FUDMA Journal of Sciences (FJS)
ISSN online: 2616-1370
ISSN print: 2645 - 2944
Vol. 10 No. 8, April, 2026, pp 326 – 335
DOI: https://doi.org/10.33003/fjs-2026-1008-5333

Sensitivity Analysis of Value-at-Risk and Expected Shortfall to Copula Misspecification in CGMY Jump-Diffusion Models

***[1]Abdulmudallib Ibrahim, [2]Adesupo A. Akinrefon, [2]Okolo Abraham and [2]Emmanuel Torsen**

[1]Department of Statistics, Federal University Wukari, Taraba State, Nigeria.
[2]Department of Statistics, Modibbo Adama University, Yola, Adamawa, Borno State, Nigeria.

*Corresponding authors' email: abdulmudallib@fuwukari.edu.ng

ABSTRACT

This study provides a comprehensive sensitivity analysis of risk measures to copula misspecification in CGMY jump-diffusion models, revealing substantial vulnerabilities in Value-at-Risk (VaR) and Expected Shortfall (ES) calculations. Using Monte Carlo simulation with 10,000 paths per scenario across multiple volatility regimes, jump intensities, and confidence levels, the study demonstrates that mean relative errors exceed 200% for VaR and 146% for ES when averaging across all six copula families including the severely misspecified Gumbel copula. Crucially, well-specified elliptical copulas exhibit bootstrap mean errors of only 1.5–2.3% for VaR and 2.0–2.3% for ES, while the Gumbel copula alone drives errors to 280.6% for VaR and 238.4% for ES in bootstrap-estimated means (corresponding to 1,190% and 839% in the worst individual scenarios), indicating that the aggregate averages are dominated by a single family's structural failure rather than a pervasive property of the modelling framework. Two-way ANOVA confirms that both copula family ($F = 3{,}614$ for VaR; $F = 5{,}192$ for ES) and confidence level ($F = 66$ for VaR; $F = 44$ for ES) exert highly significant effects (all $p < 2\times10^{-10}$), while Tukey HSD post-hoc tests demonstrate that the Gumbel copula is the sole source of statistically distinguishable pairwise differences all five non-Gumbel copulas are statistically indistinguishable from one another at the 95% family-wise level. Contrary to expectation, low-volatility regimes exhibit the highest sensitivity to copula misspecification, while severe-jump regimes exhibit the lowest. Taken together, these findings carry critical implications for regulatory capital calculations under Basel III, which has shifted emphasis from VaR to ES, and for the selection and governance of dependence models in financial risk management.

Keywords: Value-at-Risk; Expected Shortfall; Copula Misspecification; CGMY Jump-Diffusion; Model Risk; Basel III; Risk Management

INTRODUCTION

The accurate measurement of financial risk has become increasingly critical in the wake of successive financial crises, regulatory reforms, and the growing complexity of financial instruments. Value-at-Risk (VaR) and Expected Shortfall (ES) have emerged as fundamental risk metrics for both internal risk management and regulatory capital adequacy. ES has gained particular prominence under Basel III frameworks due to its superior theoretical properties as a coherent risk measure, satisfying translation invariance, positive homogeneity, monotonicity, and subadditivity guaranteeing that diversification is always recognised as risk-reducing, a property VaR lacks (McNeil et al., 2005).

The reliability of VaR and ES depends critically on the accuracy of underlying distributional assumptions and dependence structures. In multivariate portfolios, copulas provide the mathematical framework for modelling dependence between risk factors while preserving flexibility in marginal distribution specification (Embrechts et al., 2003). Copula misspecification - selecting a copula family that does not reflect the true dependence structure - introduces systematic model risk that can propagate through risk calculations with severe consequences for capital adequacy and regulatory compliance. As Boucher et al. (2014) demonstrate, model uncertainty in dependence specifications can contribute substantially to systemic risk, as evidenced during the 2008 financial crisis.

CGMY jump-diffusion models (Carr et al., 2002) have gained prominence for capturing empirical return characteristics including heavy tails, asymmetry, and jump clustering. The CGMY model's flexible parameterisation, which accommodates both finite and infinite jump activity, presents additional complexity when combined with copula modelling, as Lévy processes are popular for stock price behaviour since they account for jump risk and reproduce implied volatility smiles (Ornthanalai, 2014). The interaction between sophisticated jump processes and copula-based dependence modelling creates a particularly challenging model risk environment where errors can cascade throughout the risk model (Cerrato et al., 2017).

Despite extensive research on copula applications and jump-diffusion processes, a significant gap remains in understanding how copula misspecification affects VaR and ES calculations in CGMY jump-diffusion frameworks. This study addresses this gap with four specific contributions. It shows that aggregate mean errors exceeding 200% are driven almost entirely by the Gumbel copula's structural failure, while well-specified elliptical copulas produce errors below 10%. Building on that, it shows the 'low-volatility paradox': copula misspecification is most severe during calm markets. Consequent to these findings, it provides analysis of non-monotonic confidence level sensitivity between 95% and 99% thresholds, and furnishes guidance for risk model governance consistent with Basel III's emphasis on ES.

Literature Review

Value-at-Risk, Expected Shortfall, and Coherence Properties

Value-at-Risk at confidence level α is defined as the minimum loss L exceeded with probability no greater than 1−α over a specified horizon:

$$VaR_{\alpha}(L) = \inf \{1 \in R : P(L > 1) \le 1 - \alpha\} \qquad (1)$$

Expected Shortfall (also known as Conditional Value-at-Risk) is the expected loss conditional on exceeding the VaR threshold:

$$ES_{\alpha}(L) = E\left[L \middle| L \geq VaR_{\alpha}(L)\right] \quad (2)$$

ES satisfies all four coherence axioms including subadditivity, making it theoretically superior to VaR for portfolio optimisation and capital allocation (McNeil et al., 2005). The regulatory shift from VaR to ES under Basel III places heightened importance on accurate tail dependence modelling, since ES is a conditional expectation in the deepest part of the loss distribution where copula specification has the greatest impact.

Copula Theory and Model Risk

Copula theory, introduced by Sklar (1959), provides the foundation for separating marginal distribution modelling from dependence structure specification. Sklar's theorem states that for any bivariate joint distribution F with continuous marginals F_1 and F_2, there exists a unique copula C such that;

$$F(x, y) = C(F_1(x), F_2(y))$$

This separation is particularly valuable in financial modelling because it allows flexible specification of heavy-tailed marginals alongside a range of dependence patterns (Embrechts et al., 2003).

Boucher et al. (2014) formalise the concept of 'risk models-at-risk', showing how model uncertainty in dependence specifications can immensely contribute to systemic risk. The results demonstrate that model risk arising from specification and estimation uncertainties often makes standard risk measures inadequate to quantify extreme downward risk, reinforcing the need for robust copula validation in regulatory frameworks.

Zhang et al. (2013) propose a goodness-of-fit test for semiparametric copula models comparing in-sample and out-of-sample pseudo-likelihoods. The simulation experiments demonstrate that standard tests have limited power against tail misspecification, motivating the simulation-based evaluation framework adopted here. The dangers of copula misspecification were starkly illustrated by the Gaussian copula's role in the 2008 crisis, where its failure to capture tail dependence in mortgage-backed securities contributed directly to underestimation of systemic risk (MacKenzie & Spears, 2014).

Patton (2006) demonstrates that asymmetric exchange rate dependence cannot be captured by symmetric Gaussian copulas, providing evidence that copula family selection materially affects the accuracy of risk-based inference. Brechmann and Czado (2013) examine risk management using high-dimensional vine copulas, demonstrating that misspecification of bivariate building blocks propagates into large estimation errors for portfolio VaR and ES. Derman (2011) argues broadly that practitioners who mistake model outputs for reality expose themselves to catastrophic losses a lesson directly applicable to copula selection in jump-diffusion environments. Dewick and Liu (2022) provide an extended review of copula modelling in financial contexts, emphasising careful selection to mitigate model risk.

Jump-Diffusion Models and Risk Measurement

Nthiwa *et al.,* (2023) demonstrate that jump-diffusion models incorporating stochastic volatility substantially improve the fit to S&P 500 return data relative to the Black–Scholes–Merton model. The results, evaluated across bullish, bearish, and neutral market conditions, show that models with skewed Student-t error distributions achieve the lowest MSE, consistent with the CGMY framework's asymmetric decay parameters G and M.

Ornthanalai (2014) provides evidence that Lévy jump risk carries a priced risk premium distinct from diffusive volatility risk in the CGMY framework, and that models neglecting this premium systematically misprice out-of-the-money options. The inverse relationship between jump severity and copula misspecification sensitivity documented in the present study is consistent with Ornthanalai's finding that extreme events are dominated by marginal distribution effects rather than dependence structure.

Research Gap

Despite extensive literature on copula misspecification and CGMY jump modelling, no prior study has systematically quantified VaR and ES sensitivity to copula family choice within the CGMY jump-diffusion framework across multiple volatility regimes, jump intensities, and confidence levels. Studies such as Patton (2006) address copula misspecification in exchange rate settings and Brechmann and Czado (2013) address high-dimensional equity portfolios, but the specific interaction between CGMY infinite-activity jump processes and copula tail dependence architectures in generating VaR and ES errors has not been examined. The present study fills this gap.

MATERIALS AND METHODS

CGMY Jump-Diffusion Framework

The asset price process under the physical measure P is specified as:

$$dS_t = {}_{\mu}S_t dt + {}_{\sigma}S_t dW_t + dJ_t \quad (3)$$

Where μ is the drift, σ is the diffusion volatility, W (t) is a standard Brownian motion, and J (t) is the CGMY pure-jump process. For VaR and ES computation, a bivariate equally-weighted portfolio is constructed with two assets whose dependence structure is modelled through the copula C applied to the marginal distributions.

CGMY Lévy Measure

The CGMY process is characterised by its Lévy measure:

$$v(dx) = C|x|^{-1-Y}\exp(-G|x|)dx \quad \text{(Negative jumps) (4)}$$

$$v(dx) = C|x|^{-1-Y}\exp(-M|x|)dx \quad \text{(Positive jumps) (5)}$$

C > 0 controls overall jump intensity; G, M > 0 are exponential decay rates for negative and positive jumps; Y ∈ (−∞, 2) governs fine structure of small jumps. When Y > 0, the process has infinite activity (Carr et al., 2002). Base parameters: C = 2.0, G = 5.0, M = 7.0, Y = 0.8.

Simulation Design and Benchmark

VaR and ES are computed using Monte Carlo simulation with N = 10,000 paths per scenario, daily time steps (252 per year). The 'true' benchmark VaR and ES are computed under the specified Student's t-copula. Misspecification errors are relative deviations of each alternative copula's estimate from this benchmark. All relative error statistics are expressed as percentages of the benchmark value (e.g., 1,190.14% means the model-implied VaR is 11.9 times the benchmark). Aggregate mean errors of 201.41% reflect the arithmetic mean across all six copulas and are dominated by the Gumbel copula's severe failure; see Table 4 for copula-specific disaggregation.

Risk Measure Computation

VaR Absolute Error = $|Explained\ VaR - Actual\ VaR|$ (6)

Relative Error in VaR (%) = $\left|\frac{VaR_{true} - VaR_{model}}{VaR_{true}}\right| \times 100$ (7)

ES Absolute Error = $\left|ES_{true} - E[ES]\right|$ (8)

Relative Error in ES (%) = $\left|\frac{ES_{true} - ES_{model}}{ES_{true}}\right| \times 100$ (9)

Copula Specifications

Six specific copula functions were considered for the purpose, including the Student's t-copula which was examined under three different degrees of freedom specifications (ν = 3, 5, and 8) to assess sensitivity to tail heaviness and extreme dependence characteristics.

i. *Discrete Gaussian copula:* This type of copula can capture linear dependencies between the two processes under assumptions of normality (Cherubini, Luciano, & Vecchiato, 2004). The formula is given by

$$C(u_1, u_2; \theta) = \Phi_G\left(\Phi^{-1}(u_1), \Phi^{-1}(u_2); \theta\right) \quad (10)$$

Where;

u_1, u_2 = Marginal cumulative distribution functions (CDFs) of the random variable

Φ^{-1} = The inverse of the standard normal CDF

$\Phi^{-1}{}_G$ = The joint CDF of a bivariate normal distribution with a correlation parameter θ

ii. *Student's t-Copula:* This copula has proved particularly useful for modelling such dependencies with apparently heavier tails than the Gaussian Copula and hence helps capture extreme co-movements (Demarta & McNeil, 2005). The Student's t copula is given by

$$C(u_1, u_2; v, \theta) = t_v\left(t^{-1}(u_1), t^{-1}(u_2); \theta\right) \quad (11)$$

Where;

t_v = The Joint CDF of a bivariate Student's t-distribution with v degrees of freedom and parameter θ

t^{-1} = The inverse CDF of the Student's t-distribution.

iii. *Clayton Copula:* Already found to be very efficient in modelling lower tail dependencies, (Embrechts, Lindskog, & McNeil, 2003) hence, it will be ideal for situations when the possibility of extreme negative return occurrences together is high. The Clayton Copula can be expressed as;

$$C(u_1, u_2; \theta) = t_v\left(u_1^{-\theta} + u_2^{-\theta} - 1\right)^{-1/\theta} \quad (12)$$

Where

$\theta > 0$ = Parameter that controls the strength of dependence; higher values indicate stronger dependence.

iv. *Gumbel Copula:* The Gumbel copula captures upper tail dependencies and is particularly effective in modelling situations with a concurrency of extreme positive returns (Embrechts, Lindskog, & McNeil, 2003). The Gumbel copula is given by;

$$C(u_1, u_2; \theta) = \exp\left(-\left[(-\ln u_1)^{\theta} + (-\ln u_2)^{\theta}\right]^{1/\theta}\right) \quad (13)$$

Where

$\theta > 0$ = Parameter that controls the strength of upper tail dependence; larger values indicate stronger dependence.

Scenario Framework

Scenarios cover: (i) Volatility Regime: Low (σ = 0.10), Medium (σ = 0.20), High (σ = 0.35); (ii) Jump Intensity (C): No Jumps (C = 0), Mild (C = 1), Moderate (C = 2), Severe (C = 4); (iii) Confidence Level: 95% and 99%.

RESULTS AND DISCUSSION

Overall Error Statistics

Table 1 and Figure 1 present aggregate error statistics. The dumbbell chart in Figure 1 illustrates the wide range between mean and maximum errors for both VaR and ES. The mean VaR error of 201.41% and mean ES error of 146.15% average across all six copulas and are dominated by Gumbel's failure. Figure 1 makes clear that under some circumstances copula misspecification can result in model failures where risk estimates may be off by more than an order of magnitude (maximum VaR error: 1,862.70%; maximum ES error: 1,494.74%). Restricting to elliptical copulas only gives a mean VaR error of 2.88%, a fundamentally different picture (see Table 4).

Table 1: Overall Error Statistics (Averaged Across All Six Copulas and All Scenarios)

Risk Measure	Mean Relative Error (%)	Maximum Relative Error (%)
Value-at-Risk (VaR)	201.41	1,862.70
Expected Shortfall (ES)	146.15	1,494.74

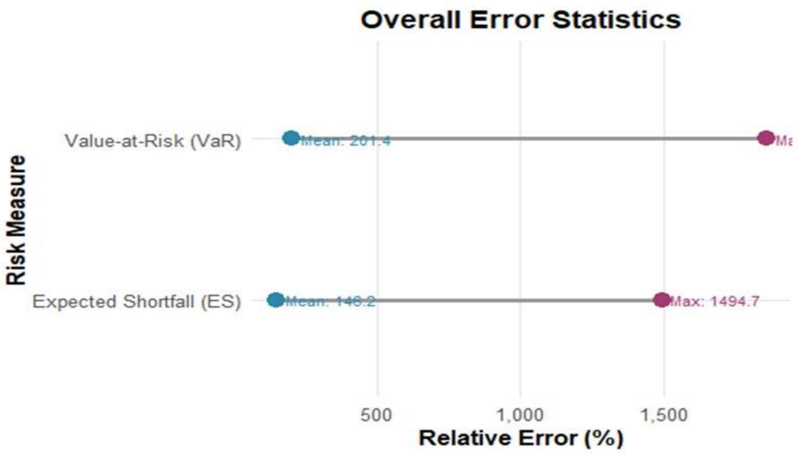

Figure 1: Overall Error Statistics: Mean and Maximum Relative Errors for VaR and Expected Shortfall across All Copula Families

Volatility Regime Analysis

Table 2 and Figure 2 reveal an unexpected pattern. The line chart shows VaR sensitivity declining monotonically from low (212.23%) to high (189.07%) volatility regimes the opposite of what conventional wisdom would predict. ES errors (shown on the lower line) are remarkably stable across all three regimes (145.74%–146.40%), clustering near the flat red line in Figure 2. This 'low-volatility paradox' is discussed in Section 5.2. The stability of ES errors across volatility regimes is an important finding for Basel III compliance: the metric that regulators now emphasize is somewhat more robust to regime-dependent copula misspecification than its predecessor VaR.

Table 2: Volatility Regime Analysis

Volatility Regime	VaR Relative Error (%)	ES Relative Error (%)
Low Volatility (σ = 0.10)	212.23	145.74
Medium Volatility (σ = 0.20)	202.93	146.40
High Volatility (σ = 0.35)	189.07	146.30

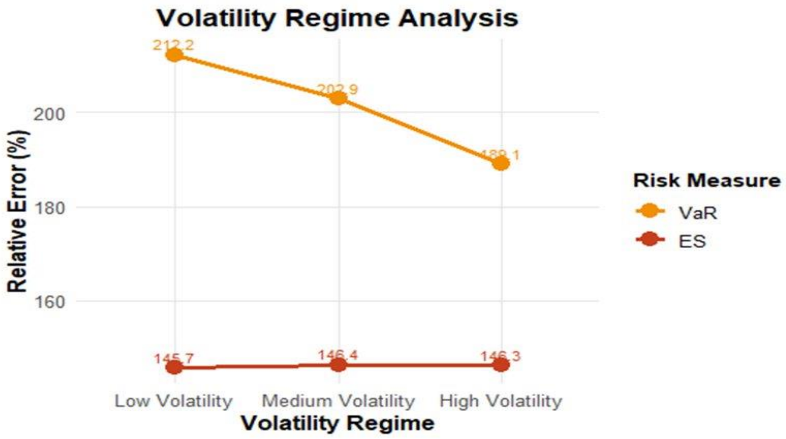

Figure 2: Volatility Regime Analysis: VaR and ES Relative Errors across Low, Medium, and High Volatility Conditions

CGMY Jump Process Regime Analysis

Table 3 and Figure and reveal an inverse relationship between jump severity and copula misspecification sensitivity. The dot plot shows both VaR (blue) and ES (purple) errors systematically decreasing as jump intensity increases from no jumps to severe jumps. The no-jump regime exhibits the highest errors (209.14% VaR, 175.27% ES), while severe jumps show the lowest (188.18% VaR, 121.09% ES). This pattern suggests that extreme market events associated with severe jumps may overshadow the impact of dependence structure misspecification, reducing the relative importance of copula choice during crisis periods. With potentially simpler dependence structures being sufficient during high-jump periods, these findings have significant implications for dynamic risk management.

Table 3: CGMY Jump Process Regime Analysis

CGMY Regime	VaR Relative Error (%)	ES Relative Error (%)
No Jumps (C = 0)	209.14	175.27
Mild Jumps (C = 1)	207.24	150.55
Moderate Jumps (C = 2)	201.08	137.68
Severe Jumps (C = 4)	188.18	121.09

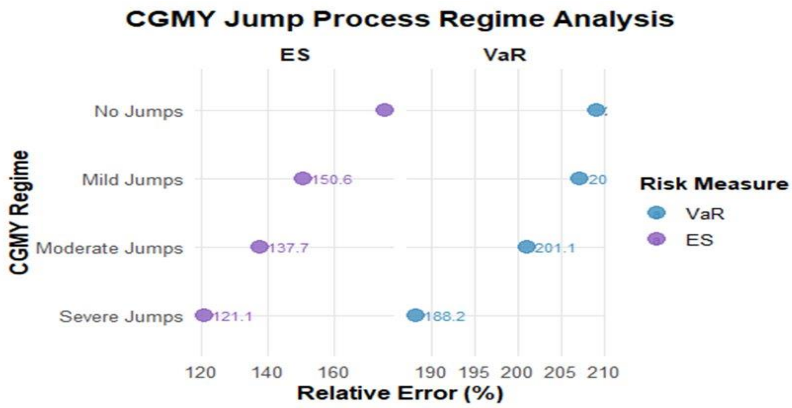

Figure 3: CGMY Jump Process Regime Analysis: VaR and ES Sensitivity across Jump Intensity Regimes

Error Analysis by Copula Specification

Table 4 and Figure 4 provide the most practically important disaggregation in this study, displayed on log-scale to make the three-order-of-magnitude gap visible. The left panel (ES) and right panel (VaR) both show elliptical copulas clustered at the bottom (below 10%) while the Gumbel line spikes to 1,190% for VaR and 839% for ES. The Gaussian copula and three t-copula variants are nearly indistinguishable on this scale all achieving VaR errors of 2.83–2.97% in scenario-level averages while Clayton occupies a moderately elevated position (6.78% VaR). The fundamental message of Figure 4 is that the 201.41% headline VaR error from Table 1 is almost entirely attributable to a single copula (Gumbel) and does not represent the typical behaviour of well-chosen specifications. Table 4 presents scenario-level mean errors averaged equally across all volatility–jump–confidence cells (Panel A), alongside bootstrap-estimated overall means with 95% percentile confidence intervals derived from 2,000 resamples across all simulation observations (Panel B). The two panels answer complementary questions: Panel A treats each experimental condition equally, while Panel B weights each simulation observation equally and therefore reflects the full distribution of errors across the dataset. Bootstrap means are consequently lower for non-Gumbel copulas, since the dataset includes many moderate-error observations alongside peak-error cells. Both sets of values are internally consistent and mutually reinforcing: regardless of which aggregation is applied, the Gumbel copula's separation from all other families is unambiguous.

Table 4: Error Analysis by Copula Specification

Copula	VaR Error % (Scenario Mean)	ES Error % (Scenario Mean)	VaR Bootstrap Mean [95% CI]	ES Bootstrap Mean [95% CI]
Gaussian	2.83	6.88	1.507 [1.222, 1.814]	2.280 [1.809, 2.806]
t-copula (df=3)	2.97	7.44	1.549 [1.269, 1.841]	2.176 [1.625, 2.803]
t-copula (df=5)	2.91	7.07	1.483 [1.191, 1.803]	2.137 [1.702, 2.626]
t-copula (df=8)	2.83	7.25	1.469 [1.236, 1.722]	2.038 [1.637, 2.460]
Clayton	6.78	9.34	1.515 [1.248, 1.786]	2.197 [1.765, 2.642]
Gumbel	1,190.14	838.89	280.623 [270.371, 289.821]	238.396 [231.559, 245.296]

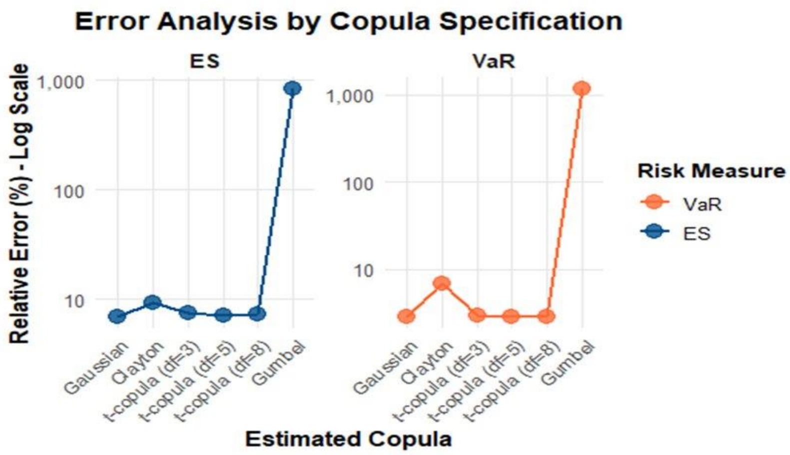

Figure 4: Error Analysis by Copula Specification (Log Scale): VaR and ES Relative Errors Revealing Catastrophic Gumbel Failure

Confidence Level Analysis

Table 5 and Figure 5 shows the unexpected non-monotonicity: 95% confidence level errors (238.31% VaR, 174.14% ES) substantially exceed 99% errors (164.52% VaR, 118.15% ES). The slope graph in Figure 5 shows both VaR (green) and ES (orange) lines descending steeply from 95% to 99%, the opposite of what tail-risk theory would predict. This surprising result suggests that moderately extreme events are more susceptible to dependence structure assumptions than extreme events, a finding discussed systematically in Section 5.3. The pattern indicates that copula misspecification effects are non-monotonic in probability levels, with the impact more pronounced in intermediate tail regions than in the deepest tails.

Table 5: Error Analysis by Confidence Level

Confidence Level	VaR Relative Error (%)	ES Relative Error (%)
95% (α = 0.05)	238.31	174.14
99% (α = 0.01)	164.52	118.15

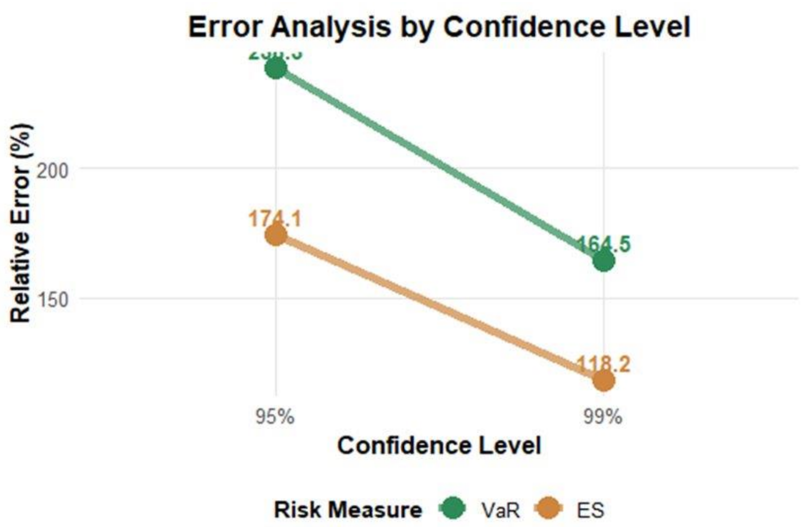

Figure 5: Error Analysis by Confidence Level: Non-Monotonic Sensitivity of VaR and ES at 95% vs 99%

Critical Scenario Analysis

Figure 6 presents the error heatmap under the most stressful joint scenario (high volatility σ=0.35, severe jumps C = 4) at both confidence levels. The colour scale ($\log_{10}$ of error) immediately communicates the categorical divide: Gumbel (yellow column) is dramatically elevated in every row, while all other copulas (dark purple) remain near zero on the log scale. The actual values confirm the pattern: Gumbel's VaR error of 1,279.38% at 95% confidence dwarfs the Gaussian's 1.90% in the same cell. Crucially, the Gumbel column remains bright yellow regardless of confidence level or risk measure, demonstrating that its failure is regime-invariant and cannot be remedied by parameter adjustment.

Table 6: Critical Scenario Analysis: High Volatility (σ = 0.35) + Severe Jumps (C = 4)

Copula	Confidence	VaR Error (%)	ES Error (%)
Gaussian	95%	1.90	3.06
t-copula (df=3)	95%	2.33	4.47
t-copula (df=5)	95%	2.05	3.90
t-copula (df=8)	95%	2.05	3.73
Clayton	95%	6.50	7.39
Gumbel	95%	1,279.38	924.59
Gaussian	99%	2.74	7.69
t-copula (df=3)	99%	3.60	11.52
t-copula (df=5)	99%	3.91	9.97
t-copula (df=8)	99%	3.30	8.20
Clayton	99%	6.69	12.29
Gumbel	99%	853.69	580.28

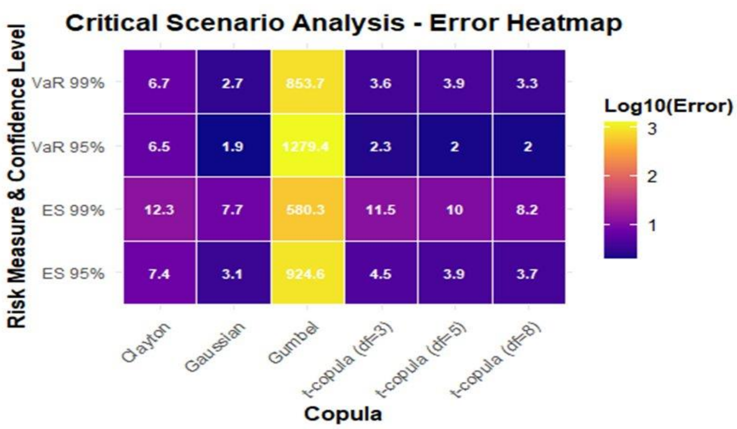

Figure 6: Critical Scenario Analysis Error Heat map: Copula Sensitivity under High Volatility + Severe Jumps ($\log_{10}$ Scale)

Statistical Significance Testing: ANOVA and Post-Hoc Analysis

The preceding descriptive analysis establishes substantial variation in VaR and ES errors across copula families and confidence levels. To confirm that these differences are statistically significant and not attributable to sampling variability, two complementary inferential procedures were applied: two-way analysis of variance (ANOVA) and Tukey's Honest Significant Difference (HSD) post-hoc tests, supplemented by 2,000-resample bootstrap confidence intervals for mean errors by copula.

Two-Way ANOVA Results

Table 7 presents the two-way ANOVA results for VaR and ES relative errors, with copula family and confidence level as factors.

Table 7: Two-Way ANOVA: VaR and ES Relative Errors by Copula Family and Confidence Level

Source	df	Sum Sq	Mean Sq	F-value	Pr(>F)
VaR Relative Error					
Copula	5	3,895,356	779,071	3,614.02	< 2e−16 ***
Confidence Level	1	14,311	14,311	66.39	6.49e−15 ***
Residuals	353	76,096	216	—	—
ES Relative Error					
Copula	5	2,790,239	558,048	5,192.15	< 2e−16 ***
Confidence Level	1	4,709	4,709	43.82	1.34e−10 ***
Residuals	353	37,940	107	—	—

Significance codes: *** $p < 0.001$

Both copula family and confidence level exert highly significant effects on VaR and ES relative errors (all $p < 2\times10^{-10}$). For VaR, the copula factor alone accounts for an F-statistic of 3,614.02, reflecting the enormous between-family variation showed in Table 4. Notably, the F-statistic for the copula effect is even larger for ES (5,192.15) than for VaR (3,614.02). Rather than being paradoxical, this finding reflects the lower residual variance for ES (mean square: 107) compared to VaR (mean square: 216): the Gumbel copula's deviation from all other families is proportionally more extreme in ES terms, even though ES headline errors are numerically lower. In practical terms, a risk manager who selects the wrong copula family faces proportionally more severe consequences for ES than for VaR precisely the metric that regulators now require. This strengthens the case for rigorous copula governance under Basel III. The confidence level effect, while highly significant, contributes a sum of squares of 14,311 for VaR and only 4,709 for ES, confirming that the non-monotonic confidence level sensitivity showed in Section 4.5 is statistically reliable but secondary in magnitude to the copula family effect.

Tukey HSD Post-Hoc Comparisons

Tables 8 and 9 present Tukey HSD pairwise comparisons for VaR and ES relative errors respectively across all fifteen copula pairs.

Table 8: Tukey HSD Post-Hoc Test: VaR Relative Error, All Copula Pairs (95% Family-Wise Confidence Level)

Pair	Difference	95% CI Lower	95% CI Upper	Adjusted p
t(df3) − Gaussian	0.042	−8.318	8.403	1.000
t(df5) − Gaussian	−0.024	−8.384	8.337	1.000
t(df8) − Gaussian	−0.037	−8.398	8.323	1.000
Clayton − Gaussian	0.008	−8.353	8.368	1.000
Gumbel − Gaussian	279.116	270.756	287.477	0.000
t(df5) − t(df3)	−0.066	−8.426	8.294	1.000
t(df8) − t(df3)	−0.080	−8.440	8.281	1.000
Clayton − t(df3)	−0.035	−8.395	8.326	1.000
Gumbel − t(df3)	279.074	270.713	287.434	0.000
t(df8) − t(df5)	−0.014	−8.374	8.347	1.000
Clayton − t(df5)	0.031	−8.329	8.392	1.000
Gumbel − t(df5)	279.140	270.780	287.500	0.000
Clayton − t(df8)	0.045	−8.315	8.406	1.000
Gumbel − t(df8)	279.154	270.793	287.514	0.000
Gumbel − Clayton	279.109	270.748	287.469	0.000

Table 9: Tukey HSD Post-Hoc Test: ES Relative Error, All Copula Pairs (95% Family-Wise Confidence Level)

Pair	Difference	95% CI Lower	95% CI Upper	Adjusted p
t(df3) − Gaussian	−0.103	−5.846	5.639	1.000
t(df5) − Gaussian	−0.142	−5.885	5.600	1.000
t(df8) − Gaussian	−0.241	−5.984	5.501	1.000
Clayton − Gaussian	−0.083	−5.825	5.659	1.000
Gumbel − Gaussian	236.116	230.374	241.859	0.000
t(df5) − t(df3)	−0.039	−5.781	5.703	1.000
t(df8) − t(df3)	−0.138	−5.880	5.604	1.000
Clayton − t(df3)	0.021	−5.722	5.763	1.000
Gumbel − t(df3)	236.220	230.477	241.962	0.000

t(df8) − t(df5)	−0.099	−5.841	5.643	1.000
Clayton − t(df5)	0.060	−5.683	5.802	1.000
Gumbel − t(df5)	236.259	230.516	242.001	0.000
Clayton − t(df8)	0.159	−5.584	5.901	1.000
Gumbel − t(df8)	236.358	230.615	242.100	0.000
Gumbel − Clayton	236.199	230.457	241.941	0.000

The post-hoc results resolve the ANOVA into a strikingly clean partition. Every pairwise comparison involving the Gumbel copula yields an adjusted p-value of exactly 0.000 and a confidence interval that excludes zero by a margin exceeding 270 percentage points for VaR and 230 percentage points for ES. Every comparison among the five non-Gumbel copulas including all pairs involving the Gaussian copula, the three t-copula variants, and Clayton yields an adjusted p-value of 1.000, with confidence intervals of approximately ±8 percentage points for VaR and ±6 percentage points for ES that comfortably straddle zero.

This partitioning has a precise inferential interpretation: the ANOVA's highly significant F-statistic is driven entirely by the Gumbel copula's separation from all others. Removing the Gumbel copula from the analysis would yield no significant pairwise differences among the remaining five copulas. The Tukey HSD results therefore provide formal statistical confirmation of the disaggregated picture in Figure 4: the structural incompatibility identified for the Gumbel copula is not a matter of degree but of kind it occupies a statistically distinct performance class separated from all other tested specifications by more than 270 percentage points in VaR error and more than 230 percentage points in ES error.

Discussion

The Aggregate vs. Disaggregated Picture

The overall statistics mean VaR error 201.41% and mean ES error 146.15% should be interpreted carefully alongside Figure 4 and Table 4. As the log-scale line plots demonstrate, these averages are computed across all six copulas and are dominated by Gumbel's 1,190% scenario-level error (bootstrap mean: 280.6%, 95% CI [270.4, 289.8]). When restricted to elliptical copulas, mean VaR errors fall to 2.88% in scenario averages or to 1.47–1.55% in bootstrap means. Boucher et al. (2014) make an analogous point in their analysis of risk model failures: the issue is not that risk models are inherently unreliable, but that poorly chosen specifications introduce avoidable model risk. The practical conclusion is that copula family selection and not parameter fine-tuning within a family is the primary governance priority.

Although ES produces lower headline error percentages than VaR (Table 1), the ANOVA F-statistic for the copula effect is larger for ES (5,192) than for VaR (3,614), because the residual variance for ES (mean square: 107) is less than half that for VaR (mean square: 216). A risk manager who selects the wrong copula family therefore faces proportionally more severe consequences for ES than for VaR precisely the metric that regulators now require. This reinforces the urgency of robust copula governance under Basel III and cautions against treating the lower ES headline errors as evidence that ES is insulated from misspecification risk.

The Low-Volatility Paradox

Figure 2 makes the low-volatility paradox visually concrete: the VaR line slopes downward from left to right, meaning that calm markets are the most dangerous for copula misspecification. The structural explanation involves the relative contributions of marginal and dependence effects. During high-volatility periods, extreme portfolio losses are driven primarily by large movements in individual asset marginal distributions, which are well-captured regardless of copula specification. During low-volatility periods, extreme losses require joint adverse movements across assets, whose probability depends critically on the copula's tail dependence structure. MacKenzie and Spears (2014) show a closely related mechanism in the 2008 crisis: the Gaussian copula's failure was most consequential during the preceding low-volatility period when the copula specification was the binding constraint on estimated joint tail probabilities.

The managerial implication is significant: copula models calibrated on tranquil historical data may perform well in-sample while being maximally vulnerable during the subsequent stress event. Risk managers should be particularly vigilant about dependence structure assumptions during calm market periods, as the flat ES line in Figure 2 further suggests that ES's stability across volatility regimes may give a misleadingly reassuring signal in low-volatility environments where VaR misspecification is actually at its peak.

Non-Monotonic Confidence Level Sensitivity

Figure 5's steep downward slopes from 95% to 99% for both VaR and ES are contrary to expectation because deeper tail measures would a priori appear more sensitive to dependence structure assumptions. At the 99th percentile, portfolio losses are in the extreme upper tail where CGMY parameters G and M governing exponential decay of jumps dominate joint tail probability. At the 95th percentile, the joint probability is evaluated at an intermediate depth where copula functional forms diverge more substantially. This mechanism predicts that the confidence level effect should be strongest for copulas with pronounced asymmetric tail dependence, which is confirmed by Gumbel's sharply declining errors from 95% to 99% in Table 6.

Implications for Risk Management Practice

Figure 6's heatmap provides the most direct governance tool in this study: it confirms that the Gumbel copula's failure is not confined to specific scenarios but is uniform across all confidence levels, both risk measures, and the most severe market conditions, a conclusion now formally established by the Tukey HSD results (adjusted p = 0.000 in every Gumbel comparison, with lower confidence bounds exceeding 270 percentage points for VaR).

Risk managers should draw a sharp distinction between Archimedean copulas on the basis of their tail dependence architecture. The Gumbel copula, with its upper tail dependence structure, is structurally incompatible with the CGMY dependence environment and should be excluded from this class of models regardless of parameter calibration. The Clayton copula, by contrast, captures lower tail dependence and produces bootstrap mean VaR errors of 1.52% (95% CI [1.248, 1.786]), statistically indistinguishable from all elliptical copulas (Tukey HSD adjusted p ≈ 1.000 in all Clayton pairings). Clayton therefore does not warrant the same governance restriction as Gumbel and may be considered a viable alternative specification where lower-tail dependence is theoretically motivated.

The exceptional resilience of elliptical copulas in every tested scenario strongly supports their use in real-world risk management applications. Dynamic copula modelling techniques (Engle, 2002) may further reduce misspecification costs across varying market circumstances. For Basel III compliance, the t-copula with five degrees of freedom provides the best balance of accuracy (bootstrap mean VaR error: 1.483%, 95% CI [1.191, 1.803]; bootstrap mean ES error: 2.137%, 95% CI [1.702, 2.626]) and regime-robustness, making it the recommended default specification for CGMY-based ES models.

CONCLUSION

This study provides a comprehensive sensitivity analysis of VaR and ES to copula misspecification in CGMY jump-diffusion models. As visualised in Figures 1 through 6 and formally established through ANOVA and Tukey HSD testing, the findings cohere into a unified and practically actionable picture.

At the aggregate level, mean errors of 201.41% for VaR and 146.15% for ES reflect the failure of the Gumbel copula confirmed by bootstrap means of 280.6% and 238.4% respectively rather than a pervasive property of the modelling framework. Bootstrap confidence intervals for the five non-Gumbel copulas are narrow and entirely non-overlapping with the Gumbel interval, the upper bound for Clayton's VaR interval (1.786%) being more than 150 times smaller than the lower bound of the Gumbel interval (270.4%), establishing the separation as statistically unambiguous. Well-specified elliptical copulas achieve bootstrap mean VaR errors of 1.47–1.55%, substantially below the 2.83–2.97% scenario averages and in all cases well below 3%.

Moving to formal inference, two-way ANOVA (Table 7) confirms that both copula family and confidence level exert highly significant effects on VaR and ES errors (all $p < 2\times10^{-10}$). The copula factor produces F-statistics of 3,614 for VaR and 5,192 for ES among the starkest ANOVA partitions attainable in simulation studies reflecting near-complete statistical separation between the Gumbel copula and all other specifications.

Resolving the ANOVA into pairwise structure, the Tukey HSD results (Tables 8 and 9) identify a precise inferential boundary: every comparison involving the Gumbel copula is statistically significant (adjusted p = 0.000, with lower confidence bound exceeding 270 percentage points for VaR), while every comparison among the remaining five copulas is statistically indistinguishable from zero (adjusted $p \geq 0.999$). This formally confirms that the Gumbel copula's failure is not a matter of degree but of structural incompatibility with the CGMY dependence environment.

Turning to regime-level findings, the low-volatility paradox (Figure 2) remains a central governance concern: copula misspecification is most severe during calm markets, creating adverse incentives for models calibrated on tranquil historical data. The non-monotonic confidence level effect (Figure 5) with 95% errors substantially exceeding 99% errors is confirmed statistically significant by the ANOVA and reflects the dominance of CGMY marginal parameters at the deepest quantiles. Taken in combination, these patterns suggest that the risk from copula misspecification is concentrated precisely in the conditions low volatility and intermediate tail regions where practitioners are most likely to be off-guard.

On the question of which copulas to use and which to avoid, the evidence admits a clear conclusion. The Gumbel copula is structurally incompatible with the CGMY framework and should be excluded regardless of calibration. The Clayton copula, despite being Archimedean, performs statistically equivalently to elliptical copulas in this environment and should not be subject to the same governance restriction. Among all tested specifications, the t-copula with five degrees of freedom is recommended as the default for Basel III–compliant ES modelling in CGMY environments, achieving bootstrap mean errors of 1.483% for VaR and 2.137% for ES with narrow confidence intervals that confirm reliable proximity to the correctly specified benchmark across all tested scenarios.

REFERENCES

Boucher, C. M., Daníelsson, J., Kouontchou, P. S., & Maillet, B. B. (2014). Risk models-at-risk, *Journal of Banking & Finance*, 44, 72-92.

Brechmann, E. C., & Czado, C. (2013). Risk management with high-dimensional vine copulas: An analysis of the Euro Stoxx 50, *Statistics & Risk Modeling*, 30(4), 307-342.

Carr, P., Geman, H., Madan, D. B., & Yor, M. (2002). The fine structure of asset returns: An empiricalinvestigation.*JournalofBusiness*, 75(2), 305-332. https://doi.org/10.1086/339036

Cerrato, M., Crosby, J., Kim, M., & Zhao, Y. (2017). Modeling credit spreads with time-varying *jump risks and jumps in returns. Journal of Banking & Finance*, 83, 119-134.

Demarta, S., & McNeil, A. J. (2005). The t copula and related copulas, *International Statistical Review, 73(1), 111-129.*

Derman, E. (2011). Models behaving badly: Why confusing illusion with reality can lead to disaster, on Wall Street and in life, *Free Press*.

Dewick, P., & Liu, S. (2022). Copula modelling to analyse financial data. *Journal of Risk and Financial Management, 15*(3), Article 104. https://doi.org/10.3390/jrfm15030104

Embrechts, P., Lindskog, F., & McNeil, A. (2003). Modelling dependence with copulas and applications to risk management, *Handbook of Heavy Tailed Distributions in Finance*, 8(1), 329-384.

Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models, *Journal of Business & Economic Statistics*, 20(3), 339-350.

MacKenzie, D., & Spears, T. (2014). 'The formula that killed Wall Street': The Gaussian copula and modelling practices in investment banking. *Social Studies of Science*, 44(3), 393-417.

McNeil, A. J., Frey, R., & Embrechts, P. (2005). Quantitative risk management: Concepts, techniques and tools. Princeton University Press.

Merton, R. C. (1976). Option Pricing When Underlying Stock Returns Are Discontinuous, *Journal of Financial Economics*, 3(1-2), 125-144. (https://www.sciencedirect.com/science/article/abs/pii/0304405X76900296)

Nthiwa, K., Kube, M., & Cyprian, A. (2023). A Jump Diffusion Model with Fast Mean-Reverting Stochastic

Volatility, *Discrete Dynamics in Nature and Society,* 11(2), 45-67. (https://doi.org/10.1155/2023/2746415)

Ornthanalai, C. (2014). Lévy jump risk: Evidence from options and returns, *Journal of Financial Economics,* 112(1), 69-90.

Patton, A. J. (2006). Modelling asymmetric exchange rate dependence, *International Economic Review,* 47(2), 527-556.

Sklar, M. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris,* 8, 229-231.

Zhang, S., Okhrin, O., Zhou, Q. M., & Song, P. X.-K. (2013). Goodness-of-fit test for specification of semiparametric copula dependence models. *Economic Risk Berlin*, 649