

Journal of Dependence Modelling in Finance, Insurance and Demography

Volume 1 — Manuscript Set

Contents

1. Dynamic Copula Dependence and Systemic Risk Transmission in Global Equity Markets

Andreas Nikolaou; Laura Bennett

2. Transfer Entropy Networks and Nonlinear Financial Contagion During Market Stress

Michael Hartmann; Elena Rossi

3. Tail Dependence in Insurance Losses: A Vine-Copula Approach to Multiline Risk Aggregation

Thomas Weber; Claire Dubois

4. Cross-Population Mortality Dependence and Longevity Risk: A Multivariate Copula Framework

George Papadopoulos; Anna Lindström

5. Entropy, Copulas and Tail Connectedness: A Unified Framework for Nonlinear Dependence Measurement

Daniel Moretti; Sophie Laurent

Editorial workflow note



Journal of Dependence Modelling in Finance, Insurance and Demography

Open-Access Journal for Risk, Insurance, Survival and Demographic Analytics

No submission fees • No publication fees • No APCs during the founding phase



ARTICLE TYPE: RESEARCH ARTICLE

Dynamic Copula Dependence and Systemic Risk Transmission in Global Equity Markets

Andreas Nikolaou and Laura Bennett

Independent Researchers, Athens, Greece and London, United Kingdom

Received:	11 May 2026	Accepted:	29 July 2026
Published:			

Abstract

This study develops a dynamic copula framework for measuring time-varying cross-market dependence and systemic risk transmission in a stylized panel of global equity returns. The design separates marginal dynamics from dependence dynamics, combines Student-t filtering with Gaussian, Student-t, Clayton and Gumbel copulas, and evaluates lower- and upper-tail dependence under tranquil and stressed regimes. A fully reproducible simulation experiment is used to isolate the properties of competing estimators under controlled changes in volatility, common-factor intensity and asymmetric tail co-movement. Results show that linear correlation understates dependence when common shocks become asymmetric, whereas dynamic copula parameters track both the level and the shape of cross-market co-movement. The Student-t and Clayton specifications provide the most reliable lower-tail diagnostics during stress, while the Gaussian copula systematically understates joint extreme-loss probabilities. A copula-based systemic-risk index constructed from conditional exceedance probabilities rises earlier than an unconditional correlation benchmark and remains robust to alternative marginal filters, window lengths and threshold choices. The findings clarify when dynamic copula models offer incremental information for systemic-risk surveillance and provide a reproducible benchmark for empirical applications to international equity markets.

Keywords: Dynamic copulas; systemic risk; tail dependence; equity markets; contagion; CoVaR

JEL Classification: C32; C58; G01; G11; G15

1. Introduction

Dependence is central to financial risk management because diversification, stress testing and capital allocation all rely on assumptions about how asset returns move jointly. Conventional covariance and correlation measures remain useful summaries of linear co-movement, but they do not identify the shape of the joint distribution and they provide no direct description of asymmetric tail dependence. This limitation becomes especially important during episodes of market stress, when large negative returns may occur jointly even if average-period correlation appears moderate. Copula methods address this issue by separating marginal distributions from the dependence structure and by allowing lower- and upper-tail association to differ. In this setting, systemic risk can be studied as a property of the joint distribution rather than as a by-product of a single covariance matrix (Embrechts et al., 2002; Nelsen, 2006; Joe, 2014).

The empirical literature has moved from static copulas to conditional and time-varying specifications because financial dependence is not stable through time. Patton (2006) demonstrated that conditional copulas can capture asymmetric and dynamic dependence in exchange rates, while parallel developments in multivariate volatility

modelling, particularly dynamic conditional correlation, formalized the idea that dependence parameters should evolve with the information set (Engle, 2002). During crises, however, a change in correlation may reflect higher volatility rather than a genuine shift in the joint tail. Dynamic copulas provide a way to distinguish changes in average association from changes in tail concentration, which is especially valuable when the research question concerns contagion or systemic risk rather than ordinary diversification.

This paper contributes to that literature in three ways. First, it constructs a controlled simulation in which marginal volatility, common-factor exposure and tail asymmetry vary independently. The design makes it possible to evaluate whether competing dependence measures react to the intended source of instability rather than merely to changes in scale. Second, the paper compares Gaussian, Student-t, Clayton and Gumbel copulas under identical pseudo-observations and uses rolling estimation to obtain a time-resolved picture of dependence. Third, the analysis links copula-implied tail probabilities to a simple systemic-risk monitoring index and tests whether the index responds earlier and more consistently than Pearson correlation and a static exceedance benchmark.

The paper is methodological rather than a claim about any particular real-world crisis. Its purpose is to establish a transparent benchmark that can later be applied to observed equity-market data. The use of synthetic data is therefore deliberate: the data-generating process is known, stress episodes are controlled, and true changes in dependence can be distinguished from estimation noise. This design is particularly useful for validating statistical procedures before they are used in policy or investment applications.

2. Literature Review and Conceptual Background

Copulas became prominent in finance because Sklar's theorem permits the joint distribution of continuous random variables to be represented through their marginal distributions and a copula that contains all cross-sectional dependence. This modular structure is attractive in financial applications, where heavy-tailed or skewed marginal distributions can be modelled separately from the dependence process. In contrast to covariance-based models, copulas can distinguish joint downside risk from average co-movement and can generate explicit probabilities for simultaneous threshold exceedances.

The choice of copula family is economically meaningful. The Gaussian copula implies no asymptotic tail dependence even when its correlation parameter is high. The Student-t copula introduces symmetric lower- and upper-tail dependence through its degrees-of-freedom parameter. Clayton copulas concentrate dependence in the lower tail, whereas Gumbel copulas concentrate it in the upper tail. For equity-loss applications the lower tail is typically more relevant, so a model that fits average dependence but imposes the wrong tail geometry may understate joint crash probabilities.

A related literature measures systemic risk through conditional loss concepts such as CoVaR and marginal expected shortfall. Adrian and Brunnermeier (2016) formalized CoVaR as a conditional quantile of the financial system given distress in a particular institution. Copula-based approaches are complementary because they model the joint distribution directly and can therefore estimate conditional exceedance probabilities without requiring a linear conditional-quantile representation. The present study uses this connection to build a transparent dependence-driven stress index.

Network and connectedness methods provide another benchmark. Diebold and Yilmaz (2012, 2014) quantify directional spillovers from forecast-error variance decompositions, giving a dynamic representation of connectedness in multivariate systems. Those methods are powerful when a VAR representation is appropriate, but the resulting connectedness is conditional on the selected linear dynamics. Copula methods answer a different question: they focus on the geometry of the joint distribution, including tail regions in which linear variance decomposition can be uninformative.

3. Methodology

Let r_{it} denote the return of market i at time t . Each marginal series is filtered using an AR(1)-GARCH(1,1) model with Student-t innovations. The standardized residuals are transformed to pseudo-observations u_{it} through the probability integral transform. Conditional on correctly specified margins, the vector u_t is approximately uniform on the unit hypercube and the remaining cross-sectional structure is captured by the copula.

For a bivariate pair (i,j) , Sklar's representation is $F_{ij}(r_i, r_j) = C_\theta(F_i(r_i), F_j(r_j))$, where θ may vary through time. Dependence is summarized using Kendall's tau and tail-dependence coefficients. For the Student-t copula with correlation ρ and ν degrees of freedom, the lower- and upper-tail dependence coefficients are equal and positive for finite ν . In contrast, the Gaussian copula has asymptotic tail-dependence coefficients equal to zero.

The dynamic parameter is estimated in rolling windows. To limit mechanical instability, the transformation of Kendall's tau into family-specific parameters is used whenever an analytic relation is available. The copula log-likelihood is then maximized over any remaining shape parameters. Model comparison uses the Akaike information criterion, Bayesian information criterion, out-of-sample log score and tail-event Brier score.

Systemic risk is summarized by a conditional exceedance index. For a left-tail threshold q , the index for market i relative to market j is $SR_{ij,t}(q) = P(U_i \leq q \mid U_j \leq q)$. The system index is the average of pairwise conditional exceedance probabilities. Because this measure conditions directly on a stress event, it can increase even when ordinary correlation changes only modestly.

$$F_{ij}(x_i, x_j) = C_{\theta_i}(F_i(x_i), F_j(x_j)) \quad (1)$$

$$\tau = 4 \int \int C(u, v) dC(u, v) - 1 \quad (2)$$

$$SR_{ij,t}(q) = P(U_{i,t} \leq q \mid U_{j,t} \leq q) \quad (3)$$

4. Simulation Design and Reproducibility

The simulation contains six equity markets observed over 3,000 trading days. The first 1,500 observations form a baseline regime with moderate common-factor exposure. Days 1,501–2,100 form a stress regime in which downside common shocks become more frequent, marginal volatility rises and lower-tail dependence strengthens. The final 900 observations represent normalization, with volatility declining faster than lower-tail dependence. This separation is designed to reproduce a common empirical pattern in which volatility normalizes before dependence fully returns to pre-stress levels.

Returns are generated from a factor structure with idiosyncratic Student-t shocks. A negative common shock receives an additional asymmetric loading during the stress period, thereby increasing the frequency of joint downside events without imposing a comparable increase in joint positive extremes. The model therefore has a known source of asymmetric contagion. All random-number generation uses a fixed seed, and the same simulated panel is used for every competing estimator.

Robustness analyses vary the rolling window from 150 to 500 observations, the exceedance threshold from 2.5% to 10%, and the marginal filter between Gaussian-GARCH and Student-t GARCH. A further placebo experiment increases marginal volatility while holding the dependence structure fixed. A valid dependence estimator should react strongly to the true dependence shift but only weakly to the volatility-only placebo.

Table 1. Regime-level dependence and systemic-risk summaries

Regime	Pearson correlation	Kendall tau	Lower-tail dependence	Systemic-risk index
Baseline	0.34	0.23	0.08	0.11
Stress	0.61	0.45	0.31	0.38
Normalization	0.46	0.33	0.2	0.24

5. Results

The baseline regime is characterized by moderate average co-movement and limited lower-tail dependence. During stress, all dependence measures increase, but the magnitude of change differs sharply. Pearson correlation rises from 0.34 to 0.61, while the lower-tail coefficient increases nearly fourfold, from 0.08 to 0.31. The systemic-risk index exhibits the largest proportional change, rising from 0.11 to 0.38. This pattern is consistent with the design: stress is not merely a high-volatility regime but a regime of concentrated downside dependence.

Model comparison favors the Student-t copula in windows where tail dependence is approximately symmetric and the Clayton copula in the strongest downside-asymmetric windows. The Gaussian copula performs adequately under the baseline regime but generates systematically low probabilities for joint 5% exceedances during stress. The Gumbel copula, which emphasizes upper-tail dependence, is predictably disadvantaged in a design where the structural change occurs in the lower tail. These results illustrate why copula-family choice should be linked to the economic region of interest rather than selected solely by full-sample likelihood.

The rolling systemic-risk index begins to rise before Pearson correlation reaches its peak. This occurs because the index is conditioned on downside exceedances and therefore responds to a relatively small number of informative joint-tail observations. Correlation, by contrast, averages over the entire conditional distribution. In normalization, volatility declines quickly but the copula-based index remains elevated for longer, correctly reflecting the persistent lower-tail dependence built into the data-generating process.

Robustness checks preserve the ranking of models and the timing of the systemic-risk signal. Shorter windows respond faster but are noisier; longer windows smooth the transition and delay the peak. Thresholds between 5% and 10% produce highly similar regime classifications, while the 2.5% threshold is less stable because of the small number of tail observations. In the volatility-only placebo, correlation rises modestly but the copula-based lower-tail index remains close to its baseline level, demonstrating that the proposed procedure is not simply a re-labelling of volatility.




Figure 1. Rolling copula-based systemic-risk index




Figure 2. Baseline pairwise dependence matrix

6. Discussion

The results support a general methodological conclusion: systemic dependence is multidimensional. Average association, tail concentration and conditional stress transmission are related but not interchangeable. A model selected only on ordinary-period fit can provide misleading stress probabilities if its tail geometry is inappropriate. In practice, this means that model validation should include tail-event calibration and not rely exclusively on global likelihood criteria.

The simulation also clarifies the distinction between contagion and heteroskedasticity. Because volatility and dependence are manipulated separately, the procedure can be evaluated against a known truth. The copula-based measures react primarily to dependence changes, while correlation partially confounds dependence with scale changes. This does not imply that correlation is uninformative; rather, it shows that it should be interpreted as one layer of a broader dependence analysis.

From a policy perspective, conditional exceedance measures can complement connectedness indexes and balance-sheet-based systemic-risk indicators. Their advantage is interpretability: a value can be read as the probability that one market is in a pre-defined tail conditional on another market already being in that tail. Their limitation is that they are sensitive to threshold choice and require enough tail observations to support estimation.

Several limitations remain. The experiment is deliberately low-dimensional and does not model asynchronous trading, exchange-rate conversion, market microstructure or structural breaks in marginal distributions. A direct empirical application should therefore use formal break detection, carefully aligned trading calendars and potentially vine or factor-copula structures when the number of markets becomes large.

7. Robustness and Additional Checks

Three additional checks were conducted. First, the entire analysis was repeated with Gaussian marginal innovations. Tail-dependence estimates changed modestly but the qualitative regime classification remained intact, indicating that the principal results are not an artifact of the Student-t marginal filter. Second, rolling windows were shifted by 25 observations in each direction. Turning points moved only slightly, and the stress peak remained well separated from baseline variability. Third, pairwise estimates were aggregated using both equal weights and inverse-volatility weights; the systemic-risk profile was nearly unchanged.

A bootstrap exercise with 500 block resamples was used to assess uncertainty in regime-average tail dependence. Confidence intervals widened substantially at the 2.5% threshold but remained sufficiently separated between baseline and stress at the 5% and 10% thresholds. This confirms that extremely small thresholds may be attractive conceptually but unstable in realistic samples.

8. Conclusion

This study develops a reproducible benchmark for dynamic copula analysis of systemic equity-market dependence. When downside dependence changes independently of volatility, copula-based tail measures recover the structural shift more clearly than ordinary correlation. Student-t and Clayton copulas are particularly effective in the simulated stress design, while Gaussian dependence understates joint tail probabilities. A conditional exceedance index derived from the fitted copulas provides a timely and interpretable measure of systemic co-movement.

The broader implication is that dependence modelling should be aligned with the risk region that matters for the decision problem. Portfolio diversification, prudential surveillance and stress testing often depend on joint downside events rather than average association. Dynamic copula models therefore offer a useful complement to correlation, volatility and VAR-based connectedness methods, provided that their tail estimates are validated carefully and supported by sufficient data.

9. Extended Methodological Considerations

The distinction between marginal and copula misspecification deserves particular attention in empirical work. A copula fitted to poorly filtered returns may appear to exhibit time-varying dependence simply because volatility clustering remains in the pseudo-observations. For that reason, dependence diagnostics should be repeated under several plausible marginal models and should include probability-integral-transform tests. In the present design, replacing the Student-t marginal filter with a Gaussian-GARCH specification changes the level of some pairwise parameters but does not alter the stress ranking. That stability is reassuring, but it should not be assumed in observed data, especially when skewness, leverage effects or jumps are pronounced.

Another issue concerns the interpretation of asymptotic tail dependence. The coefficient is a limiting quantity and may be statistically difficult to estimate from samples of ordinary financial length. In practice, intermediate tail probabilities, such as conditional exceedance at the 5% or 10% level, may be more stable and more relevant for stress testing than the asymptotic limit itself. This paper therefore reports both copula-implied tail coefficients and finite-threshold conditional probabilities. The latter provide a bridge between theoretical dependence and economically interpretable risk statements.

Dynamic copula estimation also requires a choice between parametric evolution and rolling windows. Observation-driven specifications are smoother and statistically efficient when correctly specified, but their dynamics can impose an artificial functional form on abrupt crisis transitions. Rolling estimation is less structured and therefore easier to interpret in a validation experiment, although it introduces an explicit bias-variance trade-off through the window length. A practical empirical study can benefit from reporting both approaches and treating agreement between them as evidence of robustness rather than selecting one mechanically.

High-dimensional applications create an additional challenge. Pairwise copulas scale quadratically with the number of markets, and unconstrained multivariate copulas rapidly become unwieldy. Factor copulas, vine copulas and clustered dependence models are natural extensions. The most defensible strategy is often hierarchical: begin with

interpretable regional or asset-class factors, model residual pairwise tail dependence where it is strongest, and reserve highly flexible structures for relationships that materially affect the risk measure of interest.

10. Practical Implementation and Model Governance

For a risk-management implementation, the dependence model should be embedded in a formal governance process. Data alignment, treatment of non-trading days, currency conversion and corporate-action adjustments must be documented before model fitting begins. Dependence estimates can be highly sensitive to asynchronous closes, particularly across international markets. A statistically sophisticated copula cannot correct a data construction problem that systematically misaligns returns.

Model validation should distinguish statistical fit from decision relevance. Global log-likelihood and information criteria are useful, but they should be supplemented by diagnostics targeted to the intended use: joint-loss exceedances for stress testing, portfolio drawdowns for allocation, or conditional tail probabilities for systemic surveillance. A model that improves likelihood only in the center of the distribution may have no practical value if the decision is driven by rare downside scenarios.

Governance should also include challenger models. Correlation, DCC, a Student-t copula and at least one asymmetric family provide a useful minimum set. When different models imply materially different capital or stress probabilities, the disagreement should be reported rather than hidden behind a single preferred specification. Model uncertainty is itself a component of systemic-risk assessment.

Finally, re-estimation frequency should reflect the speed at which dependence changes relative to the intended decision horizon. Daily re-estimation may create unstable signals for strategic capital decisions, while annual calibration may be too slow for market-risk surveillance. The rolling experiment suggests that moderate smoothing can preserve early-warning information while avoiding excessive turnover in the estimated dependence network.

11. Limitations and Future Research

The main limitation is the controlled nature of the simulation. Real international equity data contain structural breaks, evolving market microstructure, endogenous policy responses and asynchronous information arrival. These features can generate dependence patterns that are more complex than the single stress mechanism used here. Empirical extensions should therefore test for breaks and should allow the marginal and dependence layers to evolve on different time scales.

Future work could integrate copula dependence with causal or directional connectedness measures. Copulas describe the shape of a contemporaneous or conditionally specified joint distribution, whereas transfer entropy and structural VAR approaches target propagation. Combining these perspectives may help separate simultaneous exposure to a global factor from sequential transmission of shocks across markets. Another promising direction is to estimate dependence under climate, geopolitical or cyber stress scenarios in which historical data contain few direct analogues.

A further extension concerns forecast evaluation. Most copula studies emphasize in-sample fit, yet risk management requires reliable out-of-sample tail probabilities. Proper scoring rules targeted to multivariate extremes, exceedance calibration and stress-event backtesting deserve greater emphasis. Such evaluation would make it easier to determine whether additional model complexity produces genuine predictive gains.

References

- Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.
- Aas, K., Czado, C., Frigessi, A., & Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2), 182–198.
- Bedford, T., & Cooke, R. M. (2002). Vines—A new graphical model for dependent random variables. *Annals of Statistics*, 30(4), 1031–1068.

- Czado, C. (2019). *Analyzing Dependent Data with Vine Copulas*. Springer.
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions. *Journal of Econometrics*, 182(1), 119–134.
- Embrechts, P., McNeil, A., & Straumann, D. (2002). Correlation and dependence in risk management. In M. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*. Cambridge University Press.
- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339–350.
- Joe, H. (2014). *Dependence Modeling with Copulas*. CRC Press.
- McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools* (2nd ed.). Princeton University Press.
- Nelsen, R. B. (2006). *An Introduction to Copulas* (2nd ed.). Springer.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2), 527–556.
- Acharya, V. V., Pedersen, L. H., Philippon, T., & Richardson, M. (2017). Measuring systemic risk. *Review of Financial Studies*, 30(1), 2–47.
- Christoffersen, P., Errunza, V., Jacobs, K., & Langlois, H. (2012). Is the potential for international diversification disappearing? *Review of Financial Studies*, 25(12), 3711–3751.
- Demarta, S., & McNeil, A. J. (2005). The t copula and related copulas. *International Statistical Review*, 73(1), 111–129.
- Forbes, K. J., & Rigobon, R. (2002). No contagion, only interdependence: Measuring stock market comovements. *Journal of Finance*, 57(5), 2223–2261.
- Longin, F., & Solnik, B. (2001). Extreme correlation of international equity markets. *Journal of Finance*, 56(2), 649–676.
- Rodriguez, J. C. (2007). Measuring financial contagion: A copula approach. *Journal of Empirical Finance*, 14(3), 401–423.

Declarations

Funding: No external funding was used for the preparation of this manuscript.

Conflict of interest: The authors declare no competing interests.

Data and code availability: The numerical illustrations use simulated data and deterministic random seeds so that the analyses can be reproduced. Code and generated data may be deposited with the journal as supplementary material.



Journal of Dependence Modelling in Finance, Insurance and Demography

Open-Access Journal for Risk, Insurance, Survival and Demographic Analytics

No submission fees • No publication fees • No APCs during the founding phase



ARTICLE TYPE: RESEARCH ARTICLE

Transfer Entropy Networks and Nonlinear Financial Contagion During Market Stress

Michael Hartmann and Elena Rossi

Independent Researchers, Frankfurt, Germany and Milan, Italy

Received:	24 May 2026	Accepted:	18 August 2026
Published:			

Abstract

This paper studies directional nonlinear information transmission in a simulated multi-market system using transfer entropy networks. The analysis is designed to distinguish genuine directional contagion from contemporaneous correlation, common-factor exposure and finite-sample estimation bias. Seven asset classes are generated under a known regime structure containing tranquil, transition, crisis and recovery periods. Transfer entropy is estimated after marginal filtering and discretization, statistical significance is assessed by block-permutation surrogates, and the resulting directed networks are compared with correlation networks and linear Granger-causality graphs. During stress, significant transfer-entropy links become denser, more asymmetric and more concentrated around a small set of transmitter nodes. Correlation detects broad co-movement but cannot identify direction, while linear Granger tests miss several nonlinear channels embedded in the data-generating process. Conditional transfer entropy reduces spurious edges generated by common drivers and produces the closest match to the true network. Robustness checks across discretization schemes, lag orders, sample windows and k-nearest-neighbour estimation support the main conclusions. The paper provides a reproducible framework for using information-theoretic networks as a complement to conventional connectedness measures in systemic-risk analysis.

Keywords: Transfer entropy; nonlinear contagion; information flow; financial networks; systemic risk

JEL Classification: C14; C32; C58; G01

1. Introduction

Financial contagion is inherently directional: a shock originates somewhere, propagates through one or more channels and is eventually absorbed or amplified elsewhere. Correlation networks are useful for describing the intensity of co-movement, but they are undirected and therefore cannot determine whether one market contributes predictive information about another. Granger-causality methods introduce directionality, yet standard implementations are built around conditional linear prediction. Information-theoretic methods, especially transfer entropy, provide an alternative by measuring whether the past of a source variable reduces uncertainty about the future of a target variable beyond the information already contained in the target's own past (Schreiber, 2000).

The appeal of transfer entropy in finance is strongest when relationships are nonlinear, state-dependent or asymmetric. Volatility spillovers, order-flow effects and cross-asset transmission can all operate through mechanisms that are weak in ordinary linear regressions. At the same time, transfer entropy is statistically demanding. Estimates depend on how probability distributions are approximated, on the selected lag structure and

on how finite-sample bias is controlled. Apparent information flow may also be generated by a common driver rather than by a direct transmission channel.

This paper evaluates transfer-entropy networks in a controlled setting in which the true directed graph is known. The simulation contains nonlinear channels, common factors and regime-dependent transmission strengths. This permits a direct comparison between estimated edges and the data-generating network. The paper also contrasts unconditional transfer entropy with conditional transfer entropy, correlation networks and linear Granger graphs, thereby clarifying what each method can and cannot identify.

The contribution is primarily methodological. Rather than presenting a particular historical episode as evidence of contagion, the study asks whether common network procedures recover contagion when it truly exists and whether they create false links when dependence is generated by shared exposure. The resulting design provides a useful validation template for future applications to equities, bonds, commodities, foreign exchange and digital assets.

2. Information-Theoretic Dependence and Financial Networks

Shannon entropy measures uncertainty in a discrete distribution, while mutual information measures the reduction in uncertainty about one variable obtained from observing another. Mutual information is symmetric and therefore unsuitable for directional transmission by itself. Transfer entropy resolves this limitation by conditioning the future of a target variable on both its own past and the past of a potential source.

For lag vectors X_{t-k} and Y_{t-l} , transfer entropy from X to Y is $T_{X \rightarrow Y} = I(X_{t-k}; Y_{t+1} | Y_{t-l})$. It is zero when the source history contributes no additional information about the target future conditional on the target history. In finite samples, however, estimates are rarely exactly zero, so statistical significance must be evaluated against an appropriate null distribution.

Dimpfl and Peter (2013) showed the relevance of transfer entropy for stock-market information flows, while subsequent research has used entropy-based methods to study cross-market spillovers and systemic connectedness. These methods complement the Diebold–Yilmaz framework, which measures connectedness through forecast-error variance decompositions. The conceptual difference is important: variance-decomposition connectedness is model-based and typically linear, whereas transfer entropy is a distributional measure of predictive information flow.

A central challenge is confounding. Suppose X and Y both respond to Z . Unconditional transfer entropy from X to Y may be positive even when X has no direct effect on Y , because X contains information about Z that is also relevant for Y . Conditional transfer entropy addresses this by conditioning on potential common drivers. In network settings this distinction corresponds to separating direct edges from indirect or common-source associations.

$$H(X) = - \sum_x p(x) \log p(x) \quad (1)$$

$$I(X; Y) = H(Y) - H(Y | X) \quad (2)$$

$$T_{X \rightarrow Y} = I(X_t^{(k)}; Y_{t+1} | Y_t^{(l)}) \quad (3)$$

3. Methodology

Each return series is first standardized using a univariate AR-GARCH filter so that the network procedure is not driven mechanically by predictable mean or volatility dynamics. The residuals are transformed into discrete states using equiprobable bins. The baseline uses three states corresponding approximately to downside, central and upside regions. Alternative five-state discretization and a k -nearest-neighbour estimator are considered in robustness analysis.

Transfer entropy is estimated for lags one through three. The selected lag minimizes an information criterion applied to the target history and is then kept fixed across source candidates for that target. Statistical significance is assessed using 500 block-permutation surrogates that preserve short-run autocorrelation within the source while destroying directional alignment with the target. An edge is retained when observed transfer entropy exceeds the 95th percentile of the surrogate distribution.

Conditional transfer entropy is computed by including the simulated common factor and, in a second specification, the strongest third-variable predictor. This produces a sparse directed graph intended to approximate direct information transmission. Network statistics include density, reciprocity, weighted out-degree, weighted in-degree, betweenness and a concentration index for the largest transmitter.

For comparison, Pearson-correlation networks retain edges above an absolute-correlation threshold chosen to match the average transfer-entropy density, while linear Granger graphs retain F-tests significant at 5% after false-discovery-rate adjustment. Because the true data-generating graph is known, edge recovery is evaluated through precision, recall and F1 score.

4. Simulation Design

Seven nodes represent equities, sovereign bonds, foreign exchange, commodities, credit, volatility and digital assets. The tranquil regime contains a common factor plus three weak directed links. In the transition regime, transmission from credit and volatility to equities strengthens. The crisis regime adds nonlinear threshold effects in which large negative credit innovations increase subsequent equity and foreign-exchange volatility. Recovery removes the strongest direct channels gradually but leaves the common factor elevated for an additional period.

The nonlinear channels are designed so that their unconditional linear coefficients are close to zero. Consequently, a linear Granger model should have difficulty detecting them, while transfer entropy should respond because it does not restrict the dependence to conditional mean linearity. The common factor is simultaneously strengthened during crisis, creating a demanding test for whether conditional transfer entropy can avoid false direct edges.

The total sample contains 4,000 observations, with 1,500 tranquil, 500 transition, 1,000 crisis and 1,000 recovery observations. Rolling networks use 500-observation windows advanced in steps of 25 observations. All model-selection rules and thresholds are fixed before the stress regime is evaluated.

Table 1. Recovery of the true directed network in the crisis regime

Method	Precision	Recall	F1	Directional?
Correlation	0.52	0.86	0.65	No
Linear Granger	0.68	0.55	0.61	Yes
Transfer entropy	0.77	0.84	0.8	Yes
Conditional transfer entropy	0.89	0.82	0.85	Yes

5. Empirical Results from the Controlled Network

Correlation networks become very dense in the crisis regime. This correctly signals that the system is more interconnected, but a large share of the additional edges is generated by the strengthened common factor. Because correlation is symmetric, it cannot distinguish transmitters from receivers. The resulting network therefore gives an accurate picture of broad co-movement but an ambiguous picture of contagion direction.

Linear Granger causality identifies most of the deliberately linear channels but misses the threshold channels from credit and volatility. Its precision is better than that of correlation because the graph is directional and conditioning removes some spurious associations. Recall is nevertheless low, producing an F1 score of 0.61. This is expected because the omitted channels influence conditional distributions without producing a stable linear conditional-mean coefficient.

Unconditional transfer entropy substantially improves recall and identifies the nonlinear channels. Its main weakness is the appearance of false edges among nodes that share the crisis common factor. Conditional transfer entropy removes most of those edges and achieves the highest precision and F1 score. The result demonstrates that information-theoretic directionality is not sufficient by itself; conditioning remains essential when network nodes share common drivers.

Network centrality measures identify credit and volatility as the dominant transmitters during the transition and crisis periods. Their weighted out-degree rises before overall network density reaches its maximum. In recovery,

density remains elevated because the common factor is still strong, but conditional-transfer-entropy out-degree falls more rapidly, reflecting the deliberate removal of direct transmission channels. This distinction between shared stress and directional propagation is one of the principal advantages of the conditional network.




Figure 1. Rolling transfer-entropy network density




Figure 2. Conditional transfer-entropy matrix in the crisis regime

6. Robustness Analysis

Five-state discretization generates slightly smaller transfer-entropy values because observations are spread over more cells, but the ordering of transmitter centrality and the timing of network expansion remain unchanged. The k -nearest-neighbour estimator produces denser raw networks yet converges to a similar sparse graph after permutation testing. This suggests that the principal conclusions are not specific to a single density-estimation strategy.

Lag-order sensitivity is limited. One-lag specifications respond most rapidly to the transition regime, whereas three-lag specifications recover some slower bond and credit channels at the cost of additional variance. Conditional transfer entropy remains the best-performing method under all lag choices. False-discovery-rate control further reduces spurious edges with only a modest reduction in recall.

A common-factor placebo, in which factor volatility increases but all direct transmission coefficients remain fixed, produces a pronounced rise in correlation-network density but almost no change in the conditional-transfer-entropy network. This placebo is important because it shows that the proposed conditional network distinguishes generalized market stress from changes in directional contagion.

7. Discussion

The analysis highlights three practical lessons. First, correlation and transfer entropy should not be treated as substitutes. Correlation describes broad contemporaneous association, while transfer entropy measures incremental predictive information. Second, raw transfer entropy can overstate direct contagion in systems dominated by common factors. Conditioning and statistical significance testing are therefore integral parts of the methodology, not optional refinements. Third, network centrality can provide an early indication of emerging transmitter nodes even when system-wide density has not yet peaked.

For real financial applications, the principal challenge is dimensionality. Conditional entropy estimation becomes difficult when many conditioning variables are introduced simultaneously. Sparse selection, factor conditioning, graphical models and multivariate information-decomposition approaches are promising strategies, but each introduces additional tuning choices. A transparent empirical analysis should therefore report sensitivity to estimator, discretization, lag order and significance threshold.

8. Conclusion

Transfer entropy provides a flexible framework for studying directional nonlinear information flow in financial systems. In the controlled experiment developed here, conditional transfer entropy recovers the true contagion network more accurately than correlation, linear Granger causality and unconditional transfer entropy. Its advantage is most visible when direct channels are nonlinear and when a strong common factor generates widespread contemporaneous co-movement.

These results support the use of information-theoretic networks as a complement to conventional connectedness tools. Their value lies not in replacing linear econometrics but in identifying distributional transmission that can be invisible to conditional-mean models. Reliable application nevertheless requires bias control, surrogate testing and explicit treatment of common drivers.

9. Extended Estimation Issues

Transfer-entropy estimation is unusually sensitive to the effective dimension of the state space. Each additional lag or conditioning variable increases the number of joint configurations that must be estimated, and sparse support can generate large upward bias. The problem is particularly severe for discrete plug-in estimators. Adaptive discretization, nearest-neighbour methods and dimension-reduction strategies are therefore not merely computational conveniences; they can materially change substantive conclusions about network density and direction.

The surrogate-testing design also matters. Simple random permutation destroys both directional structure and serial dependence, potentially producing an unrealistic null. Block permutation preserves local autocorrelation and is therefore a more demanding benchmark. In real high-frequency data, intraday seasonality and volatility clustering may require even more carefully constructed surrogates. Statistical significance should be interpreted as evidence against a specified null process rather than as a universal certificate of causal transmission.

Conditional transfer entropy helps address common drivers, but conditioning on too many variables can make estimation unstable. Factor conditioning provides a pragmatic compromise: broad common variation is removed using a small number of latent or observed factors, after which pairwise directional relationships are estimated in the residual system. Sparse graphical selection can then identify which additional conditioning variables are most important. The simulation results suggest that this layered procedure recovers direct edges more reliably than either completely unconditional estimation or indiscriminate high-dimensional conditioning.

Network aggregation introduces another layer of choice. Weighted out-degree is intuitive, but a small number of strong edges and a large number of weak edges can produce similar totals. Reporting density, strength, concentration and reciprocity together gives a more complete account of system structure. During the simulated crisis, transmitter concentration rises before total density, showing that changes in network topology can contain information not visible in aggregate connectedness alone.

10. Practical Applications and Governance

In empirical surveillance, transfer-entropy networks should be treated as monitoring tools rather than automatic causal maps. A high out-degree node may reflect information leadership, market hours, liquidity differences or omitted common information. Economic interpretation therefore requires auxiliary evidence about timing, institutional structure and market design. Directional statistical dependence is a useful clue, but it is not a substitute for identification.

For operational use, the entire network pipeline should be version controlled: data transformations, filtering models, lag rules, discretization or neighbour parameters, surrogate generation and multiple-testing correction. Small undocumented changes can alter the number of significant edges. Reproducibility is especially important when network signals are used for policy communication or risk limits.

Because network estimates can be noisy, alert thresholds should be calibrated on persistent changes rather than one-window spikes. One approach is to require simultaneous increases in density, transmitter concentration and conditional edge strength. The controlled results support such multi-indicator confirmation, since the genuine crisis regime produces coordinated changes across several network statistics whereas finite-sample noise tends to affect them idiosyncratically.

Comparison with established connectedness models should be routine. If transfer entropy identifies a new transmission channel, researchers should ask whether the same relationship appears in nonlinear Granger models, quantile regressions or state-dependent VARs. Convergent evidence strengthens interpretation; disagreement can reveal which aspect of dependence drives the result.

11. Limitations and Future Research

The simulation uses a modest number of nodes, while real financial systems may contain hundreds or thousands of instruments. Direct nonparametric transfer-entropy estimation is unlikely to scale to such dimensions without structure. Future research should investigate factor-assisted, sparse and graph-neural estimators while preserving the interpretability that makes information-theoretic measures attractive.

Another limitation is the use of stationary regimes separated by known transitions. Real contagion may evolve continuously and may involve feedback loops that alter the data-generating process itself. Online estimators and change-point methods could provide a more realistic treatment of evolving information flow. High-frequency applications also require careful handling of asynchronous trading and market microstructure noise.

Finally, predictive information does not establish structural causality. Combining transfer entropy with natural experiments, institutional timing restrictions or structural identification would be an important next step. Such

integration could move entropy networks from descriptive systemic-risk monitoring toward more credible analysis of transmission mechanisms.

References

- Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.
- Aas, K., Czado, C., Frigessi, A., & Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2), 182–198.
- Bedford, T., & Cooke, R. M. (2002). Vines—A new graphical model for dependent random variables. *Annals of Statistics*, 30(4), 1031–1068.
- Czado, C. (2019). *Analyzing Dependent Data with Vine Copulas*. Springer.
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions. *Journal of Econometrics*, 182(1), 119–134.
- Embrechts, P., McNeil, A., & Straumann, D. (2002). Correlation and dependence in risk management. In M. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*. Cambridge University Press.
- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339–350.
- Joe, H. (2014). *Dependence Modeling with Copulas*. CRC Press.
- McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools* (2nd ed.). Princeton University Press.
- Nelsen, R. B. (2006). *An Introduction to Copulas* (2nd ed.). Springer.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2), 527–556.
- Dimpfl, T., & Peter, F. J. (2013). Using transfer entropy to measure information flows between financial markets. *Studies in Nonlinear Dynamics & Econometrics*, 17(1), 85–102.
- Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. *Physical Review E*, 69(6), 066138.
- Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85(2), 461–464.
- Vicente, R., Wibral, M., Lindner, M., & Pipa, G. (2011). Transfer entropy—A model-free measure of effective connectivity for the neurosciences. *Journal of Computational Neuroscience*, 30, 45–67.
- Barnett, L., Barrett, A. B., & Seth, A. K. (2009). Granger causality and transfer entropy are equivalent for Gaussian variables. *Physical Review Letters*, 103(23), 238701.
- Bossomaier, T., Barnett, L., Harré, M., & Lizier, J. T. (2016). *An Introduction to Transfer Entropy*. Springer.

Declarations

Funding: No external funding was used for the preparation of this manuscript.

Conflict of interest: The authors declare no competing interests.

Data and code availability: The numerical illustrations use simulated data and deterministic random seeds so that the analyses can be reproduced. Code and generated data may be deposited with the journal as supplementary material.



Journal of Dependence Modelling in Finance, Insurance and Demography

Open-Access Journal for Risk, Insurance, Survival and Demographic Analytics

No submission fees • No publication fees • No APCs during the founding phase



ARTICLE TYPE: RESEARCH ARTICLE

Tail Dependence in Insurance Losses: A Vine-Copula Approach to Multiline Risk Aggregation

Thomas Weber and Claire Dubois

Independent Researchers, Munich, Germany and Paris, France

Received:	7 June 2026	Accepted:	2 September 2026
Published:			

Abstract

Accurate multiline insurance capital modelling requires more than marginal severity estimation because aggregate solvency risk is highly sensitive to dependence in the extreme tail. This paper develops a vine-copula framework for aggregating simulated property, liability, motor, cyber and catastrophe losses. Heavy-tailed marginals are fitted separately and coupled through Gaussian, Student-t and regular-vine dependence structures. The experiment embeds asymmetric tail relationships between catastrophe and property losses and nonlinear dependence between cyber and liability losses, allowing competing aggregation methods to be assessed against a known data-generating process. Results show that independence and Gaussian dependence materially understate the 99.5% aggregate Value-at-Risk and Expected Shortfall, whereas Student-t and vine models provide substantially better tail calibration. The regular-vine specification captures heterogeneous pairwise dependence and delivers the most accurate diversification benefit at Solvency-II-type confidence levels. Stress tests, bootstrap intervals and alternative marginal fits confirm that dependence uncertainty becomes more important as the capital quantile increases. The study demonstrates why flexible dependence models are essential for multiline insurance risk aggregation and provides a reproducible framework for validating capital models before application to proprietary claims data.

Keywords: Insurance risk; vine copulas; tail dependence; VaR; Expected Shortfall; risk aggregation

JEL Classification: C46; C58; G22; G32

1. Introduction

Insurance groups aggregate losses across business lines that differ markedly in frequency, severity, exposure and claims development. Diversification can reduce required capital, but only if the joint distribution is represented adequately. Assuming independence is usually unrealistic, while using a single linear correlation matrix can be particularly misleading when large claims arrive jointly. Extreme weather, litigation shocks, cyber incidents and macroeconomic stress may create dependence that is weak in the center of the distribution but strong in the tail.

Copula models are attractive for insurance aggregation because they allow each marginal loss distribution to be estimated independently of the dependence structure. This is valuable when one line is well described by a lognormal body with a generalized Pareto tail and another by a gamma or Burr distribution. The dependence layer can then be selected according to empirical features such as lower- or upper-tail concentration, asymmetry and nonlinear association.

Pair-copula constructions and vines extend this flexibility to high-dimensional portfolios. Rather than imposing a single multivariate family on all business lines, a vine decomposes the joint density into bivariate copulas organized

through a sequence of trees. Different pair families can therefore be used for different loss relationships. This heterogeneity is economically plausible in insurance: catastrophe-property dependence need not resemble liability-cyber dependence, and motor losses may be nearly independent of both.

The paper evaluates the consequences of these modelling choices in a reproducible synthetic portfolio. Because the true dependence structure is known, aggregate VaR, Expected Shortfall and diversification benefit can be compared with their population targets. The exercise is designed to mimic capital-model validation rather than to estimate the risk of any named insurer.

2. Background: Insurance Aggregation and Copula Dependence

The aggregate annual loss is $S = \sum_j L_j$, where L_j is the loss in business line j . Regulatory and internal capital calculations commonly evaluate a high quantile of S , often supplemented by Expected Shortfall to capture the average loss beyond the quantile. Both measures depend on the full joint distribution. Marginal calibration alone is therefore insufficient.

Under elliptical copulas, dependence is summarized by correlation-type parameters, and the Student-t copula additionally generates symmetric tail dependence. Insurance portfolios often require greater flexibility because dependence can be asymmetric and pair-specific. Archimedean copulas such as Gumbel and Clayton provide one form of asymmetry, while vine constructions combine different bivariate copulas in a single multivariate model.

Aas et al. (2009) established the practical use of pair-copula constructions for multivariate dependence, and Czado (2019) provides a comprehensive treatment of vine-copula modelling. In insurance, the central advantage is not merely higher likelihood but the ability to represent heterogeneous tail relationships. Because solvency metrics are evaluated far into the tail, small improvements in joint-extreme calibration can have large effects on capital.

Model risk enters through both margins and dependence. A heavy-tailed marginal fit can dominate the level of capital, whereas copula choice influences the frequency with which large line-specific losses coincide. The paper therefore evaluates uncertainty from both sources and shows that their relative importance changes across confidence levels.

$$S = \sum_{j=1}^d L_j \quad (1)$$

$$VaR_\alpha(S) = \inf \{ s : P(S \leq s) \geq \alpha \} \quad (2)$$

$$ES_\alpha(S) = E[S | S > VaR_\alpha(S)] \quad (3)$$

3. Data-Generating Process and Marginal Models

The synthetic portfolio contains 5,000 annualized scenario observations for five lines: property, liability, motor, cyber and catastrophe. Marginal severities are generated from lognormal, Burr and generalized Pareto mixtures chosen to produce heterogeneous tail indexes. Property and catastrophe are the heaviest-tailed lines, while motor has the smallest coefficient of variation.

Dependence is generated by a regular vine. The first tree contains strong upper-tail dependence between property and catastrophe, moderate Student-t dependence between cyber and liability, and weak Gaussian links involving motor. Higher-tree copulas introduce a common macro-stress channel but remain deliberately weaker than the first-tree relationships. The structure therefore contains both strong local dependence and broader conditional association.

Marginal fitting is conducted without knowledge of the generating distributions. Candidate lognormal, gamma, Weibull, Burr and generalized Pareto-spliced models are compared by likelihood and tail diagnostics. Probability-integral-transform residuals are then used for copula estimation. This two-stage procedure mirrors common actuarial workflows and makes it possible to evaluate whether dependence estimates remain stable under reasonable marginal misspecification.

4. Dependence Models and Capital Estimation

Four aggregation models are considered: independence, Gaussian copula, multivariate Student-t copula and a regular vine. For the vine, pair families are selected from Gaussian, Student-t, Clayton, Gumbel and rotated variants using AIC, subject to a minimum improvement criterion that prevents overfitting weak pairs. The tree structure is chosen using absolute empirical Kendall's tau as the initial edge weight.

Capital is estimated by Monte Carlo simulation of one million annual aggregate losses from each fitted model. Reported metrics are VaR and Expected Shortfall at 95%, 99% and 99.5%, together with diversification benefit defined as one minus aggregate capital divided by the sum of stand-alone capitals. Tail calibration is assessed by comparing model-implied joint exceedance probabilities with the known simulation truth.

Uncertainty is evaluated through 500 block bootstrap replications of the scenario sample. Although the observations are generated independently at the annual level, block bootstrap is retained as a conservative procedure and as a template for extensions involving serially dependent operational or cyber-loss scenarios.

Table 1. Aggregate capital estimates and tail-calibration diagnostics

Model	VaR 99.5	ES 99.5	Diversification benefit	Tail calibration error
Independence	462	521	0.29	0.42
Gaussian	511	584	0.22	0.26
Student-t	548	636	0.17	0.1
Regular vine	556	648	0.16	0.05

5. Results

The independence model produces the lowest aggregate capital and the largest apparent diversification benefit. This result is mechanically attractive but incorrect because the generating process contains several strong dependence channels. At the 99.5% level, independence understates VaR by approximately 17% relative to the regular-vine benchmark and understates Expected Shortfall by an even larger margin. The error grows as the confidence level increases.

The Gaussian copula improves central dependence fit but remains weak in joint-tail calibration. Because asymptotic tail dependence is absent, simultaneous large losses are too infrequent. The Student-t copula performs substantially better because it allows joint extremes, yet its symmetric and homogeneous dependence structure is too restrictive for the portfolio. It tends to overstate some weak pairs and understate the specific property-catastrophe asymmetry.

The regular vine provides the best overall match to the known joint distribution. It selects an upper-tail-dependent family for property-catastrophe and a Student-t family for liability-cyber, while retaining Gaussian copulas for weak central relationships. This heterogeneity improves both joint-exceedance calibration and aggregate capital accuracy. The improvement over the Student-t model is moderate at 95% but becomes meaningful at 99.5%, illustrating that flexible dependence matters most where solvency decisions are made.

Diversification benefit declines as the dependence model becomes more realistic. This is not evidence against diversification; rather, it indicates that simplistic models overstate how much diversification is available in severe scenarios. In the generating model, diversification remains positive because not all lines are strongly dependent, but it is materially smaller than the independence calculation suggests.




Figure 1. Aggregate VaR under Gaussian and vine dependence




Figure 2. Pairwise dependence among business lines

6. Stress Testing and Robustness

A catastrophe stress doubles the frequency of upper-tail property-catastrophe co-exceedances while leaving marginal means approximately unchanged. The Gaussian model responds weakly because its dependence structure has no explicit tail parameter. Student-t capital increases, but the vine model reacts most strongly and most accurately because the stressed pair occupies a prominent position in the first tree.

Alternative marginal specifications show that dependence estimates are reasonably stable when candidate margins fit the upper tail adequately. When thin-tailed gamma margins are forced on property and catastrophe, however, all dependence models understate Expected Shortfall. This confirms that flexible dependence cannot compensate for severely misspecified margins. Capital validation must therefore consider the two layers jointly.

Bootstrap intervals widen rapidly beyond the 99% level. Dependence-family uncertainty is especially visible for pairs with few joint extremes. Model averaging across plausible vine specifications reduces variance slightly and may be useful in practice when no single pair family dominates decisively.

7. Discussion

The principal practical implication is that a single dependence number is rarely sufficient for multiline insurance portfolios. Correlation can summarize central association, but solvency capital is determined by the behavior of the joint distribution in regions containing few observations. Models that permit heterogeneous tail dependence are therefore preferable when strong economic reasons exist for certain lines to suffer together.

Vine models are not costless. They introduce structural and family-selection risk, require substantial data or expert constraints, and can be computationally intensive in high dimensions. Their flexibility should therefore be disciplined through out-of-sample checks, simplified structures and sensitivity analysis. The objective is not to maximize likelihood mechanically but to represent economically meaningful dependence with adequate parsimony.

8. Conclusion

This paper shows that dependence choice can materially affect high-quantile insurance capital even when marginal distributions are held fixed. Independence and Gaussian dependence understate joint extremes in the synthetic portfolio, whereas Student-t and regular-vine models provide better tail calibration. The vine specification performs best because it accommodates heterogeneous pairwise dependence and asymmetry.

For internal models and solvency applications, the results support a validation strategy that combines marginal diagnostics, pairwise tail checks, aggregate backtesting and stress testing of the dependence layer. Flexible copula models are most valuable when they are used transparently and when their additional complexity is justified by demonstrable improvements in tail-risk calibration.

9. Capital Model Validation Considerations

An internal capital model should be judged by more than point estimates of VaR or Expected Shortfall. The dependence component must be validated through multiple views of the data, including pairwise rank measures, joint exceedance counts, conditional tail plots and aggregate scenario backtests. A vine model can fit the sample extremely well yet still produce unstable tail extrapolation if key pair families are selected from only a handful of extreme observations. Validation should therefore focus on both goodness of fit and parameter uncertainty.

The treatment of expert judgment is particularly important in insurance. Some dependence relationships are driven by mechanisms that have not occurred frequently enough to be estimated reliably from historical claims. Catastrophe and property losses, cyber and business interruption, or liability and inflation may share plausible stress channels that are only weakly represented in the sample. Expert constraints can be incorporated by restricting copula families, imposing conservative tail parameters or stress-testing selected pair links. Such judgments should be transparent and separable from data-driven estimates.

The regular-vine structure also has governance implications. A complex vine may contain many pair copulas that contribute almost nothing to aggregate capital. Pruning weak edges can improve stability and make the model

easier to validate. The appropriate criterion is not simply statistical significance but materiality: if replacing an edge with independence has negligible effect on solvency metrics across relevant stresses, retaining a highly parameterized family may not be justified.

Monte Carlo error is another practical issue at 99.5% confidence. Capital differences of a few percentage points can be obscured if the simulation size is too small. Repeated simulation, variance reduction and convergence diagnostics should therefore accompany model comparison. In the present study, one million draws are sufficient to make numerical error small relative to dependence-model uncertainty.

10. Practical Implications for Multiline Portfolios

The strongest practical message is that diversification benefits should be stress aware. Diversification observed in ordinary years can disappear when several lines are exposed to the same macroeconomic, legal or catastrophe driver. A dependence model that allows tail concentration can translate this intuition into capital, but the result should be accompanied by scenario narratives explaining why the lines move together.

Vine structures are particularly useful when a small subset of pairs dominates aggregate risk. In the synthetic portfolio, property-catastrophe and cyber-liability relationships matter much more than motor-related pairs. This suggests a targeted modelling strategy in which the strongest economic channels receive detailed tail modelling and weak pairs remain parsimonious. Such targeting can produce most of the benefit of a flexible multivariate model without unnecessary complexity.

Risk managers should also evaluate diversification at several confidence levels. A portfolio can appear strongly diversified at the 95% level but much less diversified at 99.5%. Reporting the entire capital curve rather than a single regulatory quantile reveals whether diversification erodes gradually or abruptly in the tail. This information is valuable for reinsurance design and for assessing the resilience of capital buffers beyond the minimum regulatory standard.

Reinsurance can change the dependence structure of retained losses even when underlying gross losses are unchanged. Non-proportional covers truncate or layer marginal tails and may alter which lines dominate extreme aggregate outcomes. Extending the framework to gross and net losses jointly would therefore be especially relevant for practical capital optimization.

11. Limitations and Future Research

The present analysis uses annualized aggregate losses and does not model claim-count dependence separately from severity dependence. In some lines, common shocks may increase frequencies rather than severities, and a frequency-severity copula model could provide a richer description. Claim development and reserving uncertainty are also omitted, although they can interact with underwriting risk over a multi-year capital horizon.

Future research should examine dynamic dependence in insurance portfolios, particularly for cyber and climate-related risks where exposure and technology evolve rapidly. Static copulas assume that the dependence structure is stable across the observation period. Regime-switching vines or covariate-dependent copulas could link dependence parameters to climate indexes, inflation, interest rates or cyber-threat indicators.

Finally, capital model comparison should be extended to decision outcomes such as reinsurance purchase, portfolio limits and risk-adjusted pricing. A dependence model is valuable not because it is statistically sophisticated but because it improves the quality and stability of those decisions.

References

- Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.
- Aas, K., Czado, C., Frigessi, A., & Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2), 182–198.


- Bedford, T., & Cooke, R. M. (2002). Vines—A new graphical model for dependent random variables. *Annals of Statistics*, 30(4), 1031–1068.
- Czado, C. (2019). *Analyzing Dependent Data with Vine Copulas*. Springer.
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions. *Journal of Econometrics*, 182(1), 119–134.
- Embrechts, P., McNeil, A., & Straumann, D. (2002). Correlation and dependence in risk management. In M. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*. Cambridge University Press.
- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339–350.
- Joe, H. (2014). *Dependence Modeling with Copulas*. CRC Press.
- McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools* (2nd ed.). Princeton University Press.
- Nelsen, R. B. (2006). *An Introduction to Copulas* (2nd ed.). Springer.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2), 527–556.
- Klugman, S. A., Panjer, H. H., & Willmot, G. E. (2012). *Loss Models: From Data to Decisions* (4th ed.). Wiley.
- Embrechts, P., Klüppelberg, C., & Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer.
- Dissmann, J., Brechmann, E. C., Czado, C., & Kurowicka, D. (2013). Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics & Data Analysis*, 59, 52–69.
- Genest, C., & Favre, A.-C. (2007). Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of Hydrologic Engineering*, 12(4), 347–368.
- Frees, E. W., & Valdez, E. A. (1998). Understanding relationships using copulas. *North American Actuarial Journal*, 2(1), 1–25.

Declarations

Funding: No external funding was used for the preparation of this manuscript.

Conflict of interest: The authors declare no competing interests.

Data and code availability: The numerical illustrations use simulated data and deterministic random seeds so that the analyses can be reproduced. Code and generated data may be deposited with the journal as supplementary material.



ARTICLE TYPE: RESEARCH ARTICLE

Cross-Population Mortality Dependence and Longevity Risk: A Multivariate Copula Framework

George Papadopoulos and Anna Lindström

Independent Researchers, Thessaloniki, Greece and Stockholm, Sweden

Received:	19 June 2026	Accepted:	15 September 2026
Published:			

Abstract

Longevity-risk diversification across populations depends critically on whether mortality improvements remain correlated in adverse scenarios. This paper develops a multivariate copula framework for cross-population mortality dependence using a simulated panel of six populations observed over ages 55–90 and 45 calendar years. Population-specific mortality trends are estimated with a common-factor model, after which dependence among standardized mortality innovations is modelled using Gaussian, Student-t and factor-copula specifications. The data-generating process includes a shared longevity trend, region-specific factors and episodic joint mortality-improvement shocks, allowing the effect of tail dependence on annuity liabilities to be measured against a known benchmark. Results show that Gaussian dependence overstates diversification when extreme longevity improvements occur simultaneously, whereas the Student-t and factor-copula models reproduce joint longevity shocks more accurately. The impact is modest for one-year mortality forecasts but substantial for the upper tail of 30-year annuity present values. Bootstrap and rolling-window analyses indicate that dependence uncertainty becomes a material component of longevity capital when portfolios span highly related populations. The study offers a reproducible framework for integrating mortality forecasting and dependence modelling in longevity-risk management.

Keywords: Mortality; longevity risk; multi-population models; copulas; annuities; dependence

JEL Classification: C32; C53; G22; J11

1. Introduction

Longevity risk arises when mortality improves more rapidly than expected, increasing the duration and value of pension and annuity liabilities. Because insurers and pension funds often hold exposures to several demographic groups, cross-population diversification is potentially important. That diversification, however, depends on how mortality innovations co-move. If extreme improvements are shared across populations, a portfolio that appears diversified under Gaussian dependence can remain highly exposed in the longevity tail.

Multi-population mortality models address common trends explicitly. Extensions of the Lee–Carter framework, including coherent and common-factor approaches, allow populations to share long-run mortality dynamics while retaining population-specific deviations. These models are well suited to forecasting mean mortality paths, but the residual dependence that remains after fitting the common trend can still matter for joint liability distributions.

The present paper connects multi-population mortality modelling with copula-based dependence analysis. The objective is not to propose another central mortality trend model; rather, it is to examine how the dependence of

mortality innovations affects longevity-risk aggregation. A synthetic setting is used so that the true common trend, regional factors and episodic joint longevity shocks are known.

The study contributes by separating three sources of cross-population co-movement: deterministic common trend, persistent regional dependence and tail co-movement in innovations. This distinction helps explain why a mortality model can fit historical means well yet still understate portfolio tail risk if the dependence layer is too restrictive.

2. Mortality Forecasting and Cross-Population Dependence

The Lee–Carter model represents log mortality as an age pattern plus an age-specific sensitivity multiplied by a period index. Li and Lee (2005) extended this idea to coherent mortality forecasting, allowing related populations to share a common mortality development. Other approaches use product-ratio decompositions, common factors or multivariate time-series structures. The common objective is to prevent implausible long-run divergence among populations with related demographic experience.

Dependence matters beyond the conditional mean. Suppose two populations share a long-run period factor but also experience unusually strong mortality improvements in the same years. If those joint innovations are more frequent than Gaussian dependence implies, the upper tail of annuity values will be understated. This issue is conceptually similar to joint loss modelling in insurance, except that the risky event is unexpectedly low mortality rather than high claim severity.

Copulas provide a natural separation between population-specific innovation distributions and their dependence. A Student-t copula permits symmetric joint extremes, while factor copulas provide a scalable way to represent dependence among many populations. For longevity portfolios, dependence can be assessed directly on mortality innovations or on derived liability shocks. The former is statistically more transparent, whereas the latter is closer to the economic quantity of interest.

The paper adopts both perspectives. Copulas are fitted to mortality innovations and then propagated through annuity valuation to evaluate portfolio capital. This allows the statistical dependence model to be judged by its financial consequences rather than by goodness-of-fit alone.

$$\log m_{x,t}^{(i)} = a_x^{(i)} + b_x^{(i)} k_t^{(i)} + \varepsilon_{x,t}^{(i)} \quad (1)$$

$$k_t^{(i)} = \beta_i K_t + \gamma_i R_{g(i),t} + u_t^{(i)} \quad (2)$$

$$PV^{(i)} = \sum_{s=1}^T v^s \square_s p_x^{(i)} \quad (3)$$

3. Synthetic Mortality Panel and Estimation Strategy

The synthetic panel contains six populations grouped into three regions. Ages 55–90 are observed for 45 years. Log mortality follows a declining common period factor, age-specific sensitivities and a region factor. Population-specific innovations are heavy-tailed and include five system-wide longevity-improvement shocks. These shocks lower mortality simultaneously in all populations but with heterogeneous intensity.

Population models are estimated in two stages. First, age schedules and period factors are extracted using singular-value decomposition under standard identification constraints. Second, period-factor dynamics are modelled with random walks with drift and regional adjustments. Standardized innovations are then transformed to pseudo-observations for copula estimation.

Three dependence models are compared: Gaussian copula, Student-t copula and a one-factor Student-t copula with population-specific loadings. Forecasts are generated for 30 years, and mortality paths are converted to cohort survival probabilities. A stylized immediate annuity at age 65 is then valued under each simulated path using a fixed discount curve.

4. Evaluation Design

Model evaluation is performed at both the mortality and liability levels. Mortality diagnostics include pairwise Kendall dependence, joint lower-innovation exceedances and calibration of the number of populations simultaneously experiencing unusually strong mortality improvement. Liability diagnostics include the mean annuity present value, 95% and 99.5% liability quantiles and portfolio diversification benefit across populations.

Forecast evaluation distinguishes one-year calibration from long-horizon liability risk. This is important because cross-population dependence has relatively little effect on individual one-year point forecasts but can accumulate substantially when many future survival probabilities are combined in a long-duration annuity.

Rolling-window fits are used to assess dependence stability. Bootstrap confidence intervals are obtained by resampling years in blocks to preserve low-frequency mortality dependence. Alternative age ranges and discount rates are evaluated as robustness checks.

Table 1. Dependence-model implications for joint longevity risk

Model	Joint longevity-shock error	99.5% liability index	Diversification benefit
Gaussian copula	0.31	1.0	0.142
Student-t copula	0.12	1.046	0.111
Factor-t copula	0.07	1.058	0.103

5. Results

The estimated common mortality trend is recovered accurately by all specifications because the central trend model is identical. Differences emerge in the residual dependence layer. Pairwise Kendall dependence is moderate for populations within the same region and lower across regions, but the simulated system-wide longevity shocks create more simultaneous extremes than a Gaussian copula predicts.

The Gaussian model therefore appears adequate in ordinary years yet materially understates the frequency of years in which four or more populations experience extreme mortality improvements simultaneously. The Student-t copula improves this diagnostic, while the factor-t model performs best because it combines heavy-tailed joint shocks with heterogeneous population loadings.

At the liability level, mean annuity values are nearly identical across dependence models. This is expected because dependence does not materially change marginal mortality forecasts. The difference appears in portfolio quantiles. The 99.5% liability measure under the factor-t model is approximately 5.8% above the Gaussian benchmark, and diversification benefit falls from 14.2% to 10.3%. Thus, the principal effect of dependence modelling is on capital and diversification rather than on expected liabilities.

The effect grows with forecast horizon. At five years the capital difference is small, while at 20–30 years it becomes material because common longevity shocks propagate through many survival probabilities. This finding is particularly relevant for long-dated annuity and pension exposures.




Figure 1. Estimated common mortality period factor




Figure 2. Pairwise innovation dependence

6. Robustness and Sensitivity

Changing the age range to 60–89 slightly increases dependence estimates because older-age mortality innovations share a stronger common component in the simulation. The ranking of copula models is unchanged. Using a lower discount rate increases all annuity values but also magnifies the absolute capital difference between Gaussian and heavy-tailed dependence.

Block-bootstrap intervals for 99.5% liability capital are wider than intervals for mean liabilities, confirming that dependence uncertainty is concentrated in the tail. The factor-t model remains above the Gaussian benchmark in more than 90% of bootstrap replications. A placebo simulation with no system-wide longevity shocks removes most of the capital difference, demonstrating that the effect is genuinely driven by joint extremes rather than by ordinary correlation.

Alternative factor specifications with two global factors yield only minor improvements because the generating process contains one dominant global source plus regional components. In real data, however, model selection should allow for richer structures and for time variation in dependence.

7. Discussion

The results show why mortality diversification should not be inferred solely from historical correlation. Common trend models can eliminate long-run divergence while leaving residual tail dependence that remains important for solvency. A liability portfolio may therefore be well diversified in expectation but less diversified in extreme longevity scenarios.

The analysis also provides a practical modelling hierarchy. First, coherent mortality forecasting should capture common long-run trends. Second, residual innovations should be checked for cross-population dependence and tail concentration. Third, the dependence model should be validated using liability-level diagnostics because statistically small differences can become economically meaningful after long-horizon survival aggregation.

The synthetic design is intentionally stylized and does not include migration, cause-of-death interactions, pandemic shocks or cohort effects. Those features could alter both marginal forecasts and dependence. Nevertheless, the framework is sufficiently general to accommodate richer mortality models because the copula layer is applied to standardized innovations rather than tied to a specific mortality specification.

8. Conclusion

Cross-population dependence is a material component of longevity-risk aggregation. In the controlled panel studied here, Gaussian dependence reproduces ordinary co-movement but understates simultaneous extreme mortality improvements and therefore overstates diversification. Student-t and factor-t copulas provide more accurate joint-shock calibration and produce higher, more realistic tail liability measures.

The findings support an integrated approach in which mortality forecasting and dependence modelling are evaluated jointly. For insurers and pension funds with exposures to related populations, tail dependence should be treated as a distinct source of model risk rather than assumed to be captured automatically by a common mortality trend.

9. Dependence Modelling in Longevity Portfolios

Cross-population longevity dependence can arise from several distinct sources: shared medical technology, public-health policy, socioeconomic convergence, environmental exposure and common measurement practices. A statistical dependence parameter does not reveal which mechanism dominates. Nevertheless, separating common deterministic or stochastic trends from residual joint extremes is useful because the two components have different implications for hedging. A common trend may be partially hedgeable with broad longevity instruments, while residual population-specific tail dependence may remain as basis risk.

The choice of dependence scale is also important. Modelling innovations in period indexes is parsimonious, but it assumes that age-specific dependence is adequately summarized by the factor representation. Direct dependence

modelling across ages would be more flexible but vastly more dimensional. A practical compromise is to model a small number of mortality factors and then test residual age-specific dependence at economically important ages. This hierarchical approach mirrors factor-copula strategies in finance.

Liability aggregation can amplify small dependence differences because survival probabilities are products of many annual terms. Even if one-year mortality innovations differ only slightly across copula models, persistent joint improvements can shift long-horizon annuity quantiles. This nonlinearity explains why liability-level validation is essential. Statistical fit at the mortality-innovation level should be considered necessary but not sufficient for longevity-capital applications.

Hedging effectiveness is similarly dependence sensitive. A hedge based on another national population will perform well only if the relevant innovations remain dependent in stress states. A Gaussian correlation estimated over the full sample can overstate hedge reliability when dependence weakens or changes shape in extreme longevity scenarios. Tail-sensitive dependence measures can therefore contribute directly to basis-risk assessment.

10. Model Risk, Validation and Governance

Longevity models are long-horizon models, so parameter and structural uncertainty often dominate ordinary process risk. Dependence parameters add another source of uncertainty that should be propagated into capital rather than fixed at point estimates. Bootstrap distributions in this paper show that tail-capital intervals widen meaningfully when copula parameters are re-estimated, especially for smaller population groups.

Validation should include historical backtesting where feasible, but the scarcity of extreme longevity events limits purely empirical assessment. Scenario analysis is therefore indispensable. Examples include accelerated cardiovascular improvement, a medical breakthrough affecting several countries, or divergent improvements caused by unequal adoption of treatment. Dependence models should be tested against such narratives to determine whether they generate plausible cross-population responses.

Model governance should also distinguish demographic judgment from statistical estimation. Long-run convergence assumptions, population grouping and cohort-treatment choices often reflect substantive demographic reasoning. These choices affect residual dependence and should be documented before copula calibration. Otherwise, the dependence model may absorb structure that belongs in the mortality trend model.

For portfolio applications, reporting only a single best-estimate dependence model is insufficient. Challenger specifications, parameter stress and alternative population groupings should be maintained. The spread in liability capital across plausible models provides a transparent measure of dependence-model risk.

11. Limitations and Future Research

The synthetic panel does not include cause-of-death information, migration or explicit cohort effects. Each of these can generate additional cross-population dependence. Cause-specific improvements may spread internationally at different speeds, while cohort effects can be shared among culturally or geographically related populations. Extending the copula layer to latent cause-specific or cohort factors would be a natural next step.

Pandemic and disaster mortality also require special treatment because they represent adverse mortality shocks rather than longevity improvements. A comprehensive life-insurance portfolio may face both short-term mortality catastrophe risk and long-term longevity risk. Dependence between these components across populations is not necessarily symmetric and may require regime-dependent or mixture copulas.

Future work could also connect dependence modelling to tradable longevity instruments. Evaluating hedge ratios and basis risk under Gaussian versus tail-dependent population relationships would provide a direct economic test of model relevance. Such an extension would move the framework from liability measurement toward active longevity-risk management.

References

Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.


- Aas, K., Czado, C., Frigessi, A., & Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2), 182–198.
- Bedford, T., & Cooke, R. M. (2002). Vines—A new graphical model for dependent random variables. *Annals of Statistics*, 30(4), 1031–1068.
- Czado, C. (2019). *Analyzing Dependent Data with Vine Copulas*. Springer.
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions. *Journal of Econometrics*, 182(1), 119–134.
- Embrechts, P., McNeil, A., & Straumann, D. (2002). Correlation and dependence in risk management. In M. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*. Cambridge University Press.
- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339–350.
- Joe, H. (2014). *Dependence Modeling with Copulas*. CRC Press.
- McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools* (2nd ed.). Princeton University Press.
- Nelsen, R. B. (2006). *An Introduction to Copulas* (2nd ed.). Springer.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2), 527–556.
- Cairns, A. J. G., Blake, D., & Dowd, K. (2006). A two-factor model for stochastic mortality with parameter uncertainty. *Journal of Risk and Insurance*, 73(4), 687–718.
- Lee, R. D., & Carter, L. R. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, 87(419), 659–671.
- Li, N., & Lee, R. (2005). Coherent mortality forecasts for a group of populations: An extension of the Lee–Carter method. *Demography*, 42(3), 575–594.
- Hyndman, R. J., & Ullah, M. S. (2007). Robust forecasting of mortality and fertility rates: A functional data approach. *Computational Statistics & Data Analysis*, 51(10), 4942–4956.
- Haberman, S., & Renshaw, A. (2011). A comparative study of parametric mortality projection models. *Insurance: Mathematics and Economics*, 48(1), 35–55.
- Villegas, A. M., Kaishev, V. K., & Millossovich, P. (2018). StMoMo: An R package for stochastic mortality modeling. *Journal of Statistical Software*, 84(3), 1–38.

Declarations

Funding: No external funding was used for the preparation of this manuscript.

Conflict of interest: The authors declare no competing interests.

Data and code availability: The numerical illustrations use simulated data and deterministic random seeds so that the analyses can be reproduced. Code and generated data may be deposited with the journal as supplementary material.



ARTICLE TYPE: METHODOLOGICAL PAPER

Entropy, Copulas and Tail Connectedness: A Unified Framework for Nonlinear Dependence Measurement
Daniel Moretti and Sophie Laurent

Independent Researchers, Rome, Italy and Brussels, Belgium

Received:	1 July 2026	Accepted:	28 September 2026
Published:			

Abstract

This paper develops a unified simulation framework for comparing linear correlation, rank dependence, copula tail dependence, mutual information and transfer entropy under a broad set of nonlinear financial dependence structures. The objective is to clarify which measures answer different dependence questions rather than to identify a single universally superior statistic. Six controlled designs are considered: linear Gaussian association, nonlinear symmetric dependence, asymmetric tail dependence, regime switching, directional nonlinear transmission and common-driver confounding. Each measure is evaluated in terms of sensitivity, invariance, finite-sample stability and interpretability. The results show that Pearson correlation is efficient in the Gaussian-linear design but can be close to zero under strong nonlinear dependence; Spearman and Kendall measures improve monotonic robustness but remain nondirectional; copula tail coefficients isolate extreme co-movement; mutual information detects broad nonlinear association; and transfer entropy identifies directional predictive information when accompanied by significance testing and conditioning. A composite dependence dashboard is proposed that reports complementary measures rather than collapsing them into a single score. Monte Carlo experiments across sample sizes and heavy-tail parameters demonstrate the trade-off between flexibility and estimator variance. The framework provides a systematic methodological guide for dependence analysis in finance, insurance and demography.

Keywords: Entropy; copulas; mutual information; transfer entropy; tail dependence; nonlinear dependence

JEL Classification: C14; C46; C58

1. Introduction

Dependence is not a single statistical object. Researchers may be interested in contemporaneous linear association, monotonic co-movement, general nonlinear dependence, joint tail behavior or directional predictive information. These questions are related but they are not equivalent. Nevertheless, empirical studies often report one dependence statistic and interpret it as if it summarized the entire relationship.

This problem is especially important in financial and actuarial applications because dependence changes across the distribution and through time. A pair of assets can have modest correlation but strong lower-tail association; two insurance risks can be nonlinearly dependent with little monotonic association; and one market can transmit predictive information to another even when contemporaneous correlation is weak. No scalar statistic captures all of these features simultaneously.

The paper develops a controlled framework in which commonly used measures are evaluated according to the question they are designed to answer. Pearson correlation, Spearman rho, Kendall tau, copula tail dependence,

mutual information and transfer entropy are compared under six data-generating designs. Because the true structure is known, false negatives and false interpretations can be identified directly.

The central argument is not that flexible measures should always replace simple ones. Pearson correlation is statistically efficient and highly interpretable when the true relationship is approximately linear and elliptical. More flexible measures become valuable when the data depart from that setting, but they typically require larger samples and more tuning. A good dependence analysis should therefore match the measure to the economic question and report complementary diagnostics when several aspects of dependence matter.

2. Dependence Concepts

Pearson correlation measures standardized linear covariance. It is zero for uncorrelated variables but not, in general, for independent variables. Spearman rho and Kendall tau are rank-based and therefore invariant to strictly monotonic transformations. They capture monotonic association more broadly than Pearson correlation but can still fail when dependence is non-monotonic.

Copulas represent the full dependence structure once margins are removed. Tail-dependence coefficients summarize the probability that one variable is extreme conditional on the other being extreme. They are therefore local measures focused on a region of the joint distribution. This makes them valuable for stress analysis but inappropriate as complete summaries of central dependence.

Mutual information is zero if and only if two variables are independent under regularity conditions. It can therefore detect nonlinear and non-monotonic dependence missed by correlation. Its flexibility comes at a statistical cost: estimating a probability density or local-neighbour structure introduces finite-sample bias and tuning choices.

Transfer entropy extends conditional mutual information to time-indexed variables. It measures whether the past of a source reduces uncertainty about the future of a target after conditioning on the target's own history. Unlike mutual information, it is directional. However, a positive estimate is not automatically evidence of a direct causal channel because common drivers and finite-sample effects can generate apparent information flow.

$$\rho_p = \frac{\text{Cov}(X, Y)}{\sigma_X \sigma_Y} \quad (1)$$

$$I(X; Y) = \iint f_{XY}(x, y) \log \left[\frac{f_{XY}(x, y)}{f_X(x) f_Y(y)} \right] dx dy \quad (2)$$

$$\lambda_L = \lim_{q \rightarrow 0^+} P[V \leq q | U \leq q] \quad (3)$$

$$T_{X \rightarrow Y} = I(X_t^{(k)}; Y_{t+1} | Y_t^{(l)}) \quad (4)$$

3. Monte Carlo Designs

Design A is bivariate Gaussian with correlations ranging from zero to 0.9. It represents the environment in which Pearson correlation should perform best. Design B uses $Y = X^2 + \epsilon$ with symmetric X , producing strong nonlinear dependence but correlation close to zero. Design C uses a Student-t copula with moderate central dependence and strong symmetric tail dependence. Design D uses a Clayton-type lower-tail relationship to create asymmetry.

Design E is dynamic: X_t influences Y_{t+1} through a nonlinear threshold function, while contemporaneous correlation remains weak. This design targets transfer entropy. Design F introduces a common driver Z_t that affects both X and Y but includes no direct X -to- Y channel. It is used to evaluate whether unconditional directional measures confuse shared information with direct transmission.

Sample sizes of 250, 500, 1,000 and 5,000 are considered. For entropy-based estimators, histogram, kernel and k-nearest-neighbour variants are evaluated. Copula tail coefficients are estimated both parametrically and empirically. Each configuration is repeated 1,000 times, and performance is summarized by bias, root-mean-square error, rank stability and detection probability.

Table 1. Illustrative dependence measures across controlled designs

Design	Pearson	Kendall	Mutual information	Tail coefficient	Transfer entropy
A Linear Gaussian	0.82	0.61	0.54	0.0	0.01
B Nonlinear U-shape	0.03	0.01	0.71	0.02	0.01
C t-tail	0.48	0.34	0.39	0.29	0.02
D Lower-tail asymmetric	0.42	0.31	0.36	0.36	0.02
E Directional nonlinear	0.08	0.05	0.18	0.03	0.17
F Common driver	0.46	0.32	0.33	0.09	0.08

4. Results: What Each Measure Detects

In the Gaussian-linear design, Pearson correlation is the most efficient summary. Rank measures and mutual information all detect dependence, but they offer little additional information and exhibit slightly higher estimation variance at small samples. This result is important because methodological flexibility is not automatically beneficial; when the linear model is correct, simple estimators are hard to improve upon.

The nonlinear U-shaped design reverses the conclusion. Pearson, Spearman and Kendall measures are all close to zero because positive and negative slopes cancel across the support. Mutual information is large and stable once the sample exceeds approximately 500 observations. The result illustrates the distinction between absence of correlation and absence of dependence.

In the Student-t and lower-tail designs, rank dependence identifies general association but does not describe the concentration of joint extremes. Tail coefficients and copula families provide the relevant additional information. The Student-t design has similar lower- and upper-tail coefficients, whereas the asymmetric design has a pronounced lower-tail coefficient and negligible upper-tail dependence. Reporting a single correlation coefficient would conceal this difference.

The directional nonlinear design is intentionally difficult for contemporaneous measures. Mutual information detects some association, but only transfer entropy identifies that the source history contributes predictive information about the target future. In the common-driver design, unconditional transfer entropy produces false directional links, while conditional transfer entropy falls close to zero after conditioning on the driver. This confirms that directionality and directness are distinct concepts.




Figure 1. Finite-sample performance of mutual-information estimation




Figure 2. Similarity of dependence diagnostics across Monte Carlo designs

5. A Unified Dependence Dashboard

The simulation results motivate a dashboard rather than a universal dependence score. The proposed dashboard contains five blocks: linear association, rank association, tail dependence, general nonlinear dependence and directional information flow. Each block answers a distinct question, and interpretation begins by identifying which question matters for the application.

For portfolio diversification, correlation and rank dependence provide useful central summaries, while tail dependence should be added when extreme co-losses are important. For insurance aggregation, copula tail diagnostics are often more informative than transfer entropy because annual losses may not have a meaningful temporal direction. For systemic-risk networks, transfer entropy or variance-decomposition connectedness may be appropriate because direction is central. For demographic applications, dependence in long-run mortality innovations may be more relevant than high-frequency directional flow.

The dashboard is deliberately non-aggregative. Collapsing heterogeneous dependence concepts into a single score can obscure disagreement among measures. If Pearson correlation is low but mutual information is high, that discrepancy is itself substantive evidence of nonlinear structure. Similarly, high correlation with zero tail dependence conveys a different risk profile from moderate correlation with strong tail dependence.

6. Finite-Sample and Estimation Issues

Estimator variance is the principal cost of flexible dependence measures. Mutual-information estimates are unstable in small samples when binning is fine or kernel bandwidths are small. K-nearest-neighbour estimators perform best overall in the Monte Carlo study, but their tuning parameter still affects bias. Tail-dependence estimates are similarly unstable when the threshold is too extreme because only a small number of joint observations remain.

Transfer entropy is even more data demanding because joint histories enlarge the state space. Increasing lag order can rapidly create sparse cells or high-dimensional neighbourhoods. Significance testing with surrogate data is therefore essential. The simulations show that an apparently positive transfer-entropy estimate can persist under the null unless its finite-sample distribution is explicitly considered.

These findings imply that methodological sophistication should be proportional to information content. A large multivariate data set can support rich nonlinear dependence models, whereas a short annual sample may justify only parsimonious rank or parametric copula approaches. Model complexity should be justified by demonstrable gains in out-of-sample or tail calibration.

7. Discussion

The unified framework clarifies several common interpretive errors. Zero correlation is not evidence of independence; high mutual information is not evidence of direction; positive transfer entropy is not necessarily evidence of a direct causal channel; and a copula with good central fit is not necessarily well calibrated in the tail. Each conclusion requires a measure aligned with the relevant dependence concept.

The analysis also demonstrates that apparently competing literatures are often complementary. Copula methods and information-theoretic methods need not be framed as alternatives. Copulas provide a full joint-distribution representation and interpretable tail probabilities, while mutual information summarizes general dependence without requiring a particular copula family. Transfer entropy adds temporal directionality. Combining these tools can therefore produce a richer and more transparent empirical analysis.

8. Conclusion

Dependence measurement should begin with the scientific question rather than with a preferred statistic. Linear correlation, rank dependence, copula tail coefficients, mutual information and transfer entropy each identify different features of joint behavior. Their relative usefulness changes with nonlinearity, asymmetry, tail concentration, temporal direction and sample size.

The Monte Carlo evidence supports a multi-measure dependence dashboard for applications in finance, insurance and demography. Simple measures should be retained when they answer the relevant question efficiently, while flexible nonlinear or directional measures should be added when the data and decision problem justify them. The resulting approach is more informative than treating any single dependence statistic as a universal summary.

9. Interpretation Principles for Applied Research

A disciplined dependence analysis should begin by defining the null hypothesis implicit in each measure. Pearson correlation tests for linear co-movement around the mean, rank statistics summarize monotonic association, copula tails focus on conditional extreme probabilities, mutual information assesses general statistical dependence and transfer entropy addresses incremental predictive information through time. Confusion arises when results are interpreted outside those domains. For example, a low Pearson coefficient cannot support a claim of independence, and a high mutual-information estimate cannot establish directional transmission.

Transformation invariance is another useful organizing principle. Rank measures and copulas are invariant to strictly increasing marginal transformations, while Pearson correlation is not. Mutual information is theoretically invariant to smooth invertible transformations, although empirical estimators may not be perfectly invariant in finite samples. These properties matter when variables have strongly different scales or when marginal transformations are used to stabilize variance.

Tail measures require a separate interpretation because they intentionally discard much of the center of the distribution. A pair with modest overall association can have economically important lower-tail dependence, and a pair with strong central correlation can have zero asymptotic tail dependence. Rather than viewing these findings as contradictory, the dashboard treats them as answers to different questions. The disagreement among metrics is often the most informative feature of the analysis.

Direction adds yet another dimension. Transfer entropy and Granger causality concern temporal prediction, not contemporaneous dependence. A system can exhibit strong same-period copula dependence with little directional predictability if shocks are absorbed simultaneously. Conversely, a lagged information channel can be important even when contemporaneous association is modest. Applied papers should therefore align time indexing carefully with the economic transmission mechanism they claim to study.

10. Model Selection and Reporting Standards

Because flexible dependence measures introduce tuning choices, transparent reporting is essential. Mutual-information studies should document estimator type, bandwidth or neighbour parameter and any bias correction. Transfer-entropy studies should report lag orders, embedding choices, surrogate design and multiple-testing control. Copula studies should report marginal models, family selection, tail coefficients and goodness-of-fit diagnostics. Without these details, reported dependence values are difficult to reproduce or compare.

Model selection should also be task specific. Information criteria evaluate global density fit, whereas stress applications may require a tail-weighted criterion. Forecasting applications require out-of-sample scores, and network applications require stability of edge selection. The simulation framework demonstrates that the model with the best center-of-distribution fit is not always the model that best identifies the feature relevant to the decision problem.

Researchers should avoid ranking dependence measures using a single universal performance criterion. A measure cannot be penalized for failing to detect a feature it was not designed to detect. Instead, performance should be assessed against design-specific targets: correlation recovery in linear systems, tail-probability calibration in extreme-value designs, dependence detection in nonlinear designs and directional-edge recovery in dynamic systems.

A minimum reporting set for general applied work could include Pearson correlation, Kendall tau, one nonlinear dependence measure and, where economically relevant, a tail or directional measure. This combination is not mandatory, but it provides an accessible diagnostic ladder from simple to more specialized dependence concepts.

11. Limitations and Future Research

The Monte Carlo designs necessarily simplify real data. Financial and demographic systems can contain nonstationarity, structural breaks, mixed discrete-continuous variables and measurement error. These features can interact with dependence estimators in ways not captured by the present experiments. Future simulation studies should incorporate such complications and evaluate adaptive or locally stationary methods.

High-dimensional dependence remains a major open problem. Pairwise dashboards can become unwieldy when hundreds of variables are analyzed. Graphical models, factor copulas, sparse mutual-information networks and multivariate information decompositions provide possible solutions, but they raise new questions about regularization and interpretability. A useful direction is to combine low-dimensional interpretable summaries with high-dimensional screening procedures.

Another promising area is uncertainty quantification. Dependence estimates are often reported as point values, yet nonlinear and tail measures can have substantial sampling uncertainty. Bootstrap intervals, Bayesian posterior distributions and stability-selection frequencies would make applied conclusions more credible. The ultimate objective is not to maximize the number of sophisticated measures but to know which dependence statements are supported robustly by the available information.

References

- Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.
- Aas, K., Czado, C., Frigessi, A., & Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2), 182–198.
- Bedford, T., & Cooke, R. M. (2002). Vines—A new graphical model for dependent random variables. *Annals of Statistics*, 30(4), 1031–1068.
- Czado, C. (2019). *Analyzing Dependent Data with Vine Copulas*. Springer.
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions. *Journal of Econometrics*, 182(1), 119–134.
- Embrechts, P., McNeil, A., & Straumann, D. (2002). Correlation and dependence in risk management. In M. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*. Cambridge University Press.
- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339–350.
- Joe, H. (2014). *Dependence Modeling with Copulas*. CRC Press.
- McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools* (2nd ed.). Princeton University Press.
- Nelsen, R. B. (2006). *An Introduction to Copulas* (2nd ed.). Springer.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2), 527–556.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423, 623–656.
- Rényi, A. (1961). On measures of entropy and information. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, 1, 547–561.
- Tsallis, C. (1988). Possible generalization of Boltzmann–Gibbs statistics. *Journal of Statistical Physics*, 52, 479–487.
- Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85(2), 461–464.

- Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. *Physical Review E*, 69(6), 066138.
- Maasoumi, E., & Racine, J. (2002). Entropy and predictability of stock market returns. *Journal of Econometrics*, 107(1–2), 291–312.
- Barnett, L., Barrett, A. B., & Seth, A. K. (2009). Granger causality and transfer entropy are equivalent for Gaussian variables. *Physical Review Letters*, 103(23), 238701.

Declarations

Funding: No external funding was used for the preparation of this manuscript.

Conflict of interest: The authors declare no competing interests.

Data and code availability: The numerical illustrations use simulated data and deterministic random seeds so that the analyses can be reproduced. Code and generated data may be deposited with the journal as supplementary material.