Bayesian copula models and asymmetric dependence in intraday equity returns

Gelly Mitrodima[1] Giulia Livieri[1] Matteo Iacopini[2]

[1]Department of Statistics, London School of Economics

[2]Department of Economics, Ca' Foscari University of Venice

CRETE 2026

Outline

1. Introduction
2. The model
3. Estimation
4. Simulation
5. Empirical results
   - Marginal model comparison
   - TV-ν copula results

Outline

1. Introduction
2. The model
3. Estimation
4. Simulation
5. Empirical results
   - Marginal model comparison
   - TV-ν copula results

Motivation

- Risk measures such as Value-at-Risk (VaR) and Expected Shortfall (ES) depend critically on the *joint* tail behaviour of a portfolio, which is typically asymmetric (Ang and Chen, 2002; Longin and Solnik, 2001)
- Lower tail dependence between equity returns tends to be stronger during downturns than upper tail dependence during rallies
- Capturing this asymmetry and its time variation is essential for accurate tail risk measurement
- Copulas (Sklar, 1959) separate marginal distributions from the dependence structure — a flexible framework for multivariate financial data

Existing approaches: GARCH-type marginals and a copula

- The skew-t copula (Azzalini and Capitanio, 2003) allows for asymmetric tail dependence, richer than the symmetric Student-t or the Generalized Hyperbolic (GH) skew-t copula (Demarta and McNeil, 2005); see Deng et al. (2025a) for a formal comparison
- Deng et al. (2025a): skew-t copula via Variational Bayes (VB) — substantial heterogeneity in asymmetric dependence across intraday equity pairs
- Ito and Yoshiba (2025): dynamic correlation matrix via corrected Dynamic Conditional Correlation (cDCC) applied to the skew-t copula — captures time-varying linear dependence but leaves tail heaviness and asymmetry fixed
- Oh and Patton (2023): time-varying factor loadings via Generalised Autoregressive Score (GAS) dynamics in a GH skew-t factor copula
- Deng et al. (2025b): tractable unified skew-t (TrUST) copula — in the standard skew-t copula the skewness parameter constrains all pairs to share a common asymmetry structure; TrUST relaxes this, allowing each pair its own asymmetry pattern

Our approach

1. Flexible nonparametric marginals that capture the daily shape and scale — leaving cleaner residual dependence for the copula
2. Time-varying (TV) skew-t copula: TV skew-t estimated via a coordinate ascent variational algorithm with Kalman smoother
3. Hybrid VB–Metropolis Hastings (MH) estimator for time invariant copula that avoids the boundary issues of standard VB for the degrees of freedom

- Constant degrees of freedom represent an *average* over all intraday observations, masking genuine within-period variation in joint tail dependence
- TV skew-t copula captures this variation endogenously, within a single coherent model
- The economic motivation is well-established: joint tail dependence intensifies during market stress (Longin and Solnik, 2001; Oh and Patton, 2023)

Outline

1. Introduction
2. The model
3. Estimation
4. Simulation
5. Empirical results
   - Marginal model comparison
   - TV-ν copula results

Bayesian Joint Quantile Time Series (BJQTS) marginals: intraday extension

- **BJQTS** (Griffin and Mitrodima, 2023): models the full conditional distribution via a piecewise Linearised Inverse Transformation (LIT)
- The LIT maps the conditional distribution nonparametrically via a piecewise linear monotone transformation — no parametric distributional assumption required
- We propose an intraday extension incorporating a diurnal adjustment via lagged realised variance

BJQTS marginals: diurnal adjustment

Diurnal adjustment. Average variance share of interval τ:

$$\hat{s}_{j,\tau} = \frac{1}{|\mathcal{T}_0| M} \sum_{t \in \mathcal{T}_0} \frac{r_{j,\tau,t}^2}{\sum_{l=1}^{M} r_{j,l,t-1}^2}$$

Volatility multiplier $\hat{\phi}_{j,\tau} = \sqrt{M \hat{s}_{j,\tau}}$ satisfies $M^{-1} \sum_{\tau} \hat{\phi}_{j,\tau}^2 = 1$ by construction.

Diurnally adjusted return: $y_{j,m} = r_{j,m} / \hat{\phi}_{j,\tau(m)}$

Notation:

- $t = 1, \ldots, T$: trading days; $\tau = 1, \ldots, M$ ($M = 26$): intraday intervals
- $m = (t-1)M + \tau$: stacked index; $r_{j,m} \equiv r_{j,\tau,t}$: log-return of asset $j = 1, \ldots, d$
- $\mathcal{T}_0 = \{1, \ldots, T\}$: in-sample days; $\hat{\phi}_{j,\tau}$: diurnal volatility multiplier

BJQTS uses lagged realised variance, a predetermined variable, conditioning it out in the marginal layer.

BJQTS marginals: Bayesian Joint Symmetric Absolute Value (B-JSAV) dynamics and probability integral transforms (PITs)

B-JSAV(1,1) dynamics. Local scale parameters for quantile region $i = 1, \ldots, K$ $(K = 9)$ on each side $(\pm)$ of the distribution of $y_{j,m}$:

$$\theta^{\pm}_{j,i,m+1} = \omega^{\pm}_{j,i} + \beta^{\pm}_{j,i}\,\theta^{\pm}_{j,i,m} + \gamma^{\pm}_{j,i}\,|y_{j,m}|$$

$\theta^{\pm}_{j,i,m} > 0$ controls the width of quantile region i; larger values indicate a more dispersed conditional distribution. Gaussian process priors smooth parameters across adjacent quantile regions.

PITs:

$$U_{j,m} = F_0\big(G_{j,m}(y_{j,m})\big)$$

where $G_{j,m}$ is the LIT transformation at posterior median parameters and F_0 is the standard normal centering distribution. Stacking into vectors: $U_m = (U_{1,m}, \ldots, U_{d,m})' \in [0,1]^d$ fed directly into the copula.

The skew-t copula

Implicit copula of the Azzalini and Capitanio (2003) multivariate skew-t:

$$c(U_m;\, \bar{\Omega}, \delta, \nu_m) = \frac{f(X_m;\, \bar{\Omega}, \delta, \nu_m)}{\prod_{j=1}^{d} f_j(X_{j,m};\, \delta_j, \nu_m)}$$

where $X_{j,m} = F_j^{-1}(U_{j,m};\, \delta_j, \nu_m)$ are obtained by inverting the univariate marginals.

Parameters:

- $\bar{\Omega}$: $d \times d$ correlation matrix with factor structure $\bar{\Omega} = V_1(GG' + D)V_1$ (Creal and Tsay, 2015)
- δ: skewness vector — governs *asymmetric* tail dependence
- $\nu_m > 2$: degrees of freedom — governs *tail heaviness* (ensures finite variance and a well-defined copula density)

TV skew-t copula

We model $x_m = \log(\nu_m - 2)$ as a latent autoregressive process of order 1 (AR(1)):

$$x_m = \mu_\nu + \phi_\nu(x_{m-1} - \mu_\nu) + \sigma_\nu \varepsilon_m, \quad \varepsilon_m \overset{\text{iid}}{\sim} N(0,1), \quad |\phi_\nu| < 1$$

ensuring $\nu_m = 2 + e^{x_m} > 2$ automatically; $x_{1:n} = (x_1, \ldots, x_n)$ denotes the latent state path over all $n = MT$ intraday observations.

Why AR(1) on the log scale?

- The log transformation enforces $\nu_m > 2$ without constraints
- AR(1) dynamics allow ν_m to be persistent — consistent with the slow-moving nature of tail risk (Oh and Patton, 2023) — while remaining stationary and mean-reverting to μ_ν
- The Gaussian state equation is compatible with the Kalman filter, enabling efficient smoothing of the latent path $x_{1:n}$

Priors: designed to be coherent with a Gamma$(3, 0.2)$ prior on $(\nu_m - 2)$, placing most mass on moderate to large degrees of freedom in the absence of data information (Deng et al., 2025a)

Joint model

Stacking returns into $r_m = (r_{1,m}, \ldots, r_{d,m})'$, via Sklar's theorem (Sklar, 1959):

$$\Pr(r_m \leq x \mid \mathcal{F}_{m-1}) = C_m\big(F_{1,m}(x_1), \ldots, F_{d,m}(x_d);\ \bar{\Omega}, \delta, \nu_m\big)$$

where C_m is the TV skew-t copula, $F_{j,m}$ is the BJQTS conditional marginal, and $x_{1:n} = (x_1, \ldots, x_n)$ with $x_m = \log(\nu_m - 2)$ is the latent state path over all $n = MT$ intraday observations. The joint log-likelihood is:

$$\mathcal{L}(\bar{\Omega}, \delta, x_{1:n}) = \sum_{m=1}^{n} \log c_m\big(U_m;\ \bar{\Omega}, \delta, \nu_m\big)$$

Two-stage estimation:

1. Fit BJQTS marginals via MCMC (Griffin and Mitrodima, 2023); compute PITs U_m at posterior median parameters $\Rightarrow$ treat U_m as fixed inputs to the copula
2. Estimate $(\bar{\Omega}, \delta)$ via coordinate ascent VB and smooth the latent path $x_{1:n}$ via Kalman smoother, iterating until convergence

Outline

1. Introduction
2. The model
3. Estimation
4. Simulation
5. Empirical results
   - Marginal model comparison
   - TV-ν copula results

Standard VB: limitations for ν

Standard VB approximates the full posterior by a Gaussian factor family — well-behaved for $(\bar{\Omega}, \delta)$ but we document systematic issues for the degrees of freedom ν

		Estimator	
Period	Marginal	Standard VB	Hybrid VB-MH
2017 (Low volatility), BAC–JPM–VIX			
	GARCH	15.37 (2.26)	4.02 (0.47)
	BJQTS	12.00 (—)	4.19 (0.36)
2020 (High volatility, COVID-19), BAC–JPM–VIX			
	GARCH	10.27 (2.45)	3.98 (0.40)
	BJQTS	2.00 (0.00)	4.12 (0.44)

2017 BJQTS: standard VB overestimates relative to profile likelihood peak. 2020 BJQTS: standard VB collapses to the lower boundary $\nu = 2$. Profile likelihood confirms data support lies in $\nu \in [4, 8]$ in both periods.

Hybrid VB-MH estimator

The Gaussian variational approximation works well for $(\bar{\Omega}, \delta)$ but is problematic for ν near its lower boundary — we therefore treat ν differently

Hybrid VB-MH: retain variational updates for $(\bar{\Omega}, \delta)$ while sampling ν via a MH random walk on the log scale

- At each VB iteration, propose $\nu^{\text{prop}} = \exp(\log \nu^{\text{cur}} + \epsilon)$, $\epsilon \sim N(0, \sigma^2_{\text{MH}})$
- Accept/reject via standard MH ratio on the exact copula log-likelihood — bypassing the variational approximation for ν
- σ^2_{MH} tuned to achieve acceptance rate 0.23–0.26 in both periods
- All other parameters $(\bar{\Omega}, \delta)$ updated via adaptive stochastic gradient ascent as in Deng et al. (2025a)

Two-pass warm start: first pass initialises ν near the data-supported region; second pass warm-starts from the first, allowing $(\bar{\Omega}, \delta)$ to converge with ν stabilised

Profile log-likelihood over ν: evidence for VB-MH

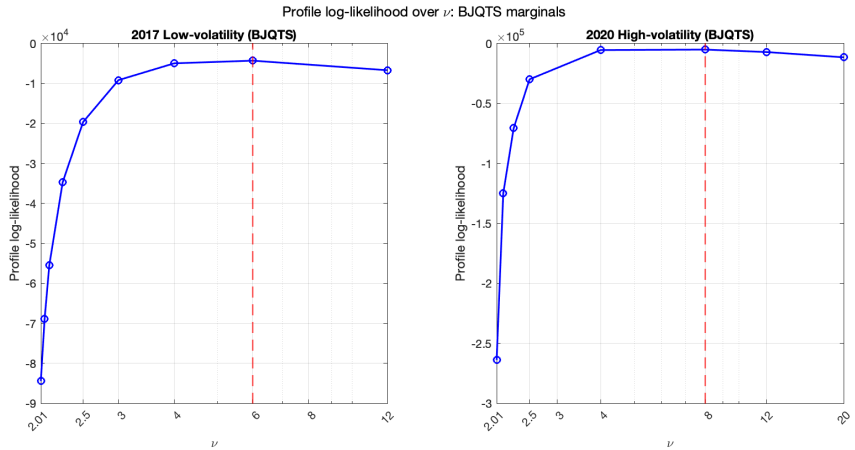

At each fixed ν, $(\bar{\Omega}, \delta)$ optimised via VB — an independent diagnostic of data-supported ν values

- Both periods: profile peaks near $\nu \in [6, 8]$, flat from 4 to 8
- Standard VB gives $\hat{\nu} = 12$ (2017) and $\hat{\nu} = 2$ (2020) — both far outside this region
- VB-MH recovers $\hat{\nu} \approx 4$ — within the region of likelihood support

Whether $\hat{\nu} \approx 4$ is the "true" value or reflects the flat likelihood region remains under investigation

Coordinate ascent algorithm for TV skew-t

Three-block coordinate ascent: optimise each parameter block in turn, holding the others fixed, until convergence. Initialised from the hybrid VB-MH static solution:

Block 1 — Static copula parameters

Update $(\bar{\Omega}, \delta)$ via adaptive stochastic gradient ascent with data augmentation for the skew-t representation

Block 2 — Latent state path $x_{1:n}$

Linearise the copula log-likelihood locally around the current state estimate. Run Kalman filter (forward pass) and Rauch–Tung–Striebel (RTS) smoother (backward pass) to obtain smoothed means and variances of x_m

Block 3 — AR(1) hyperparameters

Update $(\mu_\nu, \phi_\nu, \sigma_\nu)$ analytically from the smoothed sufficient statistics of the AR(1) state equation

Outline

1. Introduction
2. The model
3. Estimation
4. Simulation
5. Empirical results
   - Marginal model comparison
   - TV-ν copula results

Simulation: time invariant skew t — estimator comparison

Design: data generated from the skew-t copula with $d = 3$ assets, $n = 1{,}170$ observations, true $\nu_0 \in \{3, 4, 5\}$, and true $(\bar{\Omega}, \delta)$ fixed at the 2020 empirical posterior means — ensuring a realistic dependence structure. Three estimators compared: standard VB, hybrid VB-MH, and full Markov Chain Monte Carlo (MCMC) (Deng et al., 2025a) as standard benchmark.

	Standard VB		Hybrid VB-MH		Full MCMC	
ν_0	Mean	Std	Mean	Std	Mean	Std
3	6.87	0.49	4.04	0.28	5.11	0.05
4	7.86	0.78	4.04	0.37	5.09	0.05
5	8.44	0.91	4.21	0.43	5.13	0.05

- Standard VB: systematic overestimation across all designs
- Hybrid VB-MH and full MCMC both recover values within the region of likelihood support ($\nu \in [4, 8]$)
- The profile likelihood is flat from $\nu = 4$ to $\nu = 8$ — whether VB-MH recovers the true value or the likelihood provides weak discrimination in this region is an open question

Simulation: TV skew-t Kalman smoother — preliminary findings

Design: True $\phi_\nu = 0.795$, $\sigma_\nu = 0.328$, $\mu_\nu = 0.985$ taken from 2020 empirical estimates. $S = 10$ replications, $d = 3$, $n = 1{,}170$. True $(\bar{\Omega}, \delta)$ fixed at 2020 BJQTS posterior means.

Goal: assess whether the Kalman smoother recovers the latent ν_m path and the AR(1) hyperparameters $(\phi_\nu, \sigma_\nu, \mu_\nu)$.

Parameter	True	Mean	Std	RMSE
TV Kalman smoother				
$\hat{\phi}_\nu$	0.795	0.413	0.034	0.383
$\hat{\sigma}_\nu$	0.328	0.449	0.011	0.121
$\bar{\nu}_t$	—	5.048	0.147	—
Static hybrid VB-MH				
$\hat{\nu}$	—	4.102	0.106	—

- Average $\bar{\nu}_t$ well recovered — the smoother correctly identifies the level of tail thickness
- $\hat{\phi}_\nu$ underestimated — estimated persistence 0.413 well below true value 0.795
- One candidate: limited cross-sectional signal about state persistence at $d = 3$ — under investigation with larger d

Open questions

1. **Estimator for ν:** is VB-MH $\hat{\nu} \approx 4$ genuinely recovering the right region, or does the flat likelihood make this hard to assess? How does it compare to a longer MCMC run?

2. **Persistence recovery:** $\hat{\phi}_\nu$ underestimation is consistently observed in simulation — is it weak identification at $d = 3$? Does it affect the empirical ν_m path?

3. **Marginal allocation:** in a dynamic setting, how much temporal variation is allocated to the marginals vs the copula? Does the BJQTS vs GARCH distinction affect tail risk predictions or only the interpretation of dependence?

4. **Higher dimensions:** does TV skew t recovery improve with larger d? Weak identification at $d = 3$ suggests richer cross-sectional information may resolve the persistence underestimation — simulation with $d \in \{3, 10, 25\}$ planned

Outline

1. Introduction
2. The model
3. Estimation
4. Simulation
5. Empirical results
   - Marginal model comparison
   - TV-ν copula results

Data

15-minute intraday returns: Bank of America (BAC), JPMorgan Chase (JPM), CBOE VIX — from WRDS

2017: Low-volatility period

- 24 Oct – 22 Dec 2017
- 42 trading days $\times$ 26 intervals
- $n = 1{,}092$ observations
- Calm, trending market

2020: High-volatility period

- 11 Feb – 15 Apr 2020
- 45 trading days $\times$ 26 intervals
- $n = 1{,}170$ observations
- COVID-19 market crash

VIX: AR(1) fitted to VIX level series; residuals transformed via nonparametric kernel density estimator $\Rightarrow$ copula captures dependence between equity returns and *unexpected* VIX innovations

Same periods as Deng et al. (2025a) for direct comparability

Outline

1. Introduction
2. The model
3. Estimation
4. Simulation
5. Empirical results
   - Marginal model comparison
   - TV-ν copula results

What does the marginal choice imply?

Spearman correlations between equity PITs and VIX AR(1) innovation PIT

$$\hat{\rho}^{S}_{j,\mathrm{VIX}} = \mathrm{Corr}\big(\mathrm{rank}(U_{j,m}),\ \mathrm{rank}(U_{\mathrm{VIX},m})\big), j \in \{\mathrm{BAC},\ \mathrm{JPM}\}$$

Period	Asset	Raw	Model-free adj.	BJQTS	GARCH
2017	BAC	−0.347	−0.350	−0.035	−0.342
	JPM	−0.393	−0.401	−0.022	−0.388
2020	BAC	−0.581	−0.586	+0.021	−0.558
	JPM	−0.558	−0.565	+0.030	−0.540

Raw: unadjusted returns ranked. Model-free adj.: returns divided by average absolute return per interval — removes diurnal shape but not the daily volatility level. BJQTS and GARCH: full model-based PITs. $U_{\mathrm{VIX},m}$: AR(1) residual PIT (Deng et al., 2025a).

What does the marginal choice imply? Interpratation

- On high-VIX days, lagged daily variance is large $\Rightarrow$ dividing by $\hat{\phi}_{j,\tau}$ (the diurnal multiplier) conditions on the common daily volatility factor
- BJQTS absorbs the VIX–equity channel into the marginal layer — near-zero residual copula dependence
- GARCH (Engle and Sokalska, 2012) leaves this channel in the copula — a different but equally valid question
- *The two marginals answer different questions — distinguishing them is important for correct statements about tail dependence*

Outline

1. Introduction
2. The model
3. Estimation
4. Simulation
5. Empirical results
   - Marginal model comparison
   - TV-ν copula results

Time-invariant copula results: benchmark

Hybrid VB-MH estimator, BJQTS marginals, posterior standard deviations in parentheses. The static model assigns a single $\hat{\nu}$ per period — the TV skew-t copula reveals what this average conceals.

Pair	Spearman	Kendall	Major asym.	$\hat{\nu}$
2017 low-volatility				
BAC–VIX	−0.022	−0.015	0.040	
JPM–VIX	−0.013	−0.009	−0.023	4.19 (0.36)
BAC–JPM	0.674	0.502	−0.005	
2020 high-volatility				
BAC–VIX	0.009	0.006	−0.068	
JPM–VIX	0.007	0.005	−0.103	4.12 (0.44)
BAC–JPM	0.796	0.611	−0.188	

Time-invariant copula results: interpretation

Measures:

- **Spearman** $\hat{\rho}^S_{jk}$: rank correlation between $U_{j,m}$ and $U_{k,m}$
- **Kendall** $\hat{\tau}_{jk}$: fraction of concordant minus discordant pairs — more robust to tail behaviour than Spearman
- **Major asym.** $\Lambda_{\text{Major}} = \lambda^{LL} - \lambda^{UU}$: lower tail dependence minus upper tail dependence — negative means co-crashes more likely than co-rallies

Key findings:

- BAC–JPM correlation rises from 0.674 to 0.796 in 2020
- Major asymmetry for BAC–JPM shifts from -0.005 to -0.188 — co-crashes intensify during COVID-19
- VIX–equity near-zero under BJQTS, consistent with the marginal conditioning documented above
- Static $\hat{\nu} \approx 4.1$ in both periods — similar averages masking very different within-period dynamics

Dynamic tail dependence: the TV skew-t path

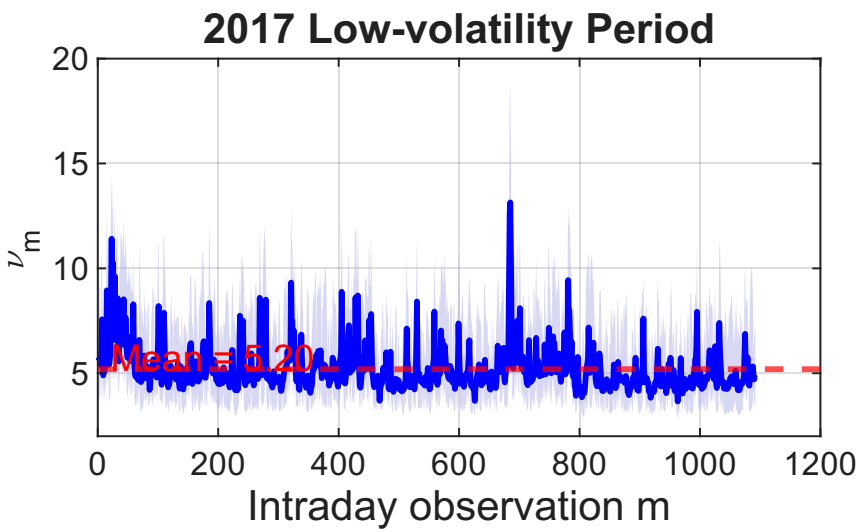

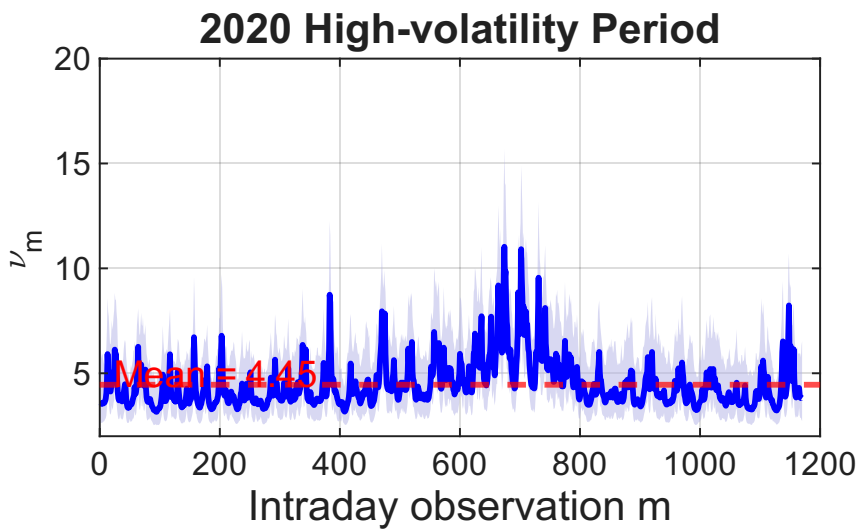

Smoothed ν_m path (dark blue) with 95% credible bands (light blue). Red dashed: posterior mean. Left: 2017. Right: 2020.

Dynamic tail dependence: what the path reveals

- Substantial within-period variation in both periods — tail thickness is not constant even within a single market regime
- Episodes of very low ν_m concentrated in 2020 — joint tail dependence intensifies during the COVID-19 crash
- Credible bands are wide, reflecting uncertainty in the latent path
- Static model assigns $\hat{\nu} = 4.12$ (2020) and $\hat{\nu} = 4.19$ (2017) — essentially identical, masking the rich within-period dynamics visible in the smoother

Dynamic tail dependence: distribution of ν_m

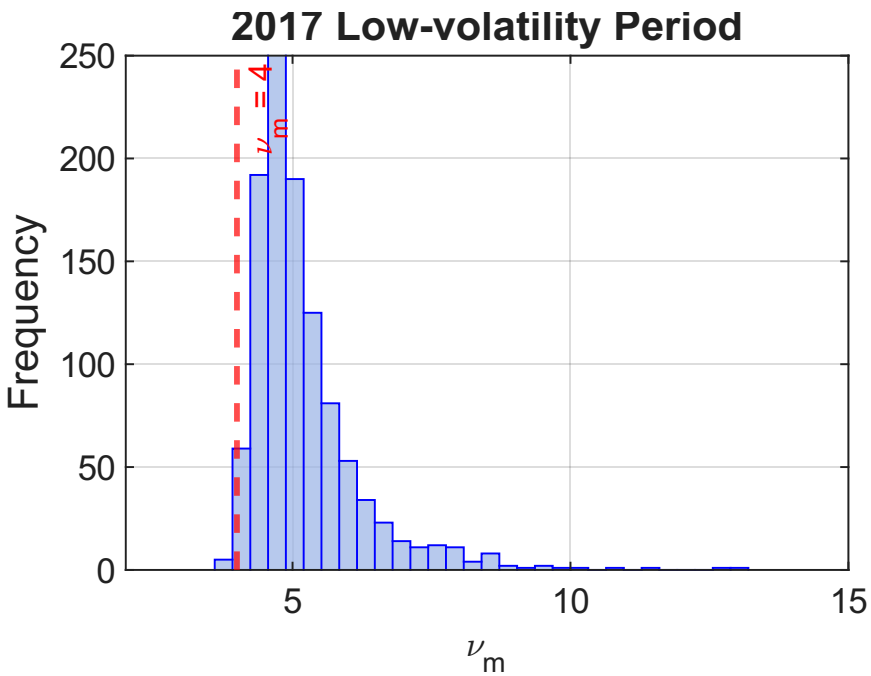

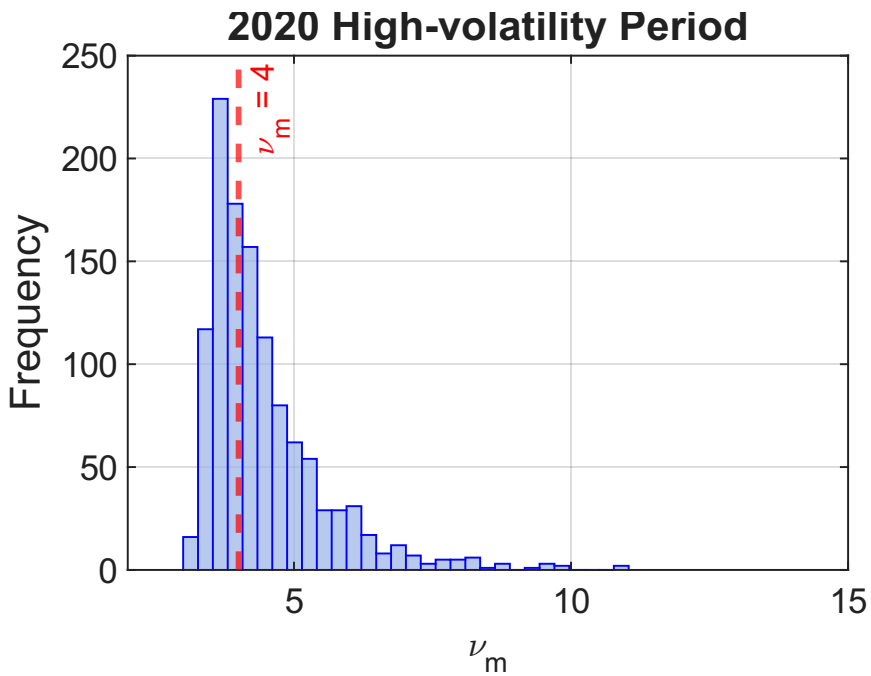

Distribution of smoothed ν_m. Dashed: $\nu_m = 4$. 2020: distribution shifted left, heavy left tail. 2017: concentrated above 4.

Dynamic tail dependence: descriptive findings, cont.

- **43% of 2020 intervals** have $\nu_m < 4$ vs only 0.7% in 2017
- The static model ($\hat{\nu} = 4.12$ vs 4.19) masks this substantial within-period variation
- Whether this variation is well-identified and what drives it remain under investigation

	2017	2020
$\bar{\nu}_t$	5.20	4.45
sd(ν_t)	1.00	1.07
ν_t range	[3.66, 13.14]	[3.16, 11.03]
Frac $\nu_t < 4$	0.7%	43%

Predictive distributions: 2020 high-volatility period

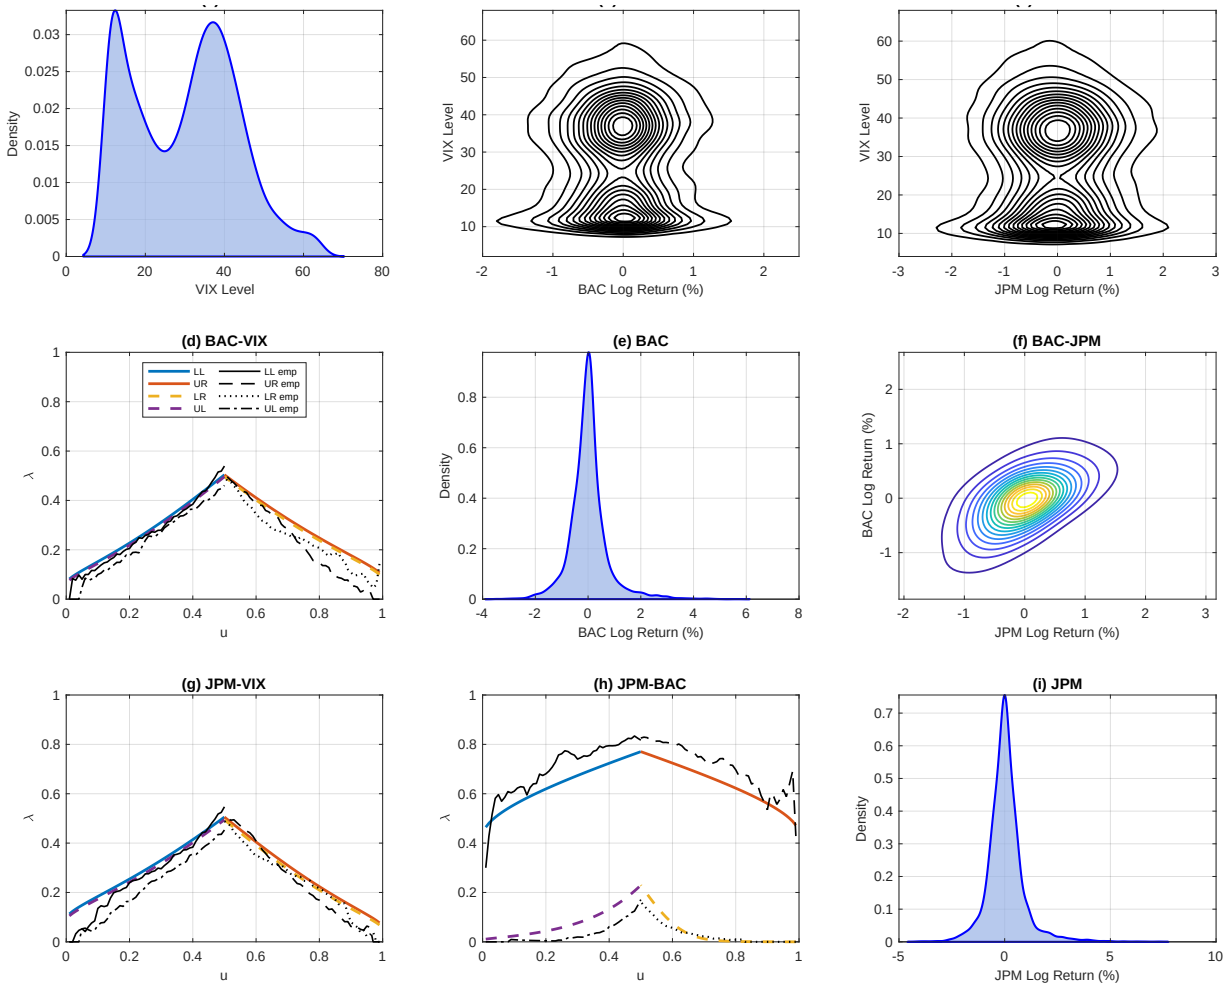

One-step-ahead predictive distributions, 2020, BJQTS marginals. Heavy joint tails and strong asymmetric dependence between BAC and JPM; near-zero residual dependence with VIX innovations, consistent with the marginal conditioning.

Summary and future work

What we propose:

1. BJQTS nonparametric marginals that condition on the daily volatility factor, leaving cleaner residual dependence for the copula
2. A hybrid VB-MH estimator that avoids the boundary issues of standard VB for the degrees of freedom ν
3. A TV skew t copula estimated via Kalman smoother revealing substantial within-period variation: 43% of 2020 intervals have $\nu_m < 4$ vs 0.7% in 2017

What the empirical results show:

The smoother reveals variation in joint tail dependence that a static model averages away. Whether this variation is well-identified and what drives it remain active research questions.

Future work:

- Risk measure evaluation and forecasting
- Simulation study to confirm theoretical intuition
- Extension to higher dimensions

LSE

References I

Ang, A. and Chen, J. (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63(3):443–494.

Azzalini, A. and Capitanio, A. (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society, Series B*, 65(2):367–389.

Creal, D. D. and Tsay, R. S. (2015). High-dimensional dynamic stochastic copula models. *Journal of Econometrics*, 189(2):335–345.

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

Deng, L., Smith, M. S., and Maneesoonthorn, W. (2025a). Large skew-t copula models and asymmetric dependence in intraday equity returns. *Journal of Business & Economic Statistics*, 43(2):269–285.

Deng, L., Smith, M. S., and Maneesoonthorn, W. (2025b). Tractable unified skew-t distribution and copula for heterogeneous asymmetries. arXiv:2505.10849.

Engle, R. F. and Sokalska, M. E. (2012). Forecasting intraday volatility in the US equity market: Multiplicative component GARCH. *Journal of Financial Econometrics*, 10(1):54–83.

References II

Griffin, J. E. and Mitrodima, G. (2023). A Bayesian quantile time series model for asset returns. *Journal of Business & Economic Statistics*. forthcoming.

Ito, K. and Yoshiba, T. (2025). Dynamic asymmetric tail dependence structure among multi-asset classes for portfolio management: Dynamic skew-t copula approach. *International Review of Economics and Finance*, 97:103724.

Longin, F. and Solnik, B. (2001). Extreme correlation of international equity markets. *Journal of Finance*, 56(2):649–676.

Oh, D. H. and Patton, A. J. (2023). Dynamic factor copula models with estimated cluster assignments. *Journal of Econometrics*, 237(2):105374.

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