

Marginal Specification and Tail Dependence in Equity Markets

Haozhou He*

Business school, Monash University, Caulfield East, Australia

**Corresponding author: Haozhou He.*

Abstract

Copula conclusions for equity markets depend on both marginal filtering and the temporal pairing of daily returns. We analyse S&P 500, FTSE 100, and Nikkei 225 returns from 2010 to 2025 in three pairwise panels. A rolling 756-day training and 63-day evaluation design compares empirical-CDF, Gaussian-GARCH, and Student-t-GARCH margins. The GARCH filters use a standard one-step-ahead recursion, and every candidate copula is estimated by pseudo maximum likelihood. The Student-t copula has the highest window-weighted density score under all three margins, with scores of 0.0924, 0.0703, and 0.0685. It is the numerical winner in eight of nine panel–margin cells, but paired moving-block intervals support a positive advantage in only three cells. A lead–lag check materially changes the S&P 500–Nikkei 225 association, while a joint-downside probability score does not select the same family as the density score. The evidence therefore supports specification-aware conditional-dependence analysis, rather than an invariant model ranking or causal contagion claim.

Keywords

copula models, cross-market dependence, tail dependence, GARCH filtering, temporal alignment

1. Introduction

Equity indices can co-move strongly in ordinary periods and in joint downturns. Linear correlation alone does not distinguish these forms of association. Copulas separate a dependence distribution from fitted univariate margins and can represent symmetric or asymmetric tail behaviour [1]. Financial applications consequently combine copulas with volatility filters, dynamic-correlation models, and rolling estimation [2–5]. Yet a dependence result can inherit properties of the marginal transformation used to create the pseudo-observations.

This issue is important for two reasons. First, volatility clustering and heavy tails change the probability-integral transform entering a copula fit [6]. A conclusion obtained after one filter may therefore not survive a different, plausible filter. Second, matching two closing returns by calendar date does not guarantee that the two returns reflect the same information set. The Japanese close precedes the US close, and the European close also precedes it. Calendar matching can therefore conceal lead–lag dependence, especially for US–Asia pairs.

Four choices are consequently entangled in a typical cross-market exercise: the return-pairing rule, the marginal filter, the copula family, and the loss function used to compare fitted models. Each choice is defensible in isolation. Their interaction is harder to assess because a change in pseudo-observations can alter the fitted parameter, the implied tail probability, and the held-out score at the same time. A density-score winner should therefore be interpreted conditional on the pairing and marginal assumptions that produced the score. The same caution applies to an apparent difference in tail dependence across regions.

This paper asks how marginal specification and temporal pairing affect rolling copula evidence for the S&P 500, FTSE 100, and Nikkei 225. It holds the candidate family set and out-of-sample protocol fixed, while comparing empirical-CDF, Gaussian-GARCH, and Student-*t*-GARCH margins. The primary contribution is not a new copula family. It is a reproducible robustness analysis with standard one-step-ahead GARCH filtering, common pseudo-maximum-likelihood estimation, uncertainty intervals for score differences, and a joint-downside prediction metric.

The analysis is designed around a simple counterfactual. If a family is genuinely robust to reasonable marginal choices, it should retain a comparable score advantage when the pseudo-observations are regenerated by another plausible pipeline. If the advantage disappears, changes sign, or is statistically unresolved, the appropriate scientific result is sensitivity rather than a forced winner. Similarly, if a lead-lag pairing changes the association, the result should be reported as temporal-alignment sensitivity. It should not be transformed into an informal causal story about market transmission.

The results are deliberately qualified. The Student-*t* copula is the numerical density-score leader after all three margins, but its lead over the closest competitor is not statistically positive in most panel-margin cells. A right-market lead-lag shift substantially changes the S&P 500–Nikkei 225 association. The study therefore concerns conditional association under stated specifications. It does not identify contagion, directed spillovers, portfolio profitability, or systemic risk. Those questions require additional identification or decision designs [7–9].

2. Related Work

Copula studies document non-linear and tail association in stock, commodity, foreign-exchange, sectoral, and technology-equity markets [10–14]. Copula–GARCH formulations model conditional marginal volatility before estimating a separate dependence layer [4]. Conditional copula work also formalises time-varying and asymmetric dependence [15]. Other work introduces time variation through GAS or DCC-type dynamics [5, 16], and flexible vine or regime-switching structures extend the family set [17]. These studies show why tail dependence is useful, but they commonly select one marginal specification before comparing dependence families.

The resulting methodological gap is narrower. Marginal transformation affects the pseudo-observations, whereas the copula comparison is usually interpreted as if only the dependence family had changed. Heavy-tailed marginal transformations can alter tail-risk conclusions even with a fixed Gaussian copula [6]. Our design therefore fixes the copula set, rolling schedule, and scoring rule, and changes only the marginal filter or temporal pairing. This isolates specification sensitivity from a claim of universal model superiority.

The comparison also requires a common estimation target. A correlation inversion for an elliptical copula and a Kendall's- τ inversion for an Archimedean copula can be useful descriptive estimators, but they do not make a family contest identical when evaluation is based on predictive density. We instead fit every family by pseudo maximum likelihood and then evaluate the same held-out density. This does not remove model misspecification, but it ensures that the family label is not confounded with a different parameter-training rule. Estimating the Student-*t* copula degrees of freedom further prevents the primary comparison from assuming the same tail thickness in every rolling block.

Forecast comparison presents a second distinction. A full copula density rewards accurate probability allocation across the entire unit square. A risk manager may instead care about a small lower-left region, such as the probability that both standardized returns fall below their fifth percentiles. These are related but non-identical tasks. Reporting both metrics makes it possible to identify whether an aggregate density result carries over to a concrete joint-downside scenario.

This distinction also separates conditional dependence from contagion. Event studies and causal-network approaches use structural breaks, shocks, or directed assumptions to study transmission [7, 8]. A rolling observational copula fit does not contain those elements. We use joint-downside probabilities as a forecast-evaluation scenario, not as evidence of a causal risk channel.

3. Method

3.1 Data and Rolling Design

We use archived public adjusted closes for the S&P 500, FTSE 100, and Nikkei 225 from 1 January 2010 to 31 December 2025. Daily returns are $r_t = \log(P_t/P_{t-1})$. Common-date pairing yields 3,946 S&P 500–FTSE 100 returns, 3,778 S&P 500–Nikkei 225 returns, and 3,803 FTSE 100–Nikkei 225 returns. The primary design fits each model on 756 returns and evaluates the next 63. Origins advance by 63 returns, producing 50, 47, and 48 blocks, respectively. Raw responses and SHA-256 hashes are retained with the experiment.

Calendar alignment is not an information-time assumption. For each pair, a temporal robustness check shifts the right-hand market by -1 , 0 , and $+1$ common trading days before calculating association and the Student- t -copula score. Thus, a positive shift pairs the left-market return at t with the right-market return at the next common trading day.

The lag convention is a robustness device, not a structural model of price discovery. It retains the same return construction and evaluates a limited set of adjacent common trading days. Holidays, exchange-specific closing conventions, and intraday news timing are not modelled. A large change under a shift nevertheless signals that the zero-lag calendar pair should not be treated as an innocuous benchmark.

3.2 Marginal Models

Three pipelines generate pseudo-observations in $(0,1)$. The empirical-CDF pipeline ranks test returns against the training distribution, using a half-rank correction and denominator $n + 1$, where n is the training-sample size. The parametric pipelines use the stationary GARCH(1,1) formulation [2],

$$r_t = \epsilon_t, \quad \epsilon_t = \sqrt{h_t} z_t, \quad h_t = \omega + \alpha \epsilon_{t-1}^2 + \beta h_{t-1}, \quad (1)$$

where h_t is known from information available at $t - 1$. At the first test date, the filter begins from $h_{T+1|T} = \omega + \alpha \epsilon_T^2 + \beta h_T$, then updates sequentially using realised test returns. No test recursion is reset to the training variance. Parameters satisfy $\omega > 0$, $\alpha, \beta \geq 0$, and $\alpha + \beta < 0.995$. The bound $\alpha + \beta < 0.995$ is imposed as a numerical stationarity buffer that keeps optimization away from the IGARCH boundary.

Gaussian GARCH uses a Gaussian likelihood and normal probability transform. Student- t GARCH uses a unit-variance Student likelihood and transform. The primary Student- t margin uses six degrees of freedom, while $\nu \in \{4,6,8,10\}$ is examined as a sensitivity check. Parameters are re-estimated in every training block.

3.3 Copula Estimation and Evaluation

Gaussian, Student- t , Clayton, and Gumbel copulas are all estimated by the same pseudo-maximum-likelihood criterion [15],

$$\hat{\theta} = \operatorname{argmax}_{\theta} \sum_{i=1}^T \log c(u_i, v_i; \theta). \quad (2)$$

Here, c is the copula density and (u_i, v_i) are training pseudo-observations. For the Student- t copula, both the correlation and degrees of freedom are estimated, with $2.05 \leq \nu \leq 40$. Gaussian, Clayton, and Gumbel parameters are optimized under their valid domains. This removes heterogeneous parameter-training rules from the family comparison.

The fitted families represent distinct tail restrictions. The Gaussian copula is asymptotically tail independent. The Student- t copula has symmetric upper and lower tail dependence, with its strength determined jointly by the fitted correlation and degrees of freedom. Clayton represents lower-tail dependence, whereas Gumbel represents upper-tail dependence. For every fitted family, the reported total tail dependence is $\lambda_{\text{total}} = \lambda_L + \lambda_U$, where λ_L and λ_U denote the lower- and upper-tail coefficients, respectively. The candidate set is

intentionally compact: it exposes which conclusions are robust across common symmetric and one-sided alternatives, but it does not claim to span all possible dependence structures.

The primary metric is the mean held-out density score,

$$\bar{s} = m^{-1} \sum_{j=1}^m \log c(u_j, v_j; \hat{\theta}), \quad (3)$$

where $m = 63$. We also evaluate a financial-tail scenario. For a fixed $q = 0.05$, each fitted copula distribution C forecasts $p = C(q, q)$ for the event $I = \mathbb{1}\{u < q, v < q\}$, and receives the binary log score $I \log p + (1 - I) \log(1 - p)$. This metric directly distinguishes copula-implied joint downside probabilities.

The two scores answer different questions. Density scoring evaluates the fitted copula at every held-out pseudo-observation. The event score discards the magnitude of non-event observations and focuses on calibration of a specified joint downside event. Neither score is a portfolio return, value-at-risk, or expected-shortfall result. In particular, the analysis does not combine indices in different currencies into an investment portfolio, so it makes no international allocation claim.

For each panel–margin cell, the highest mean density-score model is compared with its runner-up using paired moving-block resampling of rolling score differences. We report 2,000 replicates, block length four, and a 95% percentile interval. Tail-dependence paths are descriptive robustness summaries. Lower dispersion is not treated as an intrinsic model improvement because genuine dependence can change over time.

4. Results

4.1 Density-score Ranking and Its Uncertainty

At the 756-day window, the Student- t copula has the highest window-weighted density score under every marginal pipeline (Table 1). The ranking is therefore numerically robust at the aggregate level. However, the size and inferential status of the advantage depend on the panel. Student- t is the numerical winner in eight of nine panel–margin cells, whereas its paired moving-block interval is strictly positive in only three. The remaining six intervals include zero. Aggregate ranking should thus be read as a useful summary, not as proof of uniformly superior density forecasts.

The positive intervals arise in specific conditions rather than throughout the design. Under empirical-CDF margins, the Student- t lead is 0.0289 over Gumbel for S&P 500–FTSE 100, with a 95% interval of [0.0106, 0.0518], and 0.0186 over Clayton for S&P 500–Nikkei 225, with [0.0021, 0.0415]. Under Student- t GARCH, the S&P 500–FTSE 100 Student- t –Gaussian difference is 0.0063, with [0.0013, 0.0116]. These cases support a local predictive advantage. They do not justify carrying the same claim to the remaining conditions, where the closest-family interval crosses zero.

Table 1. Window-weighted mean held-out density scores over the three index pairs at the 756-day window. Higher is better. Bold marks the largest score within a marginal pipeline


Copula	Empirical CDF	Gaussian GARCH	Student- t GARCH
Gaussian	0.0651	0.0661	0.0636
Student- t	0.0924	0.0703	0.0685
Clayton	0.0652	0.0520	0.0568
Gumbel	0.0681	0.0528	0.0570

Figure 1 displays the same comparison separately for each marginal pipeline. This panel design avoids mixing scores obtained from different pseudo-observations in one bar chart. The empirical-CDF condition produces the widest separation, while the two GARCH conditions place the Gaussian and Student- t copulas much closer together.

The downside-event metric gives a complementary result. At the aggregate level, Clayton has the best joint 5% downside-event score under empirical CDF (−0.0811), Gaussian GARCH (−0.0855), and Student- t GARCH (−0.0898). A Student- t copula can therefore lead on full-density fit without being the preferred model for the particular lower-tail event. This difference illustrates why a financial scenario should be evaluated with a copula-dependent metric rather than with a portfolio quantile that is identical across copulas.

The distinction is economically intuitive but should not be over-read. The full-density score uses both central and tail observations, so it can prefer a flexible symmetric family even when a lower-tail-only event is more accurately calibrated by Clayton. Conversely, the event score is based on a rare binary outcome and is not a general test of distributional fit. The two metrics together establish a bounded result: model choice depends on the predictive target, not only on the marginal filter.


Figure 1. Mean held-out copula density score by marginal pipeline. Each panel contains an identical four-family comparison



4.2 Tail Dependence and Temporal Alignment

Figure 2 separates rolling Student-*t*-copula tail-dependence estimates by index pair. The S&P 500–FTSE 100 distribution lies above the S&P 500–Nikkei 225 distribution, while the FTSE 100–Nikkei 225 panel occupies an intermediate range. Within each pair, marginal filtering shifts the estimated distribution. The figure is descriptive, so it does not interpret a lower standard deviation as a better model.

Figure 2: Rolling total Student-*t*-copula tail dependence by index pair and marginal pipeline at the 756-day window




Temporal alignment is consequential. For S&P 500–Nikkei 225, same-date Pearson correlation is 0.1723, but pairing the S&P 500 return with the next Nikkei return raises it to 0.4469. The corresponding Student-*t*-copula mean density score rises from 0.0206 to 0.1355. The direction of the shift is consistent with the ordering of market closes, but it is not a causal estimate. It demonstrates that same-date matching is an empirical convention rather than a neutral information-time choice. Other pairs also change under lead–lag shifts, although less sharply.

The contrast is not limited to a simple correlation statistic. For the S&P 500–FTSE 100 pair, the same-date Student-*t* score is 0.1341, compared with 0.0011 under a negative shift and 0.0196 under a positive shift. For FTSE 100–Nikkei 225, the corresponding values are 0.0471, 0.0010, and 0.0428. These patterns show that the inferred dependence is sensitive to the observation schedule itself. They also motivate reporting the pair definition alongside the copula specification in cross-market applications.

The Student- t marginal degrees-of-freedom check gives modest, non-monotone score changes across $\nu \in \{4,6,8,10\}$. For example, the S&P 500–FTSE 100 score ranges from 0.1312 to 0.1344, and the S&P 500–Nikkei 225 score ranges from 0.0179 to 0.0206. Thus, fixing six degrees of freedom is not decisive for the observed numerical ranking, but it remains a modelling choice rather than a universal calibration.

Figure 3 plots tail-dependence paths against the forecast-origin date, not the training-window start. Each point therefore locates the held-out block in calendar time. The paths show persistent cross-pair differences and temporal variation, but they do not identify a crisis effect or contagion episode.

Figure 3: Rolling total tail dependence under the Student- t -GARCH and Student- t -copula configuration. Each value is associated with the origin of its subsequent 63-day evaluation block



5. Discussion

The revised analysis changes the interpretation of the earlier aggregate finding. Student- t remains the numerical density-score leader after three marginal pipelines, which is consistent with heavy-tailed joint-return behaviour. Yet numerical leadership is not a uniform statistical advantage. Six of nine winner–runner-up intervals include zero, and Clayton is preferred for the specific joint-downside score. The appropriate conclusion is specification-aware model comparison, not a claim that one copula is best for every financial purpose.

The lead–lag evidence is equally important. A same-date US–Japan return pair combines markets that close at different points in the global trading day. The strong change under a positive Nikkei shift shows that calendar alignment alone does not solve this problem. Researchers should report the temporal pairing rule and test plausible alternatives before attributing a low same-date dependence estimate to a structural difference between markets.

The evidence also reframes the role of tail-dependence dispersion. A smoother rolling tail-dependence path is not automatically preferable, because genuine joint risk can vary during stress or policy changes. The relevant robustness question is whether marginal choices materially alter the estimated path and the associated predictive ranking. In this study they do alter both, although the aggregate Student- t density ranking remains intact. This interpretation preserves the negative or inconclusive comparisons as substantive information rather than treating them as failed support for a preferred filter.

For practical use, the joint-downside exercise provides a more direct bridge from a fitted copula to a financial scenario. Its result does not select the same model as the full-density criterion. A practitioner must therefore state the decision target before translating a statistical dependence fit into a risk forecast. The present data do not include common-currency portfolio construction, transaction costs, or a trading rule. Consequently, the scenario result should be understood as probabilistic calibration of co-downside events, not an investment recommendation.

Several boundaries remain. The analysis uses daily closes and three bivariate pairs, so it does not describe an entire market network. The candidate set omits asymmetric two-tail families, regime switching, and dynamic copulas [17, 18]. The Student- t marginal sensitivity grid does not replace full marginal goodness-of-fit testing, and the bootstrap assesses pairwise ranking differences rather than a complete model-confidence

set [19]. Finally, joint downside probabilities are predictive quantities under an observational design. They do not identify spillovers, causal contagion, systemic risk, or investment gains.

6. Conclusion

This study examined how marginal filtering and temporal alignment shape copula-based cross-market dependence evidence. With standard one-step-ahead GARCH filtering and common pseudo-maximum-likelihood estimation, the Student- t copula had the largest aggregate density score under empirical-CDF, Gaussian-GARCH, and Student- t -GARCH margins. Its advantage was nevertheless statistically unresolved in most panel-margin comparisons, and the joint-downside scenario favoured Clayton. The lead-lag analysis further showed that S&P 500–Nikkei 225 dependence changes materially when the pairing respects market timing. The evidence supports specification-aware conditional dependence, not a universal ranking or causal co-movement claim.

References

- [1] Demarta, S., and McNeil, A. J., “The t copula and related copulas,” *International Statistical Review*, 73, (1), 111–129, 2005, <https://doi.org/10.1111/j.1751-5823.2005.tb00254.x>.
- [2] Bollerslev, T., “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, 31, (3), 307–327, 1986, [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1).
- [3] Münnix, M. C., and Schäfer, R., *A copula approach on the dynamics of statistical dependencies in the US stock market*, 2011, <https://arxiv.org/abs/1102.1099>.
- [4] Chen, Q., Wang, D., and Pan, M.-S., *Multivariate time-varying copula GARCH model and its application in the financial market risk measurement*, 2015, <https://doi.org/10.1155/2015/286014>.
- [5] Kim, Y., Hwang, S., Kim, J.-M., and Kim, S., *Linear time-varying regression with copula–DCC–asymmetric–GARCH models for volatility: The co-movement between industrial electricity demand and financial factors*, 2022, <https://doi.org/10.1080/00036846.2022.2086684>.
- [6] Das, B., and Fasen-Hartmann, V., *On heavy-tailed risks under gaussian copula: The effects of marginal transformation*, 2023, <https://arxiv.org/abs/2304.05004>.
- [7] Oh, M., and Kim, D., *Effect of the u.s.–china trade war on stock markets: A financial contagion perspective*, 2021, <https://arxiv.org/abs/2111.09655>.
- [8] Rigana, K., Wit, E. C., and Cook, S., *Navigating market turbulence: Insights from causal network contagion value at risk*, 2024, <https://arxiv.org/abs/2402.06032>.
- [9] Vasiljeva, A., Matvejevs, A., and Fjodorovs, J., *Dependence modeling and portfolio optimization with copula-GARCH: A european investment perspective*, 2025, <https://doi.org/10.3389/fams.2025.1675120>.
- [10] Cambriles, A. P., and Muela, S. B., *Assessing the structure dependence between the spanish stock market and some international financial markets. A time-varying copula analysis.*, 2023, <https://doi.org/10.32826/reyf.v1i1.342>.
- [11] Tian, M., Khoury, R. E., and Alshater, M. M., *The nonlinear and negative tail dependence and risk spillovers between foreign exchange and stock markets in emerging economies*, 2022, <https://doi.org/10.1016/j.intfin.2022.101712>.
- [12] Mensi, W., Hanif, W., Bouri, E., and Vo, X., *Spillovers and tail dependence between oil and US sectoral stock markets before and during COVID-19 pandemic*, 2023, <https://doi.org/10.1108/ijoem-12-2021-1799>.
- [13] Cao, S., and Sun, Z., *Extreme risk spillovers between chinese crude oil and stock markets: A CoVaR-copula approach*, 2025, <https://doi.org/10.23977/ferm.2025.080107>.
- [14] Zhou, X., and Liu, H., *Cross-market risk spillovers and tail dependence between u.s. And chinese technology-related equity markets*, 2025, <https://doi.org/10.3390/ijfs13040242>.

- [15] Patton, A. J., “Modelling asymmetric exchange rate dependence,” *International Economic Review*, 47, (2), 527–556, 2006, <https://doi.org/10.1111/j.1468-2354.2006.00387.x>.
- [16] Gong, X.-L., Liu, X.-H., and Xiong, X., *Measuring tail risk with GAS time varying copula, fat tailed GARCH model and hedging for crude oil futures*, 2019, <https://doi.org/10.1016/J.PACFIN.2019.03.010>.
- [17] Tsai, J.-J., Wang, Y.-C., and Zhang, J., *Measuring financial contagion between major european stock markets via the regime-switching vine copula model*, 2018, <https://doi.org/10.3934/dcdss.2020267>.
- [18] Chen, W., Wei, Y., Lang, Q., Lin, Y., and Liu, M., *Financial market volatility and contagion effect: A copula–multifractal volatility approach*, 2014, <https://doi.org/10.1016/J.PHYSA.2013.12.016>.
- [19] Hansen, P. R., Lunde, A., and Nason, J. M., “The model confidence set,” *Econometrica*, 79, (2), 453–497, 2011, <https://doi.org/10.3982/ECTA5771>.

Funding

This research received no external funding.

Conflicts of Interest

The authors declare no conflict of interest.

Acknowledgment

This paper is an output of the science project.

Open Access

This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

