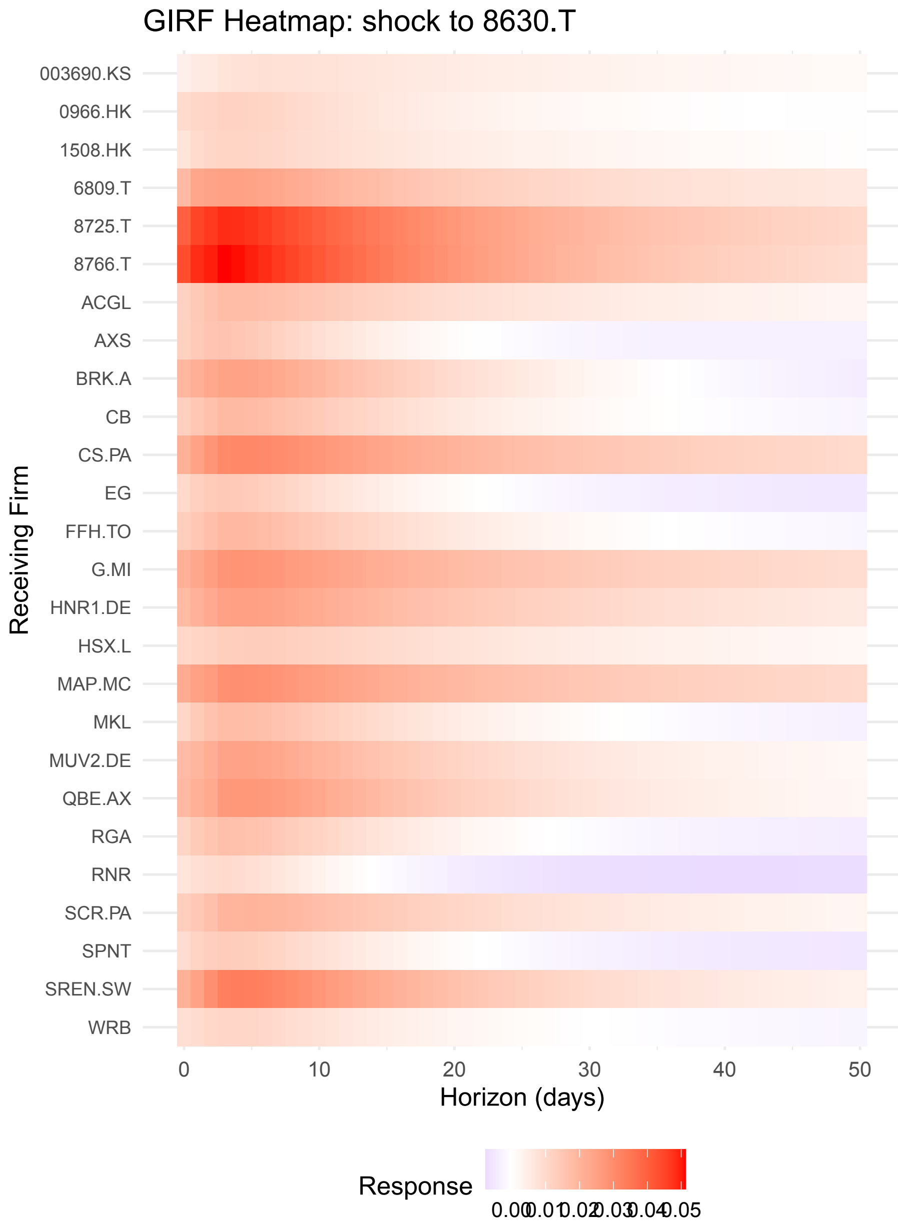

Figure 5.8: Japanese companies showing high volatility disturbance between each other - Sompo International Holdings, Ltd. (8630.T).

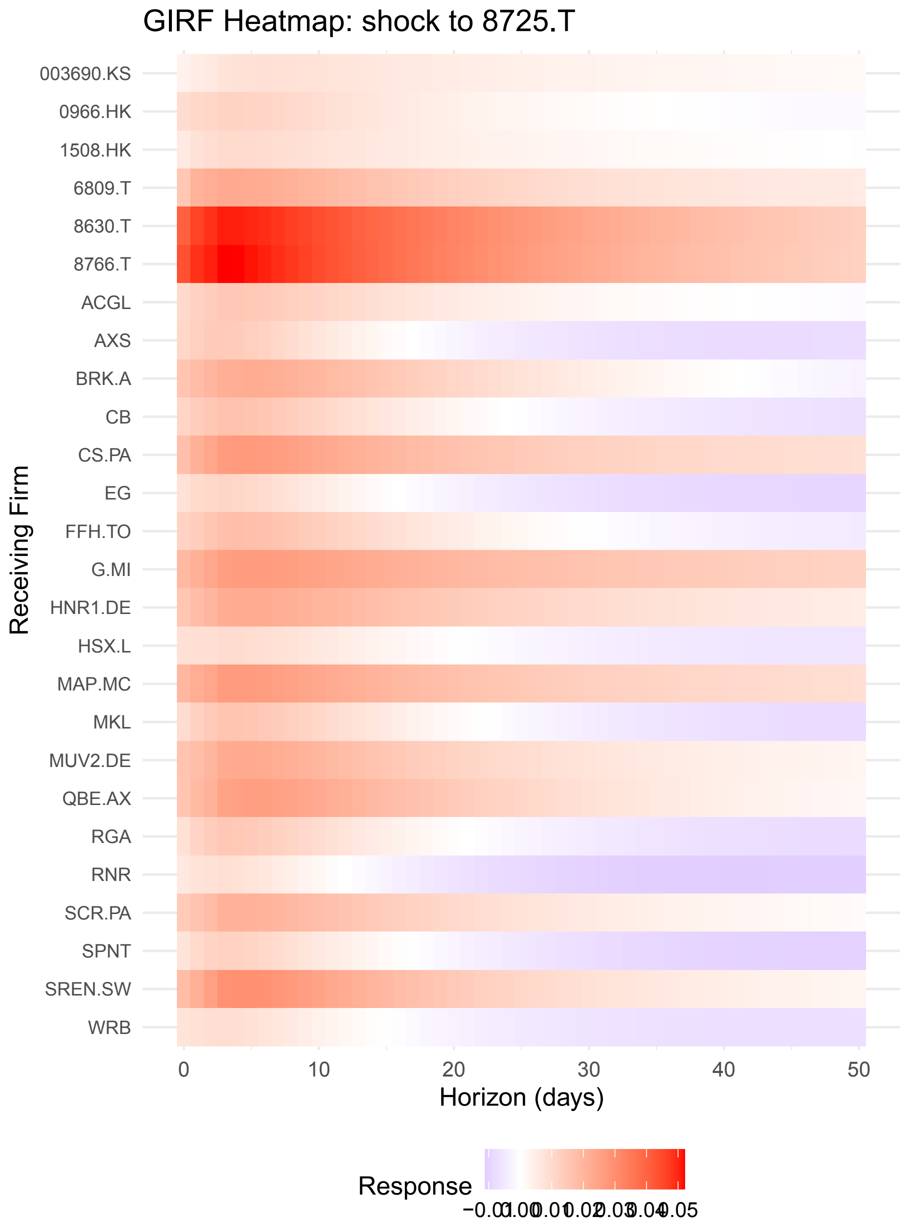

Figure 5.9: Japanese companies showing high volatility disturbance between each other - MS&AD Insurance Group Holding, Ltd. (8725.T)

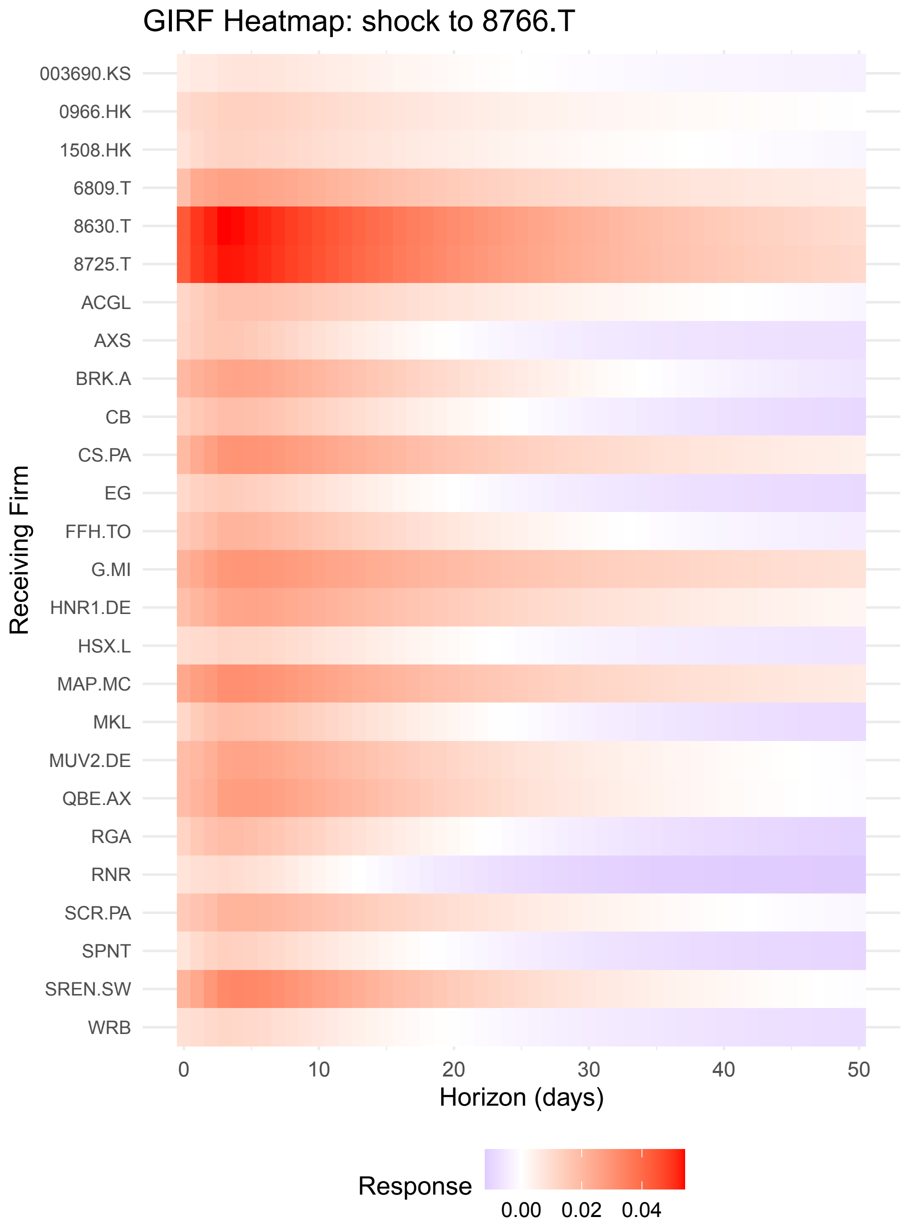

Figure 5.10: Japanese companies showing high volatility disturbance between each other - Tokio Marine & Nichido Fire Insurance Co., Ltd. (8766.T).

Table 5.3: Spearman rank correlation between Chapter 3 centrality measures and Chapter 5 systemic risk indicators.

Measure	vs CoVaR	vs MES	vs SII
Net-strength	-0.653	-0.486	0.358
In-strength	0.183	-0.039	0.681
Out-strength	-0.494	-0.558	0.702
Eigenvector	-0.577	-0.639	0.710
PageRank	0.560	0.305	0.367

transmitter) corresponds to large negative ΔCoVaR and MES (large increase in tail risk under distress), so the negative sign reflects agreement, not opposition, between the two frameworks.

We further assess the topological results of the static analysis against the systemic risk measures introduced in this chapter by examining the Out-strength and Eigenvector centrality relationships.

The correlation between the in- and out-strength and the ΔCoVaR was also addressed. Out-strength captures the directional component of a firm's influence without netting against received variance (In-strength). Despite removing the effects of the log-volatility variance received from the others, the correlation with ΔCoVaR is interestingly lower ($\rho = -0.494$, $p < 0.05$), indicating that the Net-strength serves as a stronger indicator for identifying companies that generate the highest increase in tail risk. Instead, the positive correlation with the in-strength ($\rho = 0.183$), shows that high receiving firms tend to contribute less to tail-risk. This shows that netting matters: a firm's net directional role identifies systemic contributors more sharply than gross transmission alone, since a firm that both transits and receives heavily is less systemically distinctive than a clean net exporter. The balance of flows is the better systemic signal.

We also examine the alignment between eigenvector centrality and PageRank values against the independent market-risk measures. The negative correlation between the eigenvector and both the ΔCoVaR and MES, shows that high eigenvector nodes tend to create higher tail risk, providing a similar interpretation to the net values, as its high correlation seen in Chapter 3 already suggested.

Regarding the SII, it is important to note that these measures are derived from the same BVAR model: eigenvector centrality is the principal left eigenvector of the GFEVD matrix, while the SII is the sum of absolute cumulative GIRFs, where GIRFs are constructed from the same lag matrices and covariance Σ that generated the GFEVD. The observed Spearman rank correlation ($\rho = 0.71$) thus reflects internal model consistency rather than independent empirical validation, confirming that the BVAR-based network representation is self-coherent, but is not providing out-of-sample evidence that the GFEVD centrality corresponds to real-world systemic importance. The independent cross-validation evidence consists of the ΔCoVaR and MES reported above.

5.8 Conclusion

This chapter examined the systemic importance of global reinsurers through four complementary lenses: the marginal contribution to market tail risk (ΔCoVaR), vulnerability under extreme market stress (MES and LRMES), expected capital shortfall in prolonged downturns (SRISK), and contagion capacity through the network (SII). Together, these measures extend the purely topological picture developed in Chapters 3 and 4 by introducing magnitude: not merely *who is connected*, but *how severely distress in one firm damages the system* and *whether the sector holds sufficient capital to absorb it*.

The ΔCoVaR and MES analyses converge on a consistent ranking: European reinsurers generate the largest marginal increases in system-wide tail risk and sustain the deepest expected losses when markets are in distress. This finding directly validates the network centrality results of Chapter 3, where the same firms occupied dominant positions in Out-strength and Eigenvector centrality. The cross-validation in Section 5.7 formalizes this correspondence: Net-strength is significantly rank-correlated with both ΔCoVaR ($\rho = -0.653$) and MES ($\rho = -0.486$), confirming that reinsurers identified as primary variance contributors through the GFEVD are also those whose distress generates the most material increase in tail risk.

Regarding SRISK, no firm in the sample records positive SRISK under any severity scenario or fiscal year examined. Aggregated capital shortfall equals zero across all specifications. This confirms that the sector's financial leverage is structurally modest relative to its market capitalization, and that the European contagion vectors identified through ΔCoVaR and MES are not simultaneously exposed to capital shortfall risk under standard metrics. This separation aligns with the international shift from the entity-based G-SII designation toward the activities-based Holistic Framework of the IAIS. The systemic concern raised by reinsurers is one of interconnectedness and asset-side spillover, not of bank-style solvency risk under conventional leverage metrics.

The BVAR-based stress testing reinforces the asymmetric structure of the network. Munich Re and the other European hubs generate persistent, broad-spectrum volatility disturbances that decay slowly, whereas Asian firms produce shorter-lived and more localized impulse responses. Japanese firms exhibit strong intra-cluster propagation, consistent with the domestic pairwise connections identified in Chapter 3. Asian firms, by contrast, appear to absorb rather than transmit global shocks.

Several limitations qualify these findings. The SRISK analysis covers only three fiscal years, and the balance-sheet sample coincides with a period of moderate TCI; capital adequacy risks during high-connectedness regimes identified in Chapter 4 are likely understated. The closed-form LRMES approximation assumes constant proportional leverage and log-normal returns, and the time-zone bias introduced by Asian companies' lagged price discovery may inflate their measured correlation with European markets. Finally, the GIRF framework is a reduced-form, ordering-invariant exercise: the propagation patterns it reveals reflect statistical co-movement, not structural transmission mechanisms.

Taken together, the evidence in this chapter supports a two-tier characterization of the global reinsurance system. A small group of European incumbents sits at the intersection of high topological centrality, elevated systemic contribution, and strong dynamic contagion potential. A second group, composed primarily by Asian reinsurers, occupies the network periphery and contributes little to tail-risk amplification, with capital structures that remain resilient under standard financial leverage benchmarks. This structure suggests that macro prudential oversight of reinsurance should prioritize European core for contagion monitoring, while balance-sheet surveillance across the sector should track the ratio of financial debt to market capitalization over the economic cycle, particularly during high-connectedness regimes where the current zero-shortfall finding may not hold. This chapter therefore provides a decoupling of the debate regarding the systemic properties of the reinsurance sector: global reinsurers are systemically important as interconnected, asset-side transmitters, but not as bank-like leveraged institutions prone to capital shortfall.

Chapter 6

Conclusion

This thesis investigated the systemic risk structure of the global reinsurance market by combining Bayesian econometric modeling, network topology analysis, and market-based tail-risk measures. The analysis covered 27 publicly listed reinsurers across four continents over the period October 2015 to May 2025, producing a comprehensive assessment that spans static and dynamic network characterization, firm-level systemic importance ranking, and cross-validation of complementary methodologies. This concluding chapter synthesizes the principal findings, discusses their implications for regulatory oversight, acknowledges the limitations of the framework, and identifies directions for future research.

6.1 Summary of Findings

The static network analysis of Chapter 3 established that the global reinsurance system is tightly interconnected, with a total connectedness index of $TC = 74.28$, a short average path length ($L = 0.39$), and low modularity ($Q = 0.162$). These properties indicate a globally integrated market rather than a collection of fragmented regional silos. The network exhibits disassortative mixing ($r = -0.116$) in the static estimate, meaning that highly connected European hubs link preferentially to less connected peripheral firms in Asia and the Americas. This topology implies that a shock originating at a European hub would not remain contained within the European core but would instead propagate rapidly toward the periphery. European reinsurers, in particular AXA XL, Munich Re, and Swiss Re, emerged as the dominant variance contributors, ranking highest across out-strength, eigenvector centrality, and net-strength. Asian and American reinsurers, by contrast, occupied net receiver positions, absorbing variance from the system rather than generating it.

The dynamic analysis of Chapter 4 qualified and extended these static results in several important aspects. The rolling-window BVAR estimation across 349 overlapping windows revealed that the total connectedness of the reinsurance network is structurally non-stationary. The OLS-CUSUM test formally rejected mean stationarity, and the Bai-Perron algorithm identified four structural breaks corresponding to five distinct regimes: an initial high-connectivity phase associated with the 2017 HIM hurricane season, a stabilization period, the COVID-19 peak representing the maximum connectedness observed in the sample, a prolonged elevated phase during the Russia-Ukraine conflict, and a gradual normalization. The static total connectedness value of 74.28 falls above the 75th percentile of the rolling TCI distribution, indicating that the static BVAR was skewed upwards by prolonged high-connectivity regimes.

The topology analysis uncovered an asymmetric network dynamic: the reinsurance system contracts approximately 2.2× faster than it expands. When TCI decreases, the average path length responds with a coefficient of -0.695, compared to -0.314 when TCI increases. This asymmetry implies that contagion channels form gradually during periods of increasing stress but dissolve rapidly once stress subsides. Community boundaries erode sharply during high-connectivity periods ($\rho(\text{Modularity}, \text{TCI}) = -0.78$), meaning that the slight geographical segmentation that might otherwise contain shocks provides less insulation precisely during the periods when it would be most valuable.

At the firm level, the persistent-versus-reactive centrality classification identified systemically important institutions that were invisible to the static analysis. Chubb and Mapfre ranked among the top eigenvector centrality hubs across all rolling windows despite their absence from the top positions in the static centrality ranking. Instead, firms such as RenaissanceRe and AXIS Capital, exhibited reactive behavior, rising temporarily to prominence during specific stress events before returning to peripheral positions. This finding carries the central economic message of the firm-level analysis, as the network does not simply rediscover the largest reinsurers. Some firms are central because they are large (Munich Re, Swiss Re), but others are far more systemically connected than their size or market share would imply, while the single largest firm in the sample, remains peripheral. Connectedness, albeit positively correlated with size, is therefore a distinct dimension of systemic importance.

The event analysis suggests that asset-side shocks (COVID-19, Trump tariffs) synchronize the network and amplify the European transmitter hierarchy, while liability-side shocks (HIM hurricanes) fragment it along idiosyncratic exposure lines, reducing co-movement rather than increasing it.

Chapter 5 moved beyond the purely topological perspective by introducing market-based systemic risk measures that quantify the economic magnitude of each firm's systemic contribution. The ΔCoVaR and MES analyses converged on the same ranking identified through the network analysis: European reinsurers generate the largest marginal increases in system-wide tail risk and sustain the deepest expected losses during market stress. The cross-validation formalized this correspondence quantitatively. The significant correlations between independents measures indicates that the topological hubs identified through the GFEVD framework are not merely statistical artifacts of the variance decomposition but correspond to firms whose distress would generate material tail-risk externalities on the broader system.

Note that the Eigenvector-SII correlation reflects internal model consistency rather than independent validation, as both quantities derive from the same BVAR posterior.

The SRISK analysis showed no firm in the sample with positive SRISK under any severity scenario or fiscal year examined. Firms fund operations primarily through premium float and equity capital, maintaining financial debt levels that are lower relative to their market capitalization.

The BVAR-based stress testing through Generalized Impulse Response Functions confirmed the asymmetric propagation structure. A one-standard-deviation shock to Munich Re produces persistent, broad-spectrum volatility disturbances that decay slowly across the system, while an equivalent shock to a peripheral Asian firm generates a response that is both shorter-lived and more geographically localized. Japanese firms exhibited strong intra-cluster propagation consistent with the domestic pairwise connections identified in the static analysis, reinforcing the finding that regional clustering persists even within a globally integrated network.

Taken together, these findings carry implications for how supervisory authorities might monitor and oversee the reinsurance sector.

6.1.1 Original Contribution

The most important finding and the central contribution of this thesis is not any single measure but the demonstration that they agree. Network centrality, derived purely from the GFEVD decomposition of the equity-volatility co-movement, identifies the same European core that the two methodologically independent, market-based tail-risk measures flag as systemically important. ΔCoVaR, estimated by quantile regression on macro state variables, and MES, estimated through a GJR-GARCH-DCC system, each rest on a different statistical foundation and a different notion of "systemic", yet they converge on the same ranking. The cross-validation makes this correspondence quantitative: net-strength correlates significantly with ΔCoVaR ($\rho = -0.653$) and with average MES ($\rho = -0.486$). Because these measures share no common estimator with the GFEVD network, their agreement cannot be an artifact of the variance decomposition; the BVAR-based Systemic Impact Index, by contrast, derives from the same posterior and therefore corroborates the model's internal consistency rather than providing independent validation. The implication is that the centrality rankings obtained from the network analysis are not abstract topological descriptors, but they capture an economically meaningful dimension of systemic importance, and a firm's position in the volatility network is informative about the tail-risk externality its distress would impose on the system.

A second contribution is the dissociation this thesis documents between *systemic centrality* and *financial fragility*. The firms that transmit the most risk through the network are not the firms that look fragile on their balance sheet. No reinsurer in the sample records positive SRISK under any severity scenario or fiscal year examined: the sector is float-funded, low-leverage, and well-capitalized, and a conventional capital-shortfall lens would therefore flag no systemic concern at all. Yet the same firms sit at the center of a dense, fast-propagating volatility network and transmit shocks broadly and persistently. Systemic risk in reinsurance is thus a property of interconnection and shock propagation, not of leverage or capital insufficiency: the opposite of the bank-centric paradigm in which systemic risk and under-capitalization largely coincide. This distinction reframes what supervisors should monitor in the sector: the contagion potential of the most connected transmitters, rather than the solvency of the most leveraged firms. It is the reason a connectedness aware, position-based criterion adds information that a capital-based or size-based designation cannot recover, and it motivates the supervisory recommendations developed below.

6.2 Regulatory and Policy Implications

The empirical architecture outlined above touches on some critical questions surrounding the prudential supervision of the reinsurance industry. One such question is whether large reinsurers need to be supervised in addition to the current practice of individual solvency evaluation, which defines the basis for supervision in the industry. The analysis presented in Chapter 3, Chapter 4, and Chapter 5 clearly points to a positive answer to this question. The market-based indicators of tail risk suggest that a relatively small group of incumbent companies from Europe – Munich Re, Swiss Re, Hannover Re, and SCOR – is responsible for the largest marginal contribution to aggregate tail risk, being the leading firms in terms of ΔCoVaR and Marginal Expected Shortfall, while cross-validation proves that they are also the most important nodes in terms of GFEVD networks. The stress testing via BVAR model makes the same point more explicitly: while the shock of one standard deviation to Munich Re is transmitted throughout the network and dissipated slowly, the same shock to the Asian peripheral firm is transmitted locally and dissipates quickly. A solely Solvency II based supervision that focuses on reinsurers individually is structurally blind to this externality. The systemic cost of distress at a hub is borne by the whole network, not by the firm's balance sheet, and a solo-capital standard does not price it. The findings therefore support a macroprudential overlay that complements rather than replaces the existing solo regime.

Another question is whether the international designation architecture is well suited to capture also the kind of systemic relevance that was documented here. This comparison must be drawn carefully, as the sample analyzed was purposely composed of mainly reinsurers, while the G-SII regime operated by the FSB and IAIS between 2013 and 2016 ranked across the entire insurance sector. What the comparison can actually address is the nature of the designation criterion. The G-SII indicator framework was dominated by size, non-traditional-non-insurance activity, and global cross-jurisdictional presence, whereas these results show that, within the reinsurance sector, a connectedness-based criterion re-ranks firms and elevates carriers whose systemic relevance derives from their position in the network rather than from their balance-sheet composition. The comparison between persistent and reactive firms performed in Chapter 4 identifies firms such as Mapfre and Hannover Re as persistent hubs despite their smaller market capitalization. A criterion based on size alone would systemically underweight such position-systemic firms. The above point on methodology finds further corroboration in an important fact relating to the scope of the regulatory classification: the two largest reinsurers of the world, namely, Munich Re and Swiss Re, who occupy dominant positions as transmitters of variance in the current study, have never been classified as G-SIIs, despite several primary insurers having made the cut. This is presented only as an example from a qualitative perspective of how limited the scope of entity-based classification has been within the reinsurance channel, and not as any kind of comparison between the current sample and the G-SII universe in any sense of the term. The bigger trend is certainly supportive of the above interpretation: In November 2017, the IAIS put on hold the annual publication of the G-SII list, and along with the FSB, substituted the entity-based designation system with the activities-based Holistic Framework instituted in November 2019 and enforced from 2020 (International Association of Insurance Supervisors 2019). The annual Global Monitoring

Exercise under the Holistic Framework already monitors interconnectedness directly, as one of its core systemic-risk indicator categories, and is well suited to detecting system-wide build-ups. Two features nonetheless leave room for the present analysis: its interconnectedness indicators are exposure and balance-sheet based, whereas the network constructed here is a market-implied measure that requires no contractual disclosure; and the framework operates primarily at the aggregate sector level, having suspended firm-level designation in 2017. The above results therefore complement rather than duplicate it, by supplying a connectedness-aware, firm-level dimension of the kind the 2025 recovery-and-resolution scope consultation has begun to reintroduce.

This direction is borne out by the most recent supervisory developments: in November 2025 the FSB reaffirmed its discontinuation of entity-based G-SII designation while simultaneously consulting on the scope of insurers subject to recovery and resolution planning under the Key Attributes, a firm-level selection that rests explicitly on six criteria including interconnectedness (Financial Stability Board 2025b, 2025a). The framework thus retains its activities-based backbone while reintroducing a firm-level layer keyed in part to connectedness, precisely the supervisory complement that the present analysis quantifies.

The third question concerns the systemic consequence of the sector's concentration, and answering it requires distinguishing two notions of network risk that the topology of Chapter 3 separates cleanly. The network does not present a single-point-of-failure architecture in the structural sense familiar from banking. Betweenness centrality is low across all firms, so no node functions as an indispensable bridge whose removal would fragment the system. Contagion, in fact, travels through many redundant paths rather than through a unique gateway, making the reinsurance network robust to the loss of connectivity. What it is exposed to is amplification: the low modularity, short average path length, and disassortative mixing describe a hub-and-spoke market in which out-strength is concentrated in a few European hubs. The periphery here is peripheral in strength terms, meaning that the firms are still directly wired, or easily accessible from the hubs, rather than insulated from connection. This proximity is what makes amplification fast, creating a directional and asymmetric risk. As pointed out in Chapter 3, the failure of peripheral net receivers are dampened by the system with little systemic effect, whereas the distress of a high out-strength hub is transmitted efficiently and more persistently across the network. The Generalized Impulse Response evidence of Chapter 5 confirms this asymmetry: a shock to Munich Re (high net-strength) propagates broadly and decays slowly, while an equivalent shock to a peripheral firm is localized and ephemeral. On this reading, concentration raises systemic risk, placing the capacity of the network in the hands of a handful of high-strength firms.

A first qualification follows from the dynamic analysis of Chapter 4, in which these roles are not fixed. Net-strength and assortativity vary across regimes, and while a persistent core (Swiss Re, Chubb, Mapfre) retains its centrality through time, reactive firms rise to transmitter prominence only during specific stress events. The identity of the most efficient transmitters is therefore partly state-dependent, which is itself an argument for continuous monitoring rather than a one-off designation.

Another condition, established in Chapter 5, is that this is a risk of interconnectedness and asset-side spillover rather than of bank-style capital shortfall: no firm in the sample records positive SRISK under any scenario, consistent with the sector's float-funded, low-leverage balance sheets. The two findings are complementary: the reinsurance core is systemically important as an amplifying transmitter, not as a leveraged solvency risk, and they jointly suggest that supervisory concern should be directed at the contagion potential of the core transmitters rather than at conventional capital adequacy or at any single structural choke point.

These conclusions translate into a set of indicators that supervisory authorities could monitor. The Total Connectedness Index is the natural indicator, coinciding with broad-market stress, and introducing the regime classification of Chapter 4, re-framing the supervisory question from *how connected is the system now* to *which regime is the system in and is it transitioning*.

Because the TCI is bimodal and structurally non-stationary, the Bai-Perron and the OLS-CUSUM tests are used as real-time regime-detection tools, flagging transitions into the high-connectivity state that coincided with the COVID-19 or the market crash of the Trump's tariffs episodes. At the firm level, the net-strength and eigenvector centrality rankings, together with the persistent hub list, provide a position-based complement to size and activity-based designations.

Finally, the asymmetric topology dynamic offers a distinctive early-warning logic: since the network contracts gradually under accumulating stress phases but relaxes roughly 2.2× faster once stress passes (*slow tightening, fast relaxation*), a slow, sustained rise in connectedness is a more reliable signal of building systemic fragility than a sudden move, weighting gradual upward drift in the TCI more heavily than transient spikes.

One caution carries over from the dynamic analysis: while the TCI is coincident with market stress, it is not predicted by the VIX, so broad-market instruments cannot serve as leading proxies for reinsurance connectedness, and must be monitored in their own right.

These regulatory inferences rest on a network constructed from equity price co-movement rather than from contractual risk transfer, and they should be read alongside the caveat developed in Section 6.3 below: the GFEVD connectedness captures statistical association, not the bilateral retrocession exposures that would map systemic linkages directly. The policy direction the evidence points toward is therefore robust as a supervisory orientation, but the precise identification of which nodes to designate would be strengthened materially by access to the contractual-exposure data discussed below.

6.3 Limitations

Several limitations qualify the findings presented in this thesis.

The use of stock price volatility as a proxy for underlying risk transfer creates an inherent gap between what the model captures (market sentiment and co-movement in equity returns) and what it aims to measure (the actual reinsurance exposures that create systemic linkages). The ideal foundation for this analysis would have been bilateral retrocession data — the contractual risk-transfer positions reinsurers hold against one another — which map systemic linkages directly. Such data, however, is not publicly disclosed and is held privately by the firms, leaving equity-price volatility as the only source consistently available across the sample with enough coverage for statistical estimation. This constraint is made up by the nature of the sector itself: reinsurance is a niche industry, relatively understudied, composed of few firms, many of which are privately owned and therefore outside the perimeter of public market data. The resulting scarcity of sources limits both the sample size and the granularity of the analysis. As a consequence, stock volatility may decouple from fundamental solvency conditions during non-crisis periods, and the absence of bilateral exposure data prevents direct validation of the inferred network against actual contractual linkages.

The forecast horizon of $H = 10$ days used in the GFEVD follows the convention established by Diebold and Yılmaz (2014) but is not explicitly justified for the reinsurance context, where claim settlement and capital adjustment processes operate on longer time scales. The sensitivity of the network topology to this choice has not been formally tested.

The rolling-window estimation with 250 observations and 27 variables at 4 lags produces a system with substantially more parameters than observations. The Minnesota prior regularization addresses this dimensionality challenge, but the effective degrees of freedom available for estimation remain limited, and the sensitivity of the rolling results to the window size has not been explored through alternative specifications.

The sample is conditioned on survivorship: firms that exited the market through insolvency, merger, or delisting during the sample period are excluded. Such firms may represent the most relevant cases for systemic risk analysis, and their absence may bias the results toward an overly benign view of network stability.

Finally, the SRISK analysis covers only three fiscal years, and the balance-sheet data coincides with a period of moderate total connectedness. Capital adequacy risks during high-connectedness regimes, such as those identified during the COVID-19 phase, remain untested within the SRISK framework presented here. Additionally, SRISK was calibrated for bank-like balance sheets, and the absence of reinsurance-specific measures and an objective choice of capital prudential ratio (as in banks), the results should be interpreted as illustrative of the framework's structural boundary, rather than a precise capital adequacy assessment.

6.4 Future Research

Several directions for future research emerge from this work.

A natural extension would be the application of time-varying parameter BVAR models that allow the coefficient matrices and the covariance structure to evolve continuously, avoiding the discrete jumps inherent in the rolling-window approach. Such models would provide smoother estimates of connectedness dynamics and could be combined with stochastic volatility specifications to better capture the heteroskedasticity observed in the data.

The integration of balance-sheet data with the network analysis could be strengthened by extending the SRISK computation to the full sample period, conditional on the availability of historical financial statements, and by testing whether capital shortfall risk co-varies with the high-connectedness regimes identified through the Bai-Perron structural breaks.

Should bilateral retrocession or contractual exposure data become accessible, either through regulatory disclosure initiatives or data-sharing arrangements, the network could be reconstructed directly from actual risk-transfer positions rather than inferred from equity-price co-movement. Such a dataset would allow the volatility-based network to be validated against the true contractual topology, testing how closely market-implied connectedness tracks the underlying structure of risk transfer in the sector.

The cross-validation framework could be expanded by incorporating additional systemic risk measures, such as the Conditional Expected Shortfall (CoES) used in the EIOPA stress-testing methodology, or by applying Granger causality tests to establish whether changes in network centrality precede changes in market-based risk measures.

Finally, a transmission simulation seeded by the most systemically important firms identified in Chapter 5 could provide a more direct test of whether the contagion patterns implied by the GFEVD translate into economically meaningful cascade effects under realistic stress scenarios. Such a simulation would bridge the gap between the statistical associations captured by the variance decomposition and the causal transmission mechanisms that underlie systemic risk.

Appendix A

Chapter 2

A.1 Reinsurance Companies Dataset

List of companies considered the top reinsurers by A.M. Best for the year 2025, with their estimate of Gross written premiums (GWP), Net written premiums (NWP), reported in $ billions, and Combined Ratio expressed in percentage. The first section of the table reports the companies considered in the analysis, while the second section reports the companies that were discarded from the analysis and the reason for that.

A.2 Stationarity Test Results

Results of the stationarity test performed over the volatility series of the 27 reinsurance companies. The test performed were the Augmented Dickey-Fuller test (Dickey and Fuller 1979) to check the presence of unit-root behavior, and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test (Kwiatkowski et al. 1992) for level stationarity. Although the KPSS results fail to reject level stationarity, the VAR models do not require strict-level stationarity, but rather the absence of unit-roots.

A.3 BVAR Model Robustness Check

The full grid specification of each model are reported in Table A.4. The parameters tested were the *lag* of the model, tested between 1 and 6, the parameter $\Lambda \in [0.05, 0.10, 0.20, 0.50]$, and the parameter $\alpha \in [1, 2]$. The Metropolis-Hastings scale and change remained unchanged at 0.01 for all models. The table reports than, for each model, the acceptance rate, the ESS and Geweke Diagnostic, and the residual mean squared error for the BVAR and corresponding VAR model of the same lag.

A.4 Metropolis-Hastings Algorithm

The Metropolis-Hastings (MH) algorithms proceeds following these steps:

Table A.1: List of the Companies included in the Data set.

Company Name	Ticker	Country	Mkt Cap	GWP	NWP	C.Ratio	Ceded
Korean Reinsurance Co	003690.KS	South Korea	1.34	5.39	4.27	90.9	20.75
Taiping Reinsurance Co. Ltd.	0966.HK	Hong Kong	6.92	1.11	1.03	91.5	7.25
China Reinsurance (Group) Corp	1508.HK	China	8.53	7.85	6.78	89.5	13.68
The Toa Reinsurance Co, Limited	6809.T	Japan	0.25	2.30	1.84	91.1	20.17
Sompo International Holdings, Ltd.	8630.T	Japan	26.68	26.83	24.15	83.7	9.98
MS&AD Insurance Group Holdings, Inc.	8725.T	Japan	32.04	35.34	30.59	99.4	13.45
Tokio Marine & Nichido Fire Insurance Co., Ltd.	8766.T	Japan	76.77	36.30	31.58	96.1	13.02
Arch Capital Group Ltd.	ACGL	Bermuda	33.25	11.11	7.75	83.3	30.29
AXIS Capital Holdings Limited	AXS	Bermuda	7.34	2.39	1.51	107.6	36.96
Berkshire Hathaway Inc.	BRK-A	US	1070.25	88.69	86.43	82.9	2.55
Chubb Limited	CB	Switzerland	112.18	62.00	51.47	85.8	16.99
AXA XL	CS.PA	France	74.46	8.41	6.62	79.8	21.38
Everest Group, Ltd.	EG	Bermuda	14.72	12.94	11.97	89.6	7.51
Allied World Assurance Co Holdings, Ltd.	FFH.TO	Canada	38.91	7.20	5.05	88.7	29.83
Assicurazioni Generali SpA	G.MI	Italy	43.37	81.84	78.49	107.1	4.10
Hannover Rück SE	HNR1.DE	Germany	26.60	22.68	19.80	86.6	12.68
Hiscox Ltd	HSX.L	Bermuda	3.43	1.03	0.44	65.7	56.92
MAPFRE RE, Compañía de Reaseguros S.A.	MAP.MC	Spain	10.77	6.72	4.18	94.6	37.79
Markel Corp	MKL	US	24.30	1.15	1.04	100.5	9.68
Munich Reinsurance Company	MUV2.DE	Germany	62.00	34.42	33.30	93.5	3.24
QBE Insurance Group Limited	QBE.AX	Australia	47.59	22.40	18.42	83.0	17.73
Reinsurance Group of America Inc.	RGA	US	12.81	18.49	17.84	104.0	3.51
RenaissanceRe Holdings Ltd.	RNR	Bermuda	12.04	11.73	9.95	83.3	15.18
SCOR S.E.	SCR.PA	France	4.63	17.25	13.86	86.3	19.63
SiriusPoint Ltd.	SPNT	Bermuda	2.09	1.34	1.10	88.1	17.29
Swiss Re Ltd.	SREN.SW	Switzerland	54.45	44.86	37.32	89.9	16.80
W.R. Berkley Corp	WRB	US	28.89	1.55	1.42	84.1	8.16

Table A.2: List of the Companies excluded from the Data set but that were cited in the A.M. Best top 50 reinsurer's list.

Company Name	Ticker	GWP	Revenue	C.Ratio	Reason
Lloyd's	-	23.54	58.88	0.88	bank ticker
Canada Life Re	-	3.43	14.54	0.87	not listed
Ark Insurance Holdings Ltd.	-	1.11	1.38	1.03	not listed
PartnerRe Ltd.	PREJF	9.35	9.40	0.86	shorter period
General Insurance Corporation of India	GICRE.NS	4.90	10.38	1.07	shorter period
Odyssey Group Holdings, Inc.	-	3.82	6.39	0.84	not listed
R+V Versicherung AG	-	3.55	2.50	0.96	not listed
Liberty Mutual	-	3.05	30.65	0.99	not listed
Pacific LifeCorp	-	2.84	10.15	NA	not listed
American Agricultural Insurance Company	-	2.35	0.77	0.84	not listed
Convex Group Limited	-	2.33	3.67	0.82	not listed
Deutsche Rückversicherung AG	-	2.17	0.36	0.93	not listed
Aspen Insurance Holdings Limited	AHL-PD	1.89	3.37	0.85	shorter period
Ascot Group Ltd.	-	1.75	2.45	1.00	not listed
Core Specialty Insurance Holdings, Inc.	-	1.58	1.20	1.07	not listed
DEVK Gruppe	-	1.44	3.08	0.94	not listed
Arundo Re	-	1.39	0.86	0.98	not listed
Somers Re Ltd.	-	1.30	1.31	0.96	not listed
African Reinsurance Corporation	-	1.20	1.16	0.82	not listed
Qianhai Reinsurance Co., Ltd.	-	1.19	0.51	0.96	not listed
Peak Reinsurance Company Ltd.	-	1.16	1.43	0.84	not listed
Hamilton Insurance Group, Ltd.	HG	1.15	2.33	0.88	shorter period
IRB - Brasil Resseguros S.A.	IRBR3.SA	0.98	0.80	0.80	shorter period

Table A.3: Augmented Dickey-Fuller and KPSS stationarity tests on log-volatility series.

Ticker	ADF t-stat	p-value	KPSS t-stat	p-value
003690.KS	-6.419	0.01	0.271	0.1000
0966.HK	-6.218	0.01	0.707	0.0129
1508.HK	-5.654	0.01	10.908	0.0100
6809.T	-6.193	0.01	3.765	0.0100
8630.T	-5.803	0.01	1.129	0.0100
8725.T	-5.030	0.01	1.821	0.0100
8766.T	-5.188	0.01	2.141	0.0100
ACGL	-5.090	0.01	2.727	0.0100
AXS	-5.839	0.01	4.405	0.0100
BRK.A	-6.065	0.01	4.609	0.0100
CB	-5.123	0.01	3.501	0.0100
CS.PA	-5.774	0.01	1.157	0.0100
EG	-6.482	0.01	4.417	0.0100
FFH.TO	-6.328	0.01	2.060	0.0100
G.MI	-5.489	0.01	1.882	0.0100
HNR1.DE	-5.273	0.01	0.708	0.0128
HSX.L	-5.608	0.01	4.051	0.0100
MAP.MC	-5.090	0.01	1.424	0.0100
MKL	-5.601	0.01	3.254	0.0100
MUV2.DE	-5.081	0.01	1.245	0.0100
QBE.AX	-6.511	0.01	0.669	0.0164
RGA	-4.167	0.01	3.972	0.0100
RNR	-6.079	0.01	7.492	0.0100
SCR.PA	-5.918	0.01	4.618	0.0100
SPNT	-5.325	0.01	4.108	0.0100
SREN.SW	-6.072	0.01	1.364	0.0100
WRB	-6.158	0.01	4.783	0.0100

Table A.4: Diagnostics and different parameters for the BVAR models tested.

Model ID	Lag	Λ	α	MH scale	MH change	Acc Rate	Log Like	ESS	Geweke	$RMSE_{BVAR}$	$RMSE_{VAR}$	ΔRMSE
model1	1	0.05	1	0.01	0.01	0.44	93540.18	162.16	0.40	0.62	0.62	0.00
model2	2	0.05	1	0.01	0.01	0.58	94357.62	152.69	1.23	0.61	0.62	-0.01
model3	3	0.05	1	0.01	0.01	0.46	94467.20	143.15	0.53	0.60	0.61	-0.01
model4	4	0.05	1	0.01	0.01	0.42	94693.64	193.20	-0.16	0.59	0.60	-0.02
model5	5	0.05	1	0.01	0.01	0.37	94884.42	140.51	-1.22	0.58	0.58	0.00
model6	6	0.05	1	0.01	0.01	0.36	94951.63	147.51	-1.42	0.57	0.55	0.02
model7	1	0.05	2	0.01	0.01	0.43	93540.58	175.05	0.20	0.62	0.62	-0.01
model8	2	0.05	2	0.01	0.01	0.66	94365.98	110.29	2.45	0.61	0.62	-0.01
model9	3	0.05	2	0.01	0.01	0.54	94490.06	167.24	0.58	0.59	0.61	-0.02
model10	4	0.05	2	0.01	0.01	0.52	94674.69	109.50	0.16	0.59	0.60	-0.02
model11	5	0.05	2	0.01	0.01	0.48	94821.38	181.51	1.11	0.58	0.58	0.00
model12	6	0.05	2	0.01	0.01	0.49	94886.62	161.77	0.99	0.58	0.55	0.03
model13	1	0.10	1	0.01	0.01	0.38	93541.22	108.34	0.18	0.62	0.62	-0.01
model14	2	0.10	1	0.01	0.01	0.56	94357.74	158.36	-1.33	0.61	0.62	-0.01
model15	3	0.10	1	0.01	0.01	0.44	94471.39	151.49	-0.87	0.60	0.61	-0.01
model16	4	0.10	1	0.01	0.01	0.45	94701.34	200.99	-0.50	0.59	0.60	-0.01
model17	5	0.10	1	0.01	0.01	0.39	94878.65	215.50	0.69	0.58	0.58	0.00
model18	6	0.10	1	0.01	0.01	0.36	94942.09	158.65	-2.21	0.57	0.55	0.03
model19	1	0.10	2	0.01	0.01	0.41	93540.26	184.93	0.40	0.62	0.62	-0.01
model20	2	0.10	2	0.01	0.01	0.63	94369.46	99.13	-0.41	0.60	0.62	-0.02
model21	3	0.10	2	0.01	0.01	0.56	94490.18	184.07	0.66	0.59	0.61	-0.02
model22	4	0.10	2	0.01	0.01	0.55	94673.73	173.87	-0.85	0.59	0.60	-0.02
model23	5	0.10	2	0.01	0.01	0.50	94821.49	158.22	0.95	0.58	0.58	0.00
model24	6	0.10	2	0.01	0.01	0.50	94879.44	231.12	0.69	0.58	0.55	0.03
model25	1	0.20	1	0.01	0.01	0.44	93543.98	249.64	-0.68	0.62	0.62	-0.01
model26	2	0.20	1	0.01	0.01	0.58	94361.03	164.29	0.31	0.61	0.62	-0.01
model27	3	0.20	1	0.01	0.01	0.39	94474.73	163.33	0.93	0.60	0.61	-0.01
model28	4	0.20	1	0.01	0.01	0.41	94696.20	187.63	0.17	0.59	0.60	-0.01
model29	5	0.20	1	0.01	0.01	0.33	94884.93	169.52	0.41	0.58	0.58	0.00
model30	6	0.20	1	0.01	0.01	0.32	94954.10	182.91	-3.06	0.57	0.55	0.02
model31	1	0.20	2	0.01	0.01	0.42	93542.72	130.38	1.55	0.62	0.62	-0.01
model32	2	0.20	2	0.01	0.01	0.64	94364.18	122.44	-0.77	0.60	0.62	-0.02
model33	3	0.20	2	0.01	0.01	0.52	94491.01	162.87	-0.11	0.60	0.61	-0.01
model34	4	0.20	2	0.01	0.01	0.51	94670.54	206.06	-1.28	0.59	0.60	-0.02
model35	5	0.20	2	0.01	0.01	0.52	94821.11	164.98	0.28	0.58	0.58	0.00
model36	6	0.20	2	0.01	0.01	0.48	94880.20	186.87	-1.16	0.57	0.55	0.03
model37	1	0.50	1	0.01	0.01	0.47	93547.54	173.38	0.18	0.62	0.62	0.00
model38	2	0.50	1	0.01	0.01	0.56	94362.50	172.06	-0.75	0.61	0.62	-0.01
model39	3	0.50	1	0.01	0.01	0.44	94475.96	217.19	-0.49	0.60	0.61	-0.02
model40	4	0.50	1	0.01	0.01	0.40	94704.22	188.01	-0.33	0.59	0.60	-0.02
model41	5	0.50	1	0.01	0.01	0.40	94880.54	195.77	0.71	0.58	0.58	0.00
model42	6	0.50	1	0.01	0.01	0.36	94955.66	178.22	-0.07	0.57	0.55	0.02
model43	1	0.50	2	0.01	0.01	0.42	93546.73	187.22	1.50	0.62	0.62	0.00
model44	2	0.50	2	0.01	0.01	0.62	94369.56	84.12	-0.10	0.61	0.62	-0.01
model45	3	0.50	2	0.01	0.01	0.53	94487.05	174.03	-1.76	0.60	0.61	-0.02
model46	4	0.50	2	0.01	0.01	0.53	94675.78	205.77	3.03	0.59	0.60	-0.02
model47	5	0.50	2	0.01	0.01	0.48	94822.64	108.53	0.06	0.58	0.58	0.00
model48	6	0.50	2	0.01	0.01	0.49	94889.64	193.84	1.20	0.57	0.55	0.02

- **Proposal distribution**: for each iteration $i = 1, \dots, N$, given the current state $\gamma^{(i)}$, a possible candidate $\gamma^* \sim q(\gamma^*|\gamma^{(t)})$is proposed. The `BVAR` package utilizes an adaptive random-walk Metropolis proposal of the form:

  $$\gamma^* = \gamma^{(i)} + \epsilon, \qquad \epsilon \sim \mathcal{N}(0, c \cdot H^{-1})$$

  where H is the Hessian of the log marginal likelihood evaluated at the posterior mode, and c is a scaling factor controlling the proposal scale.

- **Acceptance probability**: the candidate γ^* is accepted with probability

  $$\alpha_{MH} = min\left\{1, \frac{p(y|\gamma^*)p(\gamma^*)}{p(y|\gamma^{(i)})p(\gamma^{(i)})} \cdot \frac{q(\gamma^{(i)}|\gamma^*)}{q(\gamma^*|\gamma^{(i)})}\right\}$$

  Since the proposal is symmetric, the ratio of proposal density equals unity. The marginal likelihood $p(y|\gamma)$, as in Equation 2.4, corresponds to the closed-form Normal-Inverse-Wishart expression implied by the conjugate prior, involving the determinant of the posterior covariance matrix and the associated normalization constants.

- **Adaptive tuning**: to achieve efficient mixing, the algorithm uses adaptive scaling during the burn-in phase. The algorithm monitors the acceptance rate α_{MH} and dynamically adjust the scaling factor c:

  – if the acceptance rate falls below a lower bound α_{low}, the scale factor c is decreased to $c^* = c(1 - \Delta)$;
  – if the acceptance rate exceeds an upper bound α_{upp}, the scale factor is increased to $c^* = c(1 + \Delta)$ Adaptation is applied during the initial fraction of the burn-in phase so that the chain stabilizes before posterior sampling begins.

- **Gibbs Sampling**: conditional on the new accepted hyperparameters $\gamma^{(i+1)}$, the VAR coefficients θ are sampled via a standard Gibbs update from their conjugate Normal-Inverse-Wishart posterior:

  $$\theta|\gamma, \Sigma, y \sim \mathcal{N}(\bar{\beta}, \bar{\Omega}), \qquad \Sigma|\gamma, y \sim \mathcal{IW}(\cdot)$$

  which, under the NIW prior admits closed-form solution.

A.5 Posterior Credible Intervals for GFEVD Out-Strength

Table A.5 reports the lower and upper intervals with 90% posterior credible interval for the out-strength values derived from MCMC draws of the GFEVD matrix.

A.6 Final Model Specifications

The final BVAR model specifications used in Chapter 2 are listed in Table A.6.

Table A.5: Out-strength for all 27 reinsurers with 90% posterior credible interval derived from MCMC draws of the GFEVD matrix.

Firm	Median	Lower	Upper
003690.KS	72.91	12.54	114.41
0966.HK	83.76	29.78	128.62
1508.HK	79.86	26.21	130.09
6809.T	56.87	21.17	103.37
8630.T	78.69	37.43	122.00
8725.T	83.70	44.48	116.82
8766.T	76.09	28.63	136.28
ACGL	82.89	52.60	110.33
AXS	77.92	47.27	107.76
BRK-A	66.50	38.72	97.13
CB	86.03	55.42	111.88
CS.PA	82.66	56.41	111.24
EG	84.53	52.19	110.88
FFH.TO	76.15	44.48	105.93
G.MI	88.00	69.54	104.74
HNR1.DE	78.15	41.81	131.15
HSX.L	34.03	15.25	77.56
MAP.MC	84.75	63.80	104.94
MKL	90.53	70.60	108.36
MUV2.DE	88.48	69.55	106.16
QBE.AX	49.55	21.01	105.75
RGA	88.53	69.25	105.25
RNR	71.46	27.65	117.32
SCR.PA	73.01	44.82	96.22
SPNT	73.43	47.98	102.19
SREN.SW	78.92	52.20	112.24
WRB	90.51	69.62	108.40

Table A.6: Parameters selected for the final model.

Parameter	Prior Distribution / Type	Value / Bound Range
Minnesota Prior		
λ (Overall tightness)	$N(0.2, 0.4)$	$[0.0001, 5]$
α (Lag decay)	$N(2, 0.25)$	$[0.01, 3]$
ψ (Scale parameter)	$\Gamma(\text{scale} = 0.004, \text{shape} = 0.004)$	Automatically set by package
Sum-of-Coefficients (SOC) Prior		
μ (Sum-of-coefficients)	$N(1, 1)$	$[0.0001, 50]$
Single-Unit-Root (SUR) Prior		
δ (Single-unit-root)	$N(1, 1)$	$[0.0001, 50]$
MCMC Metropolis-Hastings Sampler		
c (Scaling factor)	Constant value	0.01
α_{MH} (Acceptance rate)	Adaptive target window	$[0.25, 0.40]$ ($\Delta = 0.01$)
Burn-in phase	50% adaptive allocation	5,000 draws
Total iterations	MCMC sampler size	15,000 draws

These hyperparameters are implemented into the BVAR model using R software and the package BVAR, developed by Kuschnig and Vashold (2021). The source code to initialize the parameters and the statistical model, and to run the diagnostical tests for the convergence of the model is reported below. Full source code is available on project's GitHub repository.

```
library(BVAR)
library(coda)
library(parallel)

# Minnesota prior
minnesota_priors <- bv_minnesota(
  lambda = bv_lambda(mode=0.2, sd=0.4, min=0.0001, max=5),
  alpha  = bv_alpha(mode=2, sd=0.25, min=0.01, max=3),
  psi    = bv_psi(scale=0.004, shape=0.004, mode="auto",
   ↪  min="auto", max="auto")
)

add_soc <- function(Y, lags, par) {
  soc <- if(lags == 1) {diag(Y[1, ]) / par} else
 ↪  {diag(colMeans(Y[1:lags, ])) / par}
  list("Y" = soc, "X" = cbind(rep(0, ncol(Y)), matrix(rep(soc,
   ↪  lags), nrow = ncol(Y))))
}
add_sur <- function(Y, lags, par) {
  sur <- if(lags == 1) {Y[1, ] / par} else {colMeans(Y[1:lags, ])
 ↪  / par}
  list("Y" = sur, "X" = c(1 / par, rep(sur, lags)))
}

soc <- bv_dummy(mode=1, sd=1, min=0.0001, max=50, fun=add_soc)
sur <- bv_dummy(mode=1, sd=1, min=0.0001, max=50, fun=add_sur)

metropolis_hastings <- bv_metropolis(
  scale_hess=0.01, adjust_acc=TRUE, adjust_burn=0.5,
  acc_lower=0.25, acc_upper=0.40, acc_change=0.01
)

priors <- bv_priors(hyper=c("lambda"), mn=minnesota_priors,
 ↪  soc=soc, sur=sur)

model <- bvar(vol, lags=4, n_draw=15000, n_burn=5000, n_thin=1,
              priors=priors, mh=metropolis_hastings,
               ↪  verbose=TRUE)

# Parallel Gelman-Rubin (1992) diagnostic
n_cores <- 4
cl <- makeCluster(n_cores)
runs <- par_bvar(cl, data=vol, lags=4, n_draw=15000, n_burn=5000,
 ↪  n_thin=1, priors=priors)
```

```r
stopCluster(cl)
gelman.diag(as.mcmc(runs, vars = "lambda"))
# Effective size diagnostic
ess <- effectiveSize(model)

# plot of the convergence of the model
plot(model)
```

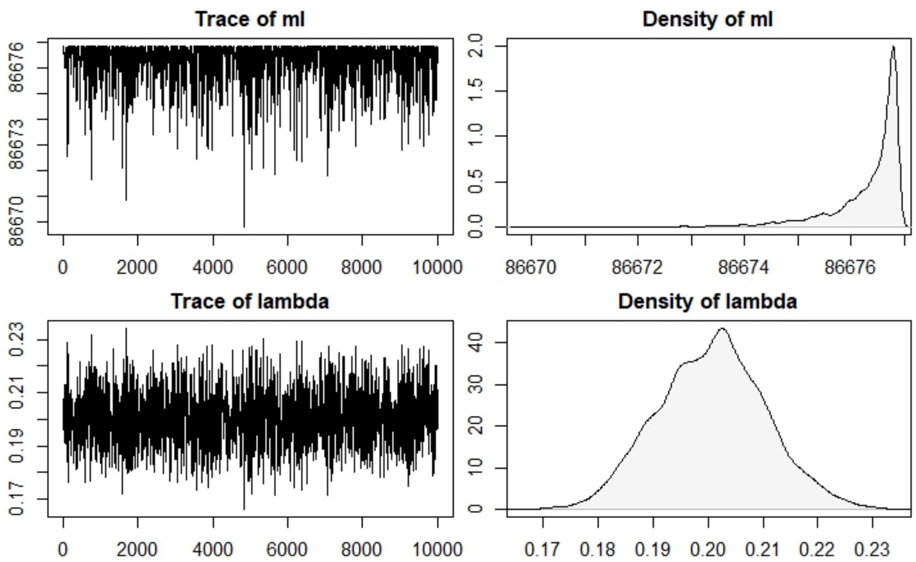

Figure A.1: Convergence of the BVAR model with the parameters specified before.

To further investigate the validity of the model and the proper convergence and exploration of the parameter space, we used the previously implemented test diagnostics:

- The ESS held a value of 1707.48, far more than the threshold value of 1000 believed to be sufficient for stable estimates and good convergence (Bürkner 2017).
- The Geweke diagnostic showed again the presence of convergence and proper mixing, with Z-score value of -0.159.
- Gelman-Rubin diagnostic (Gelman and Rubin 1992) was performed over 4 parallel MCMC chains, with a potential scale reduction factor $\hat{R}$ of 1, indicating matching between-chain and within-chain variance, exhibiting convergence of chains.

Appendix B

Chapter 3

B.1 Connectedness Table

The connectedness table below is estimated from the Generalized Forecast Error Variance Decomposition (GFEVD) for all 27 reinsurance companies. Each entry A_{ij} represents the pairwise relationship between company i and company j, or in graphical terms the weight of the edge E that quantifies the proportion of company i's forecast error variance attributable to changes in company j. The cross diagonal elements, A_{ii} represent the forecast error variance explained by its own forecast, while "From Others" row sums ($\sum_{j,j\neq i}^{n} A_{ij}$) represent the proportion of each reinsurer's forecast error variance explained by other reinsurers, and "To Others" column sums ($\sum_{i,i\neq j}^{n} A_{ij}$) represent each company's contribution to explaining the system's aggregate forecast error variance. The rightmost entry and represents the total level of interconnectedness within the network. The "Net" row captures net spillover contributions (positive = net transmitter, negative = net receiver).

Table B.1: Connectendess table estimated from the GFEVD for the 27 reinsurance companies. Rows = variance recipients (From Others), columns = variance transmitters (To Others), net = To Others - From Others.

	003690.KS	0966.HK	1508.HK	6809.T	8630.T	8725.T	8766.T	ACGL	AXS	BRK.A	CB	CS.PA	EG	FFH.TO	G.MI	HNR1.DE	HSX.L	MAP.MC	MKL	MUV2.DE	QBE.AX	RGA	RNR	SCR.PA	SPNT	SREN.SW	WRB	From Others
003690.KS	30.42	1.56	1.38	1.15	2.52	1.97	3.02	3.21	2.96	2.34	3.50	3.63	2.96	2.70	4.38	3.85	0.49	3.93	3.73	3.17	2.82	2.82	2.63	1.03	1.00	3.60	3.24	**69.58**
0966.HK	0.41	15.24	1.70	0.89	2.13	1.64	2.25	2.80	3.56	2.25	3.56	6.95	3.45	1.88	6.10	6.84	2.23	5.36	2.80	5.67	1.45	3.09	2.85	4.53	1.41	5.93	3.03	**84.76**
1508.HK	0.16	6.61	29.89	0.79	1.87	1.62	2.38	2.70	1.68	1.92	3.77	4.00	1.77	1.77	4.11	4.33	1.00	4.45	1.95	3.24	1.76	2.29	2.48	4.27	2.03	4.30	2.85	**70.11**
6809.T	0.32	0.71	0.61	46.91	1.20	0.86	1.22	2.58	2.50	1.65	2.84	3.86	3.52	1.08	3.18	3.25	0.49	3.43	1.81	3.31	1.10	1.73	2.62	3.03	1.01	3.42	1.75	**53.09**
8630.T	0.44	1.40	1.24	1.62	23.29	11.33	13.35	1.91	1.29	1.98	2.13	5.05	1.30	1.59	5.08	3.46	1.03	4.99	1.85	3.61	2.33	1.11	0.66	1.70	0.65	4.78	0.81	**76.71**
8725.T	0.60	1.80	1.18	2.35	10.72	15.92	12.01	1.75	1.49	3.01	2.26	5.38	1.15	1.82	6.02	4.58	0.44	5.70	2.04	4.70	3.14	1.67	0.81	2.55	1.17	4.96	0.77	**84.08**
8766.T	0.30	1.29	1.10	2.18	17.42	18.16	30.21	0.97	0.86	1.47	1.28	2.44	0.87	1.33	3.21	1.95	0.80	3.21	1.15	2.15	1.25	1.06	0.80	1.00	0.56	2.31	0.66	**69.79**
ACGL	0.58	1.18	0.99	2.29	2.39	1.71	2.38	17.22	5.02	3.78	4.62	5.03	5.19	2.41	4.54	4.76	0.75	4.69	4.12	4.65	2.53	3.75	3.31	2.24	2.20	5.05	2.64	**82.78**
AXS	0.59	0.66	0.60	2.02	1.49	1.05	1.49	4.68	21.95	3.31	4.85	4.64	5.65	3.58	3.55	4.48	1.50	3.44	5.62	3.96	1.99	5.10	3.39	1.81	2.34	3.99	2.27	**78.05**
BRK.A	0.52	1.11	0.52	0.97	1.79	1.03	1.83	3.01	2.25	36.34	4.74	5.18	2.97	1.72	3.64	4.53	1.23	3.60	3.24	3.77	1.33	3.94	1.97	1.39	1.56	3.29	2.52	**63.66**
CB	0.42	1.18	0.90	1.16	2.18	1.79	2.40	4.71	5.43	4.45	13.32	5.63	5.54	2.40	4.49	5.07	1.20	4.61	4.36	4.51	1.26	4.36	4.26	3.65	1.79	5.08	3.85	**86.68**
CS.PA	0.26	1.34	0.82	1.38	3.56	2.93	3.49	1.61	1.70	2.78	2.51	17.37	1.67	1.33	8.92	8.53	1.77	8.22	2.49	7.72	2.07	1.70	0.86	4.62	1.60	7.55	1.19	**82.63**
EG	0.39	1.29	0.95	1.91	2.82	2.21	3.04	5.52	5.25	3.16	4.84	4.48	14.76	4.04	4.08	4.49	0.96	4.43	3.95	4.41	1.19	4.19	4.85	2.60	2.28	4.86	3.04	**85.24**
FFH.TO	0.50	1.25	1.17	0.73	2.86	2.54	3.40	2.81	2.91	3.28	4.68	4.51	4.83	24.25	4.42	3.95	1.60	4.58	3.72	4.25	1.94	2.79	2.52	2.10	1.97	4.38	2.05	**75.75**
G.MI	0.58	1.55	1.30	2.30	4.09	3.30	4.21	2.43	2.62	3.27	3.11	8.37	2.29	2.31	11.70	7.67	1.17	7.67	2.99	6.93	2.50	2.54	1.85	3.51	1.78	6.40	1.57	**88.30**
HNR1.DE	0.35	0.62	0.53	1.45	0.96	0.68	0.94	1.07	1.20	1.71	1.38	9.02	1.26	0.65	8.55	29.29	1.46	8.38	1.49	10.01	1.25	1.44	1.26	5.66	1.48	7.17	0.71	**70.71**
HSX.L	0.15	0.27	0.41	0.52	0.72	0.61	0.75	0.47	0.41	0.88	0.78	2.44	0.48	0.38	1.74	2.61	77.03	2.00	0.84	1.75	0.47	0.45	0.28	0.90	0.38	1.91	0.35	**22.97**
MAP.MC	0.34	2.21	1.54	2.54	3.35	2.75	3.81	2.32	1.79	2.37	2.50	8.30	1.79	2.19	7.83	7.50	1.71	14.75	2.61	7.05	2.35	2.35	1.79	3.97	1.97	7.07	1.25	**85.25**
MKL	0.69	1.58	1.32	2.93	3.15	2.60	3.41	3.22	3.84	4.42	4.60	6.46	3.31	3.30	5.56	6.03	1.37	5.78	9.09	5.50	2.32	3.90	1.97	3.19	2.38	5.94	2.15	**90.91**
MUV2.DE	0.42	1.55	1.05	2.32	2.62	2.01	2.71	3.01	3.32	3.49	3.88	7.75	3.24	2.02	6.61	8.21	1.49	7.05	3.53	11.06	1.71	2.95	2.49	3.77	2.21	7.10	2.43	**88.94**
QBE.AX	0.69	0.59	0.52	0.65	1.68	1.34	1.73	1.11	1.29	1.15	1.32	1.86	1.00	0.89	2.57	1.46	0.76	1.90	1.05	1.29	67.91	0.70	1.15	1.25	0.94	2.25	0.96	**32.09**
RGA	0.65	1.46	1.38	2.42	3.24	2.51	3.51	3.39	3.53	4.50	4.59	6.48	3.34	2.27	5.75	6.19	1.50	5.74	4.04	5.53	2.02	10.95	2.17	2.61	2.00	5.62	2.61	**89.05**
RNR	0.36	0.61	0.59	0.71	1.41	1.04	1.64	4.58	4.56	1.87	5.34	2.51	6.76	3.43	2.70	3.71	0.78	2.48	3.19	3.12	1.03	3.09	33.21	1.27	1.50	2.83	5.68	**66.79**
SCR.PA	0.30	1.32	0.98	1.10	3.10	3.54	3.70	1.68	1.99	2.12	2.50	6.64	2.86	1.48	5.80	6.42	1.36	6.37	1.93	5.06	2.15	1.69	1.35	26.42	1.02	6.34	0.78	**73.58**
SPNT	0.30	1.16	0.81	1.81	1.60	1.13	1.41	2.46	2.45	3.82	3.51	5.13	2.46	2.10	3.99	5.26	1.48	4.21	2.67	4.09	1.61	3.13	1.98	2.18	33.69	3.74	1.83	**66.31**
SREN.SW	0.29	0.89	0.66	0.98	2.26	1.61	2.24	2.29	1.76	3.55	2.98	8.21	2.45	0.76	6.31	7.83	1.88	7.25	2.22	8.43	1.62	1.83	1.66	4.09	1.37	23.23	1.34	**76.77**
WRB	0.69	1.47	1.21	1.77	2.95	2.43	3.17	4.29	4.21	3.89	5.81	5.52	4.39	2.89	5.06	5.31	1.53	5.20	4.52	4.93	1.70	3.83	4.13	2.68	2.18	5.18	9.07	**90.93**
To Others	**11.29**	**36.65**	**25.49**	**40.97**	**84.10**	**74.38**	**85.51**	**70.57**	**69.90**	**72.41**	**87.89**	**139.49**	**76.52**	**52.33**	**128.18**	**132.25**	**31.98**	**128.67**	**73.94**	**122.82**	**46.87**	**67.49**	**56.09**	**71.61**	**40.77**	**125.03**	**52.31**	**74.28**
Net	**-58.29**	**-48.11**	**-44.62**	**-12.12**	**7.39**	**-9.71**	**15.72**	**-12.21**	**-8.15**	**8.75**	**1.20**	**56.86**	**-8.72**	**-23.42**	**39.89**	**61.54**	**9.01**	**43.41**	**-16.98**	**33.87**	**14.79**	**-21.56**	**-10.70**	**-1.97**	**-25.54**	**48.26**	**-38.62**	**74.28**

Appendix C

Chapter 4

C.1 Rolling BVAR Diagnostic

Effective Sample Size (ESS) and Gelman-Rubin (GR) diagnostics run for the lambda and the maximum likelihood from the Monte Carlo Markov Chain (MCMC) process used to estimate the parameters of the BVAR model. Values for the ESS greater than $1,000$ indicate great convergence of the model, while for the GR diagnostic, $\hat{R} < 1.10$ are symptoms of a satisfactory convergence of the model.

The diagnostics were performed every 20 windows, to the hardware limitations at disposal and the burdensome calculations needed to extract these metrics. However, the results reported in Table C.1 show that each of the window that were subjected to the test returned positive results, indicating proper convergence throughout the series.

C.2 Structural Break Test

C.2.1 Bai-Perron

The Bai-Perron algorithm (Bai and Perron 1998) tests for m structural breaks in a linear regression model. For our scenario, we tested for regime shifts in the unconditional mean of the connectedness index. For a time series y_t (TCI), the model with m breaks (so $m + 1$ regimes) is:

$$y_t = \mu_j + u_t, \quad \text{for } t = T_{j-1} + 1, \dots, T_j$$

The algorithm identifies the break dates $T_1, \dots, T_j$ by minimizing the global residual sum of squares (SSR) via dynamic programming (Bai and Perron (2003)):

$$\min_{T_1,\dots,T_m} \sum_{j=1}^{m+1} \sum_{t=T_{j-1}+1}^{T_j} (y_t - \mu_j)^2$$

The algorithm successfully identified 4 optimal breaks, translating to 5 distinct systemic regimes. The dates of the breaks, with their confidence interval at 95% significance level is reported in Table C.2.

Table C.1: Effective Sample Size (ESS) and Gelman-Rubin diagnostic values for the 18 windows computed.

Date	ESS - ml	ESS - Λ	GR - ml	GR - Λ
2017-02-07	1513.25	1264.98	1.00	1
2017-07-26	1005.57	985.03	1.01	1
2018-01-09	1262.34	1122.70	1.00	1
2018-06-26	1014.97	1152.48	1.00	1
2018-12-07	1022.76	1027.44	1.01	1
2019-06-14	1250.65	982.65	1.00	1
2019-12-02	1244.93	1134.34	1.00	1
2020-05-29	1285.57	1162.18	1.00	1
2020-11-17	1154.08	1189.75	1.00	1
2021-05-06	1371.64	1295.01	1.00	1
2021-10-26	1410.26	1265.43	1.02	1
2022-04-14	1148.39	1207.70	1.00	1
2022-10-19	1096.27	1066.27	1.01	1
2023-04-03	1054.01	1040.85	1.00	1
2023-10-05	1069.13	1203.48	1.01	1
2024-03-21	1297.07	1238.25	1.00	1
2024-09-13	1258.36	1221.42	1.00	1
2025-03-14	1253.76	1422.78	1.00	1

Table C.2: Dates of the regime breaks identified by the Bai-Perron algorithm

Date	Lower bound	Upper bound
2018-04-09	2018-03-28	2018-06-19
2020-03-02	2020-02-21	2020-03-09
2021-06-17	2021-06-09	2021-06-24
2023-01-18	2022-07-22	2023-02-14

The first regime coincides with high connectedness period. This coincides with the 2017 HIM hurricanes, but as the previous parts note, the NatCat event happened during a phase that was already in a high connectedness state and moving towards a declining trend, independent on the event, which did not alter significantly the moving of the TCI. The second regime coincides with a stabilization phase ($\mu = 62.15$); the third identified regime represents the maximum connectedness average value ($\mu = 79.50$), corresponding to the COVID-19 pandemic and the subsequent crisis. The fourth regime maintained medium-level connectedness, with the first part characterized by a significative decline until 2022, where the trend reverted to an increasing phase, leading to an average TCI value of $\mu = 69.01$. The last regime identified saw a gradual normalization of total connectedness ($\mu = 63.79$), although the high spike in the last windows might not be fully captured by the Bai-Perron analysis, due to their recent event.

C.2.2 OLS-CUSUM

The cumulative sum of ordinary least squares residuals tests the null hypothesis (H_0) that the mean μ is constant over time (Zeileis et al. (2002)). The empirical fluctuation process is defined as:

$$W_n(t) = \frac{1}{\hat{\sigma}\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \hat{u}_i \quad \text{for } t \in [0,1]$$

where $\hat{u}_i$ are the OLS residuals from the regression on the entire sample, and $\hat{\sigma}$ is the estimated standard deviation (Brown, Durbin, and Evans (1975)). Under H_0, $W_n(t)$ converges in distribution to a standard Brownian bridge. Instead, if the mean shifts, the residuals accumulate systemically, causing $W_n(t)$ to drift and cross the asymptotic confidence boundaries.

The p-value $= 0.000$ obtained rejects the null hypothesis, meaning that the total connectedness within the reinsurance network is empirically non-stationary. The blue line (CUSUM process) in Figure C.1 repeatedly and severely breaches the confidence interval (gray area), proving that the mean connectedness of the market fundamentally changes over time, advocating the assessment of reinsurance network through dynamic analysis and not only with the static ones.

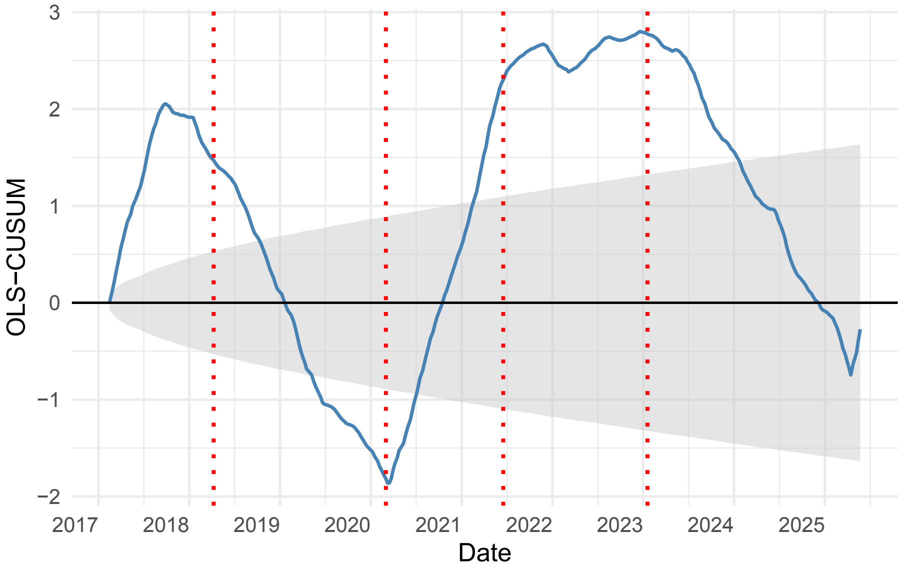

Figure C.1: CUSUM plot: The empirical fluctuation process crosses the 5% critical boundary, indicating non-stationarity of the dynamic connectedness.

Appendix D

Chapter 5

D.1 Balance Sheet Data

Balance sheet data used to calculate the SRISK in Chapter 5. Total debt and market figures are reported in $ billions after FX conversion to USD. Data via the `yfinance` API for the fiscal year 2022, 2023, and 2024. Each firm's year-end balance sheet serves as the input for the SRISK formula:

$$SRISK = k \cdot D - (1 - k) \cdot MV \cdot (1 - LRMES)$$

where $k = 0.08$ is the prudential capital ratio.

D.2 CoVaR

CoVaR (Conditional Value at Risk) and ΔCoVaR measure the marginal contribution of each firm to system-wide tail risk via quantile regression at $q = 0.05$. The system proxy is the MSCI World ETF (`IWDA.AS`), while the iShares US insurance ETF (`IAK`) served as comparison for a regional proxy. Six macro state variables from FRED are included as controls: VIX log-change, 10Y−3M term spread, Baa−10Y credit spread, 3-month T-bill rate, S&P 500 return, and VNQ REIT return; the final aligned sample spans approximately 2,002 daily observations. ΔCoVaR isolates the pure marginal contribution by differencing the $q = 0.05$ and $q = 0.50$ CoVaR estimates, so that macroeconomic conditions are held constant and only firm-specific distress is identified.

D.3 GJR-GARCH and DCC Parameters

The GJR-DCC framework uses a two-stage quasi-maximum likelihood estimation procedure to exploit the model's hierarchical structure. It proceeds in two stages: in the first one, for each asset the GJR-GARCH parameters are estimated by maximizing the univariate quasi-log-likelihood, yielding the conditional volatility series and the standardized residuals.

In the second stage, the standardized residuals from the previous step are used to estimate the DCC parameters by maximizing the correlation component of the quasi-log-likelihood.

Table D.1: Balance sheet used to calculate the SRISK for the year 2022.

Ticker	Date	Debt	Market Cap
0966.HK	2022-12-31	11.531	3.935
1508.HK	2022-12-31	2.304	2.328
ACGL	2022-12-31	2.725	22.225
AXS	2022-12-31	1.496	4.255
BRK-A	2022-12-31	122.744	684.213
CB	2022-12-31	14.877	88.496
CS.PA	2022-12-31	67.864	50.604
EG	2022-12-31	3.084	12.381
FFH.TO	2022-12-31	7.172	13.579
G.MI	2022-12-31	38.618	23.690
HNR1.DE	2022-12-31	5.874	20.805
HSX.L	2022-12-31	0.863	471.224
MAP.MC	2022-12-31	3.467	4.953
MKL	2022-12-31	4.104	17.802
MUV2.DE	2022-12-31	5.062	38.579
QBE.AX	2022-12-31	2.060	11.520
RGA	2022-12-31	3.961	8.932
RNR	2022-12-31	1.170	7.978
SCR.PA	2022-12-31	3.502	3.274
SPNT	2022-12-31	0.778	0.973
SREN.SW	2022-12-31	10.094	21.466
WRB	2022-12-31	2.837	17.880
6809.T	2023-03-31	0.029	0.174
8630.T	2023-03-31	4.613	11.894
8725.T	2023-03-31	5.366	14.559
8766.T	2023-03-31	1.673	34.753

Table D.2: Balance sheet used to calculate the SRISK for the year 2023.

Ticker	Date	Debt	Market Cap
0966.HK	2023-12-31	12.984	2.892
1508.HK	2023-12-31	1.975	2.128
ACGL	2023-12-31	2.726	26.296
AXS	2023-12-31	1.523	4.532
BRK-A	2023-12-31	128.271	785.608
CB	2023-12-31	14.495	88.913
CS.PA	2023-12-31	66.265	61.181
EG	2023-12-31	3.386	14.734
FFH.TO	2023-12-31	8.163	21.274
G.MI	2023-12-31	39.121	29.643
HNR1.DE	2023-12-31	5.396	25.672
HSX.L	2023-12-31	0.960	463.542
MAP.MC	2023-12-31	3.378	5.906
MKL	2023-12-31	3.780	18.541
MUV2.DE	2023-12-31	5.216	49.959
QBE.AX	2023-12-31	2.107	13.266
RGA	2023-12-31	4.427	10.220
RNR	2023-12-31	1.959	10.132
SCR.PA	2023-12-31	3.583	4.300
SPNT	2023-12-31	0.786	1.974
SREN.SW	2023-12-31	9.444	27.074
WRB	2023-12-31	2.837	17.150
6809.T	2024-03-31	0.020	0.207
8630.T	2024-03-31	4.568	19.376
8725.T	2024-03-31	4.722	25.654
8766.T	2024-03-31	1.482	57.894

Table D.3: Balance sheet used to calculate the SRISK for the year 2024.

Ticker	Date	Debt	Market Cap
003690.KS	2024-12-31	0.005	0.859
0966.HK	2024-12-31	10.578	5.341
1508.HK	2024-12-31	1.799	4.547
ACGL	2024-12-31	2.728	34.587
AXS	2024-12-31	1.488	7.235
BRK-A	2024-12-31	124.762	975.288
CB	2024-12-31	15.179	109.318
CS.PA	2024-12-31	64.986	69.171
EG	2024-12-31	3.587	15.188
FFH.TO	2024-12-31	9.100	30.611
G.MI	2024-12-31	39.068	41.406
HNR1.DE	2024-12-31	4.859	27.939
HSX.L	2024-12-31	0.932	457.951
MAP.MC	2024-12-31	2.781	7.451
MKL	2024-12-31	4.330	21.720
MUV2.DE	2024-12-31	6.578	62.203
QBE.AX	2024-12-31	1.801	16.737
RGA	2024-12-31	5.042	13.675
RNR	2024-12-31	1.887	12.406
SCR.PA	2024-12-31	3.693	3.855
SPNT	2024-12-31	0.639	1.653
SREN.SW	2024-12-31	8.039	38.070
WRB	2024-12-31	2.841	21.635
6809.T	2025-03-31	0.020	0.175
8630.T	2025-03-31	4.620	27.596
8725.T	2025-03-31	3.947	31.249
8766.T	2025-03-31	1.519	71.461

Table D.4: Values of the Delta CoVaR for the firms with their standard errors and the t-stat significance level. The value is considered significant for t>1.96.

Ticker	DeltaCoVaR	Std. Error	t-stat
MUV2.DE	-0.006	0.025	10.424
CS.PA	-0.006	0.027	9.022
HNR1.DE	-0.006	0.032	7.883
SREN.SW	-0.006	0.028	9.031
G.MI	-0.005	0.034	6.959
MAP.MC	-0.004	0.033	5.410
SCR.PA	-0.004	0.019	7.795
8630.T	-0.004	0.029	4.550
8766.T	-0.004	0.024	5.354
0966.HK	-0.004	0.010	9.090
HSX.L	-0.003	0.019	6.036
8725.T	-0.003	0.030	3.459
QBE.AX	-0.003	0.034	2.982
003690.KS	-0.003	0.025	4.502
1508.HK	-0.003	0.021	4.282
ACGL	-0.003	0.040	2.770
6809.T	-0.003	0.022	3.596
RGA	-0.002	0.027	3.323
AXS	-0.002	0.042	2.026
BRK.A	-0.002	0.069	2.096
MKL	-0.002	0.042	2.068
FFH.TO	-0.002	0.043	1.619
SPNT	-0.002	0.026	2.083
CB	-0.001	0.038	1.910
EG	-0.001	0.036	1.398
WRB	-0.001	0.048	1.110
RNR	0.000	0.032	-0.253

The use of a GJR-DCC model allows all three MES drivers to vary, producing estimates that react to changing market conditions. Following V. V. Acharya et al. (2016), this decomposition separates a firm's expected returns on bad market days into the systematic component and the idiosyncratic tail component. Mathematically, the first step GJR-GARCH (1,1) variance estimation can be expressed as:

$$\sigma_{i,t}^2 = \omega + (\alpha + \gamma \times I_{\epsilon_{t-1}<0}) \times \epsilon_{t-1}^2 + \beta^{GJR} \times \sigma_{t-1}^2 \tag{D.1}$$

where ω is the baseline long-run variance component, α captures the intensity of the reaction to all shocks, while γ captures the additional volatility generated by negative shocks, and β^{GJR} captures the persistence of volatility clustering in the volatility of tomorrow.

This first step volatility estimation allows the model to capture the asymmetry that financial firms exhibit to positive and negative shocks, reflecting their sensitivity to market stress. To account for fat tails typically exhibited by financial return distributions, the innovations of the GJR-GARCH model were estimated using a *Student-t* distribution. This allows the model to treat extreme outliers as part of the natural distribution, preventing the underestimation of volatility that would occur under Normal distribution.

The second step of the DCC Q-matrix evolution is described as:

$$Q_t = (1 - a - b)\bar{Q} + a(z_{t-1}z_{t-1}') + bQ_{t-1} \tag{D.2}$$

where a and b represent the sensitivity and persistence of the correlation. If the sum of the two terms is close to 1, the correlation shocks are highly persistent throughout the whole sample, and during high stress periods the higher correlation Q_t will persist for a longer period. To ensure numerical stability and convergence, the DCC model was fitted using a *Multivariate Normal* distribution. This approach follows the Quasi-Maximum Likelihood framework, which maintains statistical consistency for the correlation parameters even when the underlying innovations are non-normal. An alternative specification using a multivariate Student-t distribution was attempted but resulted in numerical instabilities that prevented stable convergence. The use of a Normal distribution rather than a Student-t implies that the model does not account for as much tail dependence, meaning that extreme co-movement during market crisis are understated. The resulting MES and LRMES estimates for the crisis periods should therefore be interpreted as conservative, with true losses during simultaneous reinsurer distress potentially exceeding those reported.

Table D.5 reports the estimated GJR-GARCH and DCC parameters for each of the 27 firms. The asymmetry coefficient $\hat{\gamma}$ is positive and statistically meaningful for the majority of the sample, confirming that negative volatility shocks generate disproportionately larger conditional variance than positive shocks of equal magnitude — the central motivation for the GJR specification over symmetric GARCH. The DCC component models the dynamic correlation between each firm and the market proxy, feeding time-varying ρ_t into the MES formula.

Table D.5: Estimated GJR-GARCH(1,1) parameters (ω,α,γ,β) and DCC parameters (a,b) for each firm. Persistence $= \alpha + \gamma/2 + \beta$.

Ticker	ω	α	β	γ	persistence	DCC_a	DCC_b
003690.KS	0	0.112	0.844	-0.017	0.949	0.000	0.957
0966.HK	0	0.083	0.873	0.008	0.960	0.006	0.989
1508.HK	0	0.102	0.885	0.001	0.988	0.014	0.943
6809.T	0	0.159	0.612	0.136	0.838	0.027	0.731
8630.T	0	0.030	0.857	0.118	0.946	0.006	0.866
8725.T	0	0.048	0.904	0.069	0.987	0.005	0.990
8766.T	0	0.063	0.875	0.042	0.959	0.006	0.842
ACGL	0	0.018	0.927	0.091	0.990	0.023	0.958
AXS	0	0.019	0.911	0.112	0.986	0.009	0.830
BRK.A	0	0.054	0.846	0.105	0.953	0.003	0.992
CB	0	0.043	0.894	0.091	0.983	0.008	0.984
CS.PA	0	0.029	0.853	0.163	0.963	0.015	0.965
EG	0	0.034	0.855	0.164	0.971	0.010	0.983
FFH.TO	0	0.107	0.780	0.045	0.909	0.012	0.967
G.MI	0	0.036	0.875	0.133	0.977	0.034	0.936
HNR1.DE	0	0.004	0.944	0.074	0.985	0.016	0.972
HSX.L	0	0.041	0.817	0.135	0.925	0.068	0.612
MAP.MC	0	0.007	0.957	0.052	0.990	0.009	0.982
MKL	0	0.065	0.828	0.122	0.954	0.025	0.608
MUV2.DE	0	0.000	0.951	0.084	0.993	0.022	0.966
QBE.AX	0	0.063	0.846	0.059	0.938	0.013	0.850
RGA	0	0.021	0.941	0.068	0.996	0.019	0.960
RNR	0	0.006	0.951	0.076	0.995	0.008	0.989
SCR.PA	0	0.017	0.964	0.036	0.999	0.043	0.908
SPNT	0	0.005	0.953	0.064	0.990	0.008	0.988
SREN.SW	0	0.053	0.846	0.121	0.960	0.023	0.952
WRB	0	0.032	0.912	0.075	0.982	0.008	0.989

Appendix E

AI usage

For the writing of the thesis, the following generative AI tools were used:

- Claude Code (Anthropic)
  - Sonnet 4.6
  - Opus 4.6
  - Opus 4.8
- Gemini CLI (Google)
  - Gemini 3.1 Pro
  - Gemini 3.1 Flash

The models were used for code development and debugging of R and Python scripts used to build the analysis tools and the statistical models, for linguistic revision and structural refinement of chapter sections, and for the methodological review for identification of potential inconsistencies in econometric specifications.

All outputs generated by the LLM were independently verified for accuracy and reliability before being incorporated in the final draft. No AI tools were used to generate research conclusions, produce statistical outputs, or generate images or infographics included in this work. All econometric analyses are executed directly via R and Python.

The use of these tools was conducted in accordance with the principles of academic responsibility, originality, and integrity required by Ca' Foscari University of Venice.

References

Acharya, Viral V., Lasse H. Pedersen, Thomas Philippon, and Matthew Richardson. 2016. "Measuring Systemic Risk." *The Review of Financial Studies* 30 (1): 2–47. https://doi.org/10.1093/rfs/hhw088.

Acharya, Viral, Robert Engle, and Matthew Richardson. 2012. "Capital Shortfall: A New Approach to Ranking and Regulating Systemic Risks." *American Economic Review* 102 (3): 59–64. https://doi.org/10.1257/aer.102.3.59.

Acharya, Viral, Lasse Pedersen, Thomas Philippon, and Matthew Richardson. 2010. "Measuring Systemic Risk." *Federal Reserve Bank of Cleveland, Working Paper*, January. https://doi.org/10.2139/ssrn.1573171.

Adrian, Tobias, and Markus K. Brunnermeier. 2016. "CoVaR." *American Economic Review* 106 (7): 1705–41. https://doi.org/10.1257/aer.20120555.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys. 2003. "Modeling and Forecasting Realized Volatility." *Econometrica* 71 (2): 579–625. https://doi.org/https://doi.org/10.1111/1468-0262.00418.

Bai, Jushan, and Pierre Perron. 1998. "Estimating and Testing Linear Models with Multiple Structural Changes." *Econometrica* 66 (1): 47–78. http://www.jstor.org/stable/2998540.

———. 2003. "Computation and Analysis of Multiple Structural Change Models." *Journal of Applied Econometrics* 18 (1): 1–22. https://doi.org/https://doi.org/10.1002/jae.659.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin. 2010. "Large Bayesian Vector Auto Regressions." *Journal of Applied Econometrics* 25 (1): 71–92. https://doi.org/https://doi.org/10.1002/jae.1137.

Banerjee, Chandan, Lucia Bevere, Michael Ewald, Erik Lindgren, and Mahesh Puttaiah. 2025. "Natural Catastrophes: Insured Losses on Trend to USD 145 Billion in 2025." sigma 1/2025. Swiss Re Institute. https://www.swissre.com/institute/research/sigma-research/sigma-2025-01-natural-catastrophes-trend.html.

Battiston, Stefano, J. Doyne Farmer, Andreas Flache, Diego Garlaschelli, Andrew G. Haldane, Hans Heesterbeek, Cars Hommes, Carlo Jaeger, Robert May, and Marten Scheffer. 2016. "Complexity Theory and Financial Regulation." *Science* 351 (6275): 818–19. https://doi.org/10.1126/science.aad0299.

Battiston, Stefano, Domenico Delli Gatti, Mauro Gallegati, Bruce Greenwald, and Joseph E. Stiglitz. 2012. "Default Cascades: When Does Risk Diversification Increase Stability?" *Journal of Financial Stability* 8 (3): 138–49. https://doi.org/https://doi.org/10.1016/j.jfs.2012.01.002.

Billio, Monica, Getmansky Mila, Andrew W. Lo, and Loriana Pelizzon. 2012. "Econometric Measures of Connectedness and Systemic Risk in the Finance and Insurance Sectors." *Journal of Financial Economics* 104 (3): 535–59.

https://doi.org/10.1016/j.jfineco.2011.12.010.

Blondel, Vincent D, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. "Fast Unfolding of Communities in Large Networks." *Journal of Statistical Mechanics: Theory and Experiment* 2008 (10): P10008. https://doi.org/10.1088/1742-5468/2008/10/p10008.

Boyer, Martin, and Théodora Dupont-Courtade. 2015. "The Structure of Reinsurance Contracts." *The Geneva Papers on Risk and Insurance Issues and Practice* 40 (May). https://doi.org/10.1057/gpp.2015.2.

Brin, Sergey, and Lawrence Page. 1998. "The Anatomy of a Large-Scale Hypertextual Web Search Engine." *Computer Networks* 30: 107–17. https://snap.stanford.edu/class/cs224w-readings/Brin98Anatomy.pdf.

Brown, R. L., J. Durbin, and J. M. Evans. 1975. "Techniques for Testing the Constancy of Regression Relationships over Time." *Journal of the Royal Statistical Society. Series B (Methodological)* 37 (2): 149–92. http://www.jstor.org/stable/2984889.

Brownlees, Christian, and Robert F. Engle. 2016. "SRISK: A Conditional Capital Shortfall Measure of Systemic Risk." *The Review of Financial Studies* 30 (1): 48–79. https://doi.org/10.1093/rfs/hhw060.

Bürkner, Paul-Christian. 2017. "Brms: An r Package for Bayesian Multilevel Models Using Stan." *Journal of Statistical Software* 80 (1): 1–28. https://doi.org/10.18637/jss.v080.i01.

Chen, Chen, Garud Iyengar, and Ciamac C. Moallemi. 2013. "An Axiomatic Approach to Systemic Risk." *Management Science* 59 (6): 1373–88. https://doi.org/10.1287/mnsc.1120.1631.

Chib, Siddhartha, and Edward Greenberg. 1995. "Understanding the Metropolis-Hastings Algorithm." *The American Statistician* 49 (4): 327–35. https://doi.org/10.1080/00031305.1995.10476177.

Csárdi, Gábor, Tamás Nepusz, Vincent Traag, Szabolcs Horvát, Fabio Zanini, Daniel Noom, Kirill Müller, David Schoch, and Maëlle Salmon. 2025. *igraph: Network Analysis and Visualization in r*. https://doi.org/10.5281/zenodo.7682609.

Cummins, J. David, and Mary A. Weiss. 2014. "Systemic Risk and the u.s. Insurance Sector." *Journal of Risk and Insurance* 81 (3): 489–528. https://doi.org/https://doi.org/10.1111/jori.12039.

Dickey, David A., and Wayne A. Fuller. 1979. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." *Journal of the American Statistical Association* 74 (366a): 427–31.

Diebold, Francis X., and Kamil Yilmaz. 2012. "Better to Give Than to Receive: Predictive Directional Measurement of Volatility Spillovers." *International Journal of Forecasting* 28 (1): 57–66. https://doi.org/https://doi.org/10.1016/j.ijforecast.2011.02.006.

Diebold, Francis X., and Kamil Yılmaz. 2014. "On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms." *Journal of Econometrics* 182 (1): 119–34. https://doi.org/https://doi.org/10.1016/j.jeconom.2014.04.012.

Doherty, N. A., and S. M. Tinic. 1981. "Reinsurance Under Conditions of Capital Market Equilibrium: A Note." *The Journal of Finance* 36 (4): 949–53. http://www.jstor.org/stable/2327559.

Engle, Robert. 2002. "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity

Models." *Journal of Business & Economic Statistics* 20 (3): 339–50. http://www.jstor.org/stable/1392121.

Evans, J. D. 1996. *Straightforward Statistics for the Behavioral Sciences*. Thomson Brooks/Cole Publishing.

Financial Stability Board. 2025a. "Scope of Insurers Subject to the Recovery and Resolution Planning Requirements in the FSB Key Attributes: Consultation Report." Financial Stability Board. https://www.fsb.org/2025/11/scope-of-insurers-subject-to-the-recovery-and-resolution-planning-requirements-in-the-fsb-key-attributes-consultation-report.

———. 2025b. "The FSB Reaffirms Its Endorsement of the IAIS Holistic Framework for the Assessment of Systemic Risk in the Insurance Sector." Financial Stability Board. https://www.iais.org/2025/11/the-fsb-reaffirms-its-endorsement-of-the-iais-holistic-framework-for-the-assessment-of-systemic-risk-in-the-insurance-sector.

Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein. 1993. "Risk Management: Coordinating Corporate Investment and Financing Policies." *The Journal of Finance* 48 (5): 1629–58. https://doi.org/https://doi.org/10.1111/j.1540-6261.1993.tb05123.x.

Fruchterman, Thomas M. J., and Edward M. Reingold. 1991. "Graph Drawing by Force-Directed Placement." *Software: Practice and Experience* 21 (11): 1129–64. https://doi.org/https://doi.org/10.1002/spe.4380211102.

Gallagher Re. 2024. "Reinsurance Market Report: Mid-Year 2025 Analysis." Gallagher Re. https://www.ajg.com/gallagherre/.

Gelman, Andrew, and Donald B. Rubin. 1992. "Inference from Iterative Simulation Using Multiple Sequences." *Statistical Science* 7 (4): 457–72. https://doi.org/10.1214/ss/1177011136.

Geweke, John. 1992. "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments." In *Bayesian Statistics 4: Proceedings of the Fourth Valencia International Meeting, Dedicated to the Memory of Morris h. DeGroot, 1931–1989*. Oxford University Press. https://doi.org/10.1093/oso/9780198522669.003.0010.

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri. 2015. "Prior Selection for Vector Autoregressions." *The Review of Economics and Statistics* 97 (2): 436–51. https://doi.org/10.1162/REST_a_00483.

Granger, C. W. J., and P. Newbold. 1974. "Spurious Regressions in Econometrics." *Journal of Econometrics* 2 (2): 111–20. https://doi.org/https://doi.org/10.1016/0304-4076(74)90034-7.

International Association of Insurance Supervisors. 2019. "Holistic Framework for Systemic Risk in the Insurance Sector." International Association of Insurance Supervisors. https://www.iaisweb.org/uploads/2022/01/191114-Holistic-Framework-for-Systemic-Risk.pdf.

———. 2025. "Global Insurance Market Report (GIMAR)." International Association of Insurance Supervisors.

Kleinberg, Jon M. 1999. "Authoritative Sources in a Hyperlinked Environment." *J. ACM* 46 (5): 604–32. https://doi.org/10.1145/324133.324140.

Koop, Gary, M.Hashem Pesaran, and Simon M. Potter. 1996. "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics* 74 (1): 119–47. https://doi.org/https://doi.org/10.1016/0304-4076(95)01753-4.

Kuschnig, Nikolas, and Lukas Vashold. 2021. "BVAR: Bayesian Vector Autoregressions with Hierarchical Prior Selection in R." *Journal of Statistical*

Software 100 (14): 1–27. https://doi.org/10.18637/jss.v100.i14.

Kwiatkowski, Denis, Peter C. B. Phillips, Peter Schmidt, and Yongcheol Shin. 1992. "Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root." *Journal of Econometrics* 54 (1-3): 159–78.

Litterman, Robert B. 1986. "Forecasting with Bayesian Vector Autoregressions: Five Years of Experience." *Journal of Business & Economic Statistics* 4 (1): 25–38. http://www.jstor.org/stable/1391384.

Mayers, David, and Clifford W. Smith. 1982. "On the Corporate Demand for Insurance." *The Journal of Business* 55 (2): 281–96. http://www.jstor.org/stable/2352704.

Miranda-Agrippino, Silvia, and Giovanni Ricco. 2018. "Bayesian Vector Autoregressions." Working Paper 18. Sciences Po OFCE. www.ofce.sciences-po.fr/pdf/dtravail/WP2018-18.pdf.

Modigliani, Franco, and Merton H. Miller. 1958. "The Cost of Capital, Corporation Finance and the Theory of Investment." *The American Economic Review* 48 (3): 261–97. http://www.jstor.org/stable/1809766.

Munich Re. 2025. "Natural Disaster Risks — Rising Trend in Losses." https://www.munichre.com/en/risks/natural-disasters.html.

National Association of Insurance Commissioners. 2025. "Report on the Cybersecurity Insurance Market." National Association of Insurance Commissioners.

Newman, M. E. J. 2006. "Modularity and Community Structure in Networks." *Proceedings of the National Academy of Sciences* 103 (23): 8577–82. https://doi.org/10.1073/pnas.0601602103.

Outreville, J. 1998. *Theory and Practice of Insurance*. https://doi.org/10.1007/978-1-4615-6187-3.

Park, Sojung Carol, and Xiaoying Xie. 2014. "Reinsurance and Systemic Risk: The Impact of Reinsurer Downgrading on Property-Casualty Insurers." *Journal of Risk and Insurance* 81 (3): 587–622. https://doi.org/10.1111/jori.12045.

Pesaran, H.Hashem, and Yongcheol Shin. 1998. "Generalized Impulse Response Analysis in Linear Multivariate Models." *Economics Letters* 58 (1): 17–29. https://doi.org/https://doi.org/10.1016/S0165-1765(97)00214-0.

Plantin, Guillame, and Jean-Charles Rochet. 2007. *When Insurers Go Bust*. Princenton University Press.

Rejda, George E., and Michael J. McNamara. 2021. *Principles of Risk Management and Insurance*. 14th ed. Boston: Pearson.

Sims, Christopher. 1993. "A Nine-Variable Probabilistic Macroeconomic Forecasting Model." In *Business Cycles, Indicators, and Forecasting*, 179–212. National Bureau of Economic Research, Inc. https://EconPapers.repec.org/RePEc:nbr:nberch:7192.

SwissRe. 2015. "The Essential Guide to Reinsurance." Swisse Re Institution.

The Geneva Association. 2010. "Systemic Risk in Insurance: An Analysis of Insurance and Financial Stability." Geneva: The Geneva Association.

Thompson, K. R. 1966. *Reinsurance: A Digest on Some Aspects of the Practice of Reinsurance and Excess Insurance and a Legal Treatise on the Subject*. The Spectator. https://books.google.it/books?id=Rr7uAAAAMAAJ.

Webb, B. L., American Institute for Property, and Liability Underwriters. 1984. *Insurance Company Operations*. Insurance Company Operations, v. 1. American Institute for Property; Liability Underwriters. https://books.google.it/books?id=Z5Sa8ZaG7gUC.

Yang, Dennis, and Qiang Zhang. 2000. "Drift-independent Volatility Estimation Based on High, Low, Open, and Close Prices." *The Journal of Business* 73 (3): 477–92. http://www.jstor.org/stable/10.1086/209650.

Yuan, Yelie, Jun Yan, and Panpan Zhang. 2021. "Assortativity Measures for Weighted and Directed Networks." *Journal of Complex Networks* 9 (2): cnab017. https://doi.org/10.1093/comnet/cnab017.

Zeileis, Achim, Friedrich Leisch, Kurt Hornik, and Christian Kleiber. 2002. "Strucchange: An r Package for Testing for Structural Change in Linear Regression Models." *Journal of Statistical Software* 7 (2): 1–38. https://doi.org/10.18637/jss.v007.i02.