



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.