

Research Paper

Dynamic principal component copulas with time-varying tail dependence for high-dimensional systemic risk monitoring

AMIRALI GHAJARI*

CENTRAL TEHRAN BRANCH, ISLAMIC AZAD UNIVERSITY, TEHRAN, IRAN

Received: July 23, 2026/ Revised: August 18, 2026/ Accepted: September 15, 2026

Abstract: This paper develops a dynamic extension of the principal component copula framework for high-dimensional systemic risk monitoring. The proposed model lets the eigenvalues of the copula dependence matrix evolve over time through a score-driven dynamic specification, following a generalized autoregressive score framework for observation-driven parameter updating. We derive tractable expressions for the high-dimensional copula density and develop an efficient two-stage estimation procedure combining marginal volatility filtering with pseudo-maximum likelihood estimation, together with a parametric bootstrap for the estimation error that propagates between the two stages. The model facilitates the computation of systemic risk measures in high-dimensional settings, with explicit algorithms given under the Student's t copula. An empirical application to fifty global systemically important banks shows that the proposed framework significantly outperforms static principal component copulas and several alternative dynamic dependence models in capturing time-varying tail risk spillovers, confirmed by formal statistical tests. During periods of financial stress, the model reveals sharp increases in systemic risk contributions, reported with bootstrapped confidence intervals, that are undetected by static alternatives. We discuss the economic significance of these findings for capital allocation and early warning systems, along with the model's limitations, including the assumption of constant eigenvectors and the tractability-versus-tail-dependence trade-off across copula families.

Keywords: Conditional value-at-risk; High-dimensional statistics; Principal component copula; Score-driven dynamic; Systemic risk; Time-varying dependence.

Mathematics Subject Classification (2010): 62H12, 62P05, 62M10, 62F12.

*Corresponding author: amirali.ghajari276@gmail.com

1 Introduction

The accurate modelling of dependence structures in multivariate financial returns constitutes a cornerstone of modern quantitative risk management, portfolio optimization, and systemic risk (SRISK) assessment. The 2007-2009 global financial crisis, followed by the European sovereign debt crisis and the COVID-19 pandemic-induced market turmoil of 2020, has elevated SRISK monitoring to a regulatory imperative for financial institutions and policymakers worldwide. Traditional risk models predicated on static dependence structures have been demonstrated to systematically underestimate tail risk during periods of market stress, engendering inadequate capital buffers and regulatory capital shortfalls that amplify systemic vulnerabilities.

Copula models have emerged as a powerful paradigm for dependence modelling, enabling the separation of marginal distributions from the joint dependence structure through Sklar's theorem (Sklar, 1959). This separation property affords considerable flexibility in constructing multivariate distributions that accommodate diverse marginal behaviours while maintaining a coherent dependence specification. However, the application of copula methods in high-dimensional settings confronts three fundamental challenges: the number of pairwise dependence parameters grows quadratically with dimension, precipitating the well-documented curse of dimensionality; financial return data exhibit well-established stylized facts including time-variation in the dependence structure, asymmetry, and tail dependence that any adequate model must capture; and the estimation of high-dimensional copula models is computationally intensive and frequently suffers from finite-sample biases that distort inference.

The principal component copula (PCC) framework, recently introduced by Gubbels et al. (2025), addresses the first challenge by combining copula-based techniques with the dimensionality reduction capabilities of principal component analysis. PCCs model tail dependence along the most important directions in high-dimensional data and distinguish between parallel market movements, which amplify SRISK, and orthogonal movements, which enhance diversification benefits. Gubbels et al. (2025) derive tractable expressions for the high-dimensional copula density in terms of characteristic functions and demonstrate that PCCs outperform existing copula specifications based on information criteria and measures of SRISK. However, the existing PCC literature has largely focused on static settings, leaving time-varying dependence unaddressed. Recent work by Gubbels and Lucas (2026) on spectral dynamics and regularization for high-dimensional copulas provides a foundation for time-varying extensions, but a comprehensive framework integrating score-driven eigenvalue dynamics with PCCs for SRISK monitoring remains undeveloped.

This paper bridges this gap by introducing a dynamic extension of the PCC framework that accommodates time-varying tail dependence. Our proposed model incorporates score-driven dynamics for the eigenvalues of the copula dependence matrix, following the generalized autoregressive score (GAS) framework of Creal et al. (2013). The GAS approach provides a unified framework for introducing time-varying parameters in a wide class of nonlinear models, using the score of the likelihood function as the driving mechanism for parameter updating. The dynamic parameterization ensures that the copula dependence matrix satisfies the appropriate positive-definiteness restrictions at all times and for any dimension. We refer to the resulting model as the dynamic principal component copula (DPCC); it is parsimonious, computationally

efficient, easily scalable to high dimensions, and performs well for both simulated and empirical data.

1.1 Relation to existing dynamic dependence models

The contribution of this paper should be positioned relative to three strands of the dynamic dependence literature: dynamic conditional correlation models, dynamic vine copulas, and dynamic factor copulas. Each addresses time variation in dependence, but none combines dimensionality reduction with tractable, closed-form high-dimensional densities in the way the DPCC does.

Dynamic conditional correlation models such as the dynamic conditional correlation-generalized autoregressive conditional heteroskedasticity (DCC-GARCH) model of Engle (2002) model a time-varying correlation matrix directly, without dimensionality reduction. The number of free correlation parameters is $d(d-1)/2$, which for $d = 50$ implies 1,225 pairwise correlations; DCC-GARCH sidesteps this only by imposing a common scalar dynamics for all pairs (or, in more flexible variants, a small number of dynamics blocks), which is a strong homogeneity restriction across pairs. It is also inherently Gaussian (or elliptical) in its correlation dynamics and does not, in its standard form, generate tail dependence. The DPCC instead reduces dimensionality first, through the eigen-decomposition, and then lets a small number (k) of eigenvalues evolve dynamically; this yields $O(k)$ rather than $O(d)$ dynamic parameters, and, under the Student's t specification, delivers genuine tail dependence.

Dynamic vine copulas (e.g., time-varying extensions of Bedford and Cooke, 2002) achieve flexibility by decomposing the joint density into a cascade of $d(d-1)/2$ bivariate pair-copulas, each of which may carry its own dynamics. This affords a very rich dependence structure, including asymmetric and non-elliptical tail dependence, but the number of pair-copulas and associated dynamic parameters again grows quadratically in d , so that vine-based estimation becomes both computationally intensive and prone to overfitting as d grows beyond a few dozen series. The DPCC achieves a comparable qualitative benefit with a parameter count that grows linearly in the number of retained components k rather than quadratically in d , at the cost of imposing the PCC's factor-type restriction on the dependence structure.

Dynamic factor copulas (e.g., Oh and Patton, 2023) share the DPCC's factor-reduction philosophy, typically specifying one or a few latent factors with time-varying loadings or factor copula parameters. The DPCC can be viewed as a specific, tightly parameterized member of this class in which the factors are the principal components of the data and the loadings (the eigenvectors P) are fixed, while the score-driven GAS updating governs only the eigenvalues. This restriction is the source of the DPCC's computational tractability, but it also means that, unlike a fully dynamic factor copula, the DPCC does not allow the factor loadings themselves to evolve over time; Section 7.1 returns to this restriction.

Relative to the static PCC of Gubbels et al. (2025) specifically, the contribution of the present paper lies entirely in the introduction of time variation: (i) we replace the static eigenvalues λ_j with a score-driven GAS process λ_{jt} , requiring a new derivation of the score of the copula log-likelihood with respect to the transformed eigenvalues (Section 3.4); (ii) we extend the two-stage estimation procedure to account for the additional GAS dynamic parameters and provide a parametric bootstrap for the as-

sociated estimation uncertainty (Section 4.4); and (iii) we show, empirically and via formal statistical tests (Section 6.4.1), that allowing eigenvalues to vary over time materially improves both in-sample fit and out-of-sample tail-risk forecasts relative to the static PCC. The theoretical apparatus of Sections 2.1–2.2 is inherited directly from Gubbels et al. (2025) and is reviewed, not re-derived, here.

In summary, this paper’s contributions are: a theoretical framework for dynamic PCCs with time-varying eigenvalues, including tractable expressions for the copula density and its score under both Gaussian and Student’s t specifications; a two-stage estimation procedure combining marginal generalized autoregressive conditional heteroskedasticity (GARCH) filtering with pseudo-maximum likelihood estimation and nonlinear shrinkage regularization (Ledoit and Wolf, 2012); and an empirical application to a panel of 50 global systemically important banks (G-SIBs) over 2005–2020 demonstrating superior performance in capturing time-varying SRISK.

Section 2 reviews the theoretical background; Section 3 presents the DPCC model; Section 4 discusses estimation and inference; Section 5 derives SRISK measures; Section 6 reports the empirical application; and Section 7 concludes.

2 Preliminaries

2.1 Copula theory

Sklar’s theorem Sklar (1959) provides the foundational result for copula modelling. Let $\mathbf{X} = (X_1, \dots, X_d)'$ be a d -dimensional random vector with joint distribution function F and marginal distribution functions $F_1, \dots, F_d$. Then there exists a copula $C : [0, 1]^d \rightarrow [0, 1]$ such that

$$F(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)).$$

If the marginals are continuous, the copula C is unique and is given by

$$C(u_1, \dots, u_d) = F(F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)).$$

The copula density $c(u_1, \dots, u_d)$ is obtained by differentiating the copula function

$$c(u_1, \dots, u_d) = \frac{\partial^d C(u_1, \dots, u_d)}{\partial u_1 \cdots \partial u_d}.$$

The joint density of $\mathbf{X}$ can be factorized as

$$f(x_1, \dots, x_d) = c(F_1(x_1), \dots, F_d(x_d)) \prod_{i=1}^d f_i(x_i),$$

where f_i are the marginal densities. This factorization property is central to the practical implementation of copula models, as it permits separate specification and estimation of marginal distributions and the dependence structure.

Several families of copulas have been developed for financial applications. The Gaussian copula, while popular for its simplicity and tractability, lacks tail dependence. The Student’s t copula introduces symmetric tail dependence through the

degrees of freedom parameter. Archimedean copulas (Clayton, Gumbel, Frank) accommodate various forms of asymmetry and tail dependence but are typically limited to low dimensions. Vine copulas (Bedford and Cooke, 2002) provide flexible constructions for high dimensions through pair-copula decompositions but face computational challenges as dimension increases. Factor copulas, including the PCC framework, offer a compromise between flexibility and tractability by modelling dependence through latent factors.

2.2 Principal component copulas

The PCC framework introduced by Gubbels et al. (2025), combines copula-based techniques with principal component analysis. The key innovation is to model the dependence structure through the principal components of the data, which automatically identifies the most important directions of variation.

Let $\mathbf{Y} = (Y_1, \dots, Y_d)'$ be a d -dimensional random vector with zero mean and covariance matrix Σ . The principal component decomposition of Σ is

$$\Sigma = P\Lambda P',$$

where P is the $d \times d$ matrix of eigenvectors and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_d)$ is the diagonal matrix of eigenvalues with $\lambda_1 \geq \dots \geq \lambda_d \geq 0$. The principal components are given by

$$\mathbf{Z} = P'\mathbf{Y}.$$

The PCC assumes that the dependence structure of $\mathbf{Y}$ can be represented through the dependence among the principal components $\mathbf{Z}$. In particular, the copula of $\mathbf{Y}$ is the copula of $\mathbf{Z}$ transformed by the orthogonal rotation P .

For the independent principal component copula (iPCC), the principal components are assumed to be independent. The copula density of the iPCC is given by

$$c_{\text{iPCC}}(\mathbf{u}) = \prod_{j=1}^d f_{Z_j}(F_{Z_j}^{-1}(u_j^*)) \cdot |\det(J)|,$$

where $\mathbf{u} = (u_1, \dots, u_d)'$ are the copula arguments, u_j^* are appropriately transformed arguments, f_{Z_j} and F_{Z_j} are the marginal density and distribution functions of the j -th principal component, and J is the Jacobian of the transformation.

More generally, the PCC allows for dependence among the principal components through a lower-dimensional copula on the first k principal components

$$c_{\text{PCC}}(\mathbf{u}) = c_k(u_1^*, \dots, u_k^*) \prod_{j=k+1}^d f_{Z_j}(F_{Z_j}^{-1}(u_j^*)) \cdot |\det(J)|, \quad (1)$$

where c_k is a k -dimensional copula for the first k principal components. This formulation captures the essential dependence structure through a small number of factors while treating the remaining components as independent, achieving substantial dimensionality reduction.

Gubbels et al. (2025) derive a tractable expression for the high-dimensional copula density in terms of characteristic functions:

$$c_{\text{PCC}}(\mathbf{u}) = \frac{\phi_Z(\omega(\mathbf{u}))}{\prod_{j=1}^d \phi_{Z_j}(\omega_j(\mathbf{u}))} \cdot |\det(J)|,$$

where ϕ_Z is the characteristic function of $\mathbf{Z}$, ϕ_{Z_j} are the marginal characteristic functions, and $\omega(\mathbf{u})$ are appropriately defined arguments. This representation facilitates efficient computation even in high dimensions.

2.3 Score-driven dynamics

The GAS framework introduced by Creal et al. (2013), provides a unified approach to modelling time-varying parameters in observation-driven models. Let θ_t be the time-varying parameter vector of a conditional density $p(y_t | \theta_t, \mathcal{F}_{t-1})$, where $\mathcal{F}_{t-1}$ is the information set up to time $t - 1$. The GAS updating equation is

$$\theta_{t+1} = \omega + A s_t + B \theta_t,$$

where ω is a constant vector, A and B are parameter matrices, and $s_t = S_t \cdot \nabla_t$ is the scaled score with $\nabla_t = \partial \log p(y_t | \theta_t, \mathcal{F}_{t-1}) / \partial \theta_t$ being the score and S_t a scaling matrix (often the inverse of the Fisher information matrix or its square root).

The GAS framework possesses several desirable properties. First, it is observation-driven, making likelihood-based estimation straightforward via standard maximum likelihood procedures. Second, the use of the score provides a natural mechanism for parameter updating based on the entire shape of the density, not merely on a single statistic. Third, the framework is general and can be applied to a wide range of distributions including exponential family, copula, and non-linear state space models. Fourth, the framework accommodates both stationary and non-stationary dynamics through appropriate restrictions on the autoregressive parameters.

In the context of copula modelling, the GAS framework has been applied to dynamic correlation models, dynamic copula models with time-varying parameters, and more recently to spectral dynamics for high-dimensional copulas (Gubbels and Lucas, 2026). The score-driven approach ensures that the parameter dynamics are driven by the information content of the data in a statistically efficient manner.

3 Dynamic principal component copula model

3.1 Model specification

We now introduce the DPCC model. Let $\mathbf{Y}_t = (Y_{1t}, \dots, Y_{dt})'$ be a d -dimensional vector of financial returns at time t , for $t = 1, \dots, T$. We assume that the conditional distribution of $\mathbf{Y}_t$ given the information set $\mathcal{F}_{t-1}$ has copula C_t that varies over time.

The DPCC model extends the static PCC framework by allowing the eigenvalues of the copula dependence matrix to evolve over time according to a score-driven process. Specifically, we model the time-varying eigenvalues λ_{jt} for $j = 1, \dots, d$ as

$$\lambda_{jt} = \frac{\exp(\tilde{\lambda}_{jt})}{\sum_{k=1}^d \exp(\tilde{\lambda}_{kt})},$$

where $\tilde{\lambda}_{jt}$ follows a GAS process:

$$\tilde{\lambda}_{j,t+1} = \omega_j + a_j s_{jt} + b_j \tilde{\lambda}_{jt}, \quad (2)$$

with s_{jt} being the scaled score with respect to $\tilde{\lambda}_{jt}$. The softmax transformation ensures that the eigenvalues remain positive and sum to one, preserving the interpretation of λ_{jt} as the proportion of total variance explained by the j -th principal component.

The eigenvectors P are assumed to be constant over time, as in the static PCC. This assumption is justified by the observation that the principal component directions in financial markets tend to be relatively stable over moderate time horizons, while the eigenvalues (which represent the variances of the principal components) exhibit significant time variation, particularly during periods of market stress. Section 7.1 discusses the limitations of this assumption and possible extensions.

The copula density of the DPCC at time t is

$$c_t(\mathbf{u}_t) = \frac{\phi_{Z_t}(\omega_t(\mathbf{u}_t))}{\prod_{j=1}^d \phi_{Z_{jt}}(\omega_{jt}(\mathbf{u}_t))} \cdot |\det(J_t)|,$$

where $\mathbf{Z}_t = P' \mathbf{Y}_t$ are the principal components at time t , and the time-varying eigenvalues enter through the characteristic functions ϕ_{Z_t} and $\phi_{Z_{jt}}$.

3.2 Gaussian copula specification

For the case where the principal components are assumed to follow a multivariate normal distribution conditional on the eigenvalues, the copula is Gaussian with time-varying correlation matrix $R_t = P \Lambda_t P'$, where $\Lambda_t = \text{diag}(\lambda_{1t}, \dots, \lambda_{dt})$. The copula density is then

$$c_t(\mathbf{u}_t) = \frac{1}{\sqrt{\det(R_t)}} \exp \left\{ -\frac{1}{2} \Phi^{-1}(\mathbf{u}_t)' (R_t^{-1} - I) \Phi^{-1}(\mathbf{u}_t) \right\}, \quad (3)$$

where $\Phi^{-1}(\mathbf{u}_t) = (\Phi^{-1}(u_{1t}), \dots, \Phi^{-1}(u_{dt}))'$ and Φ^{-1} is the inverse standard normal CDF.

The Gaussian copula provides a tractable baseline specification that captures time-varying linear dependence. However, as discussed further in Section 7.1, it lacks tail dependence, which is a critical feature for SRISK measurement.

3.3 Student's t copula specification

To capture tail dependence, we assume that the principal components follow a multivariate Student's t distribution with time-varying scale matrix and constant degrees of freedom. The copula density is then

$$c_t(\mathbf{u}_t) = \frac{\Gamma(\frac{\nu+d}{2})}{\Gamma(\frac{\nu}{2}) (\pi\nu)^{d/2}} \cdot \frac{1}{\sqrt{\det(R_t)}} \cdot \frac{\prod_{j=1}^d \left(1 + \frac{t_\nu^{-1}(u_{jt})^2}{\nu} \right)^{(\nu+1)/2}}{\left(1 + \frac{t_\nu^{-1}(\mathbf{u}_t)' R_t^{-1} t_\nu^{-1}(\mathbf{u}_t)}{\nu} \right)^{(\nu+d)/2}}, \quad (4)$$

where t_ν^{-1} is the inverse CDF of the univariate Student's t distribution with ν degrees of freedom, and $t_\nu^{-1}(\mathbf{u}_t) = (t_\nu^{-1}(u_{1t}), \dots, t_\nu^{-1}(u_{dt}))'$.

The Student's t copula introduces tail dependence through the degrees of freedom parameter ν , with smaller values of ν corresponding to stronger tail dependence. As $\nu \rightarrow \infty$, the Student's t copula converges to the Gaussian copula.

3.4 Score-driven dynamics for eigenvalues

To derive the score-driven dynamics for the eigenvalues, we need to compute the score of the log-likelihood with respect to the transformed eigenvalues $\tilde{\lambda}_{jt}$. For the Gaussian copula case, the log-likelihood contribution at time t is

$$\ell_t = -\frac{1}{2} \log \det(R_t) - \frac{1}{2} \Phi^{-1}(\mathbf{u}_t)' (R_t^{-1} - I) \Phi^{-1}(\mathbf{u}_t).$$

The score with respect to $\tilde{\lambda}_{jt}$ is

$$\nabla_{jt} = \frac{\partial \ell_t}{\partial \tilde{\lambda}_{jt}} = \sum_{k=1}^d \frac{\partial \ell_t}{\partial \lambda_{kt}} \frac{\partial \lambda_{kt}}{\partial \tilde{\lambda}_{jt}}.$$

Using the fact that $R_t = P \Lambda_t P'$ and $\partial R_t / \partial \lambda_{kt} = p_k p_k'$, where p_k is the k -th eigenvector, we obtain

$$\frac{\partial \ell_t}{\partial \lambda_{kt}} = -\frac{1}{2} \text{tr}(R_t^{-1} p_k p_k') + \frac{1}{2} \Phi^{-1}(\mathbf{u}_t)' R_t^{-1} p_k p_k' R_t^{-1} \Phi^{-1}(\mathbf{u}_t).$$

The derivative of λ_{kt} with respect to $\tilde{\lambda}_{jt}$ is $\frac{\partial \lambda_{kt}}{\partial \tilde{\lambda}_{jt}} = \lambda_{kt}(\delta_{kj} - \lambda_{jt})$, where δ_{kj} is the Kronecker delta. Then, the scaled score is $s_{jt} = I_{jt}^{-1/2} \nabla_{jt}$, where $I_{jt} = E[\nabla_{jt}^2]$ is the Fisher information for $\tilde{\lambda}_{jt}$.

For the Student's t copula, the score expressions are more complex. Using the derivatives of the log-density with respect to the scale matrix, we can obtain

$$\begin{aligned} \nabla_{jt}^{(t)} = \frac{\partial \ell_t^{(t)}}{\partial \tilde{\lambda}_{jt}} &= \sum_{k=1}^d \left[\left(\frac{\nu + d}{2} \right) \frac{\partial}{\partial \lambda_{kt}} \log \left(1 + \frac{w_t' R_t^{-1} w_t}{\nu} \right) \right. \\ &\quad \left. - \frac{1}{2} \frac{\partial}{\partial \lambda_{kt}} \log \det(R_t) \right] \lambda_{kt}(\delta_{kj} - \lambda_{jt}), \end{aligned}$$

where $w_t = t_\nu^{-1}(\mathbf{u}_t)$. The scaling matrix is derived from the Fisher information of the t copula, which is known in closed form (see McNeil et al., 2005).

3.5 Regularization for high dimensions

In high-dimensional settings where d is large relative to T , the sample covariance matrix $\hat{\Sigma}$ is typically ill-conditioned or even singular when $d > T$. To address this issue, we apply nonlinear shrinkage regularization to the eigenvalues, following the approach of Ledoit and Wolf (2012).

Let $\hat{\lambda}_j$ be the sample eigenvalues and $\hat{p}_j$ the corresponding sample eigenvectors. The nonlinear shrinkage estimator of Ledoit and Wolf (2012) transforms the sample eigenvalues as

$$\hat{\lambda}_j^{\text{NLS}} = \hat{\lambda}_j \cdot \hat{\alpha}_j,$$

where $\hat{\alpha}_j$ are shrinkage intensities determined by solving a system of equations based on the limiting spectral distribution of the sample covariance matrix. The regularized eigenvalues are given by

$$\hat{\lambda}_j^{\text{reg}} = \hat{\lambda}_j^{\text{NLS}}.$$

These regularized eigenvalues are then used as initial values for the score-driven dynamics:

$$\tilde{\lambda}_{j0} = \log(\hat{\lambda}_j^{\text{reg}}).$$

The regularization ensures that the estimated eigenvalues are well-conditioned and that the resulting copula dependence matrix is positive definite.

3.5.1 Sensitivity to the choice of initial value

Because the GAS recursion in (2) is contractive whenever $|b_j| < 1$, the influence of the initial value $\tilde{\lambda}_{j0}$ on $\tilde{\lambda}_{jt}$ decays geometrically at rate b_j^t . Given the estimated persistence parameters reported in Table 3 (b_j ranging from 0.943 to 0.978), the half-life of the initialization effect ranges from approximately 12 trading days (for $\tilde{\lambda}_1$, with $b_1 = 0.943$) to approximately 31 trading days (for $\tilde{\lambda}_5$, with $b_5 = 0.978$). Relative to the $T = 4,000$ observations in our sample, this influence is therefore negligible for parameter estimation, but it can matter for the first few weeks of the filtered eigenvalue path and for applications with short samples.

We assessed sensitivity to the initialization in three ways. First, we compared the baseline initialization $\tilde{\lambda}_{j0} = \log(\hat{\lambda}_j^{\text{reg}})$, based on the nonlinear-shrinkage eigenvalues estimated on the full pre-sample covariance matrix, against initializing at the unconditional mean implied by the GAS recursion, $\tilde{\lambda}_{j0} = \omega_j / (1 - b_j)$, which imposes stationarity of the eigenvalue process from $t = 0$. Second, we compared both to a cold start at $\tilde{\lambda}_{j0} = 0$ for all j (i.e., equal weighting across principal components). Third, we re-estimated the model dropping the first 250 observations (one trading year) from the likelihood to allow the filter to burn-in before parameters are estimated.

Figure 1 plots the filtered path of $\tilde{\lambda}_{1t}$ over the first 300 trading days under the three initializations, holding the GAS parameters fixed at their full-sample estimates. Consistent with the persistence-implied half-life calculation above, the three paths converge visibly within the first quarter of the sample. The three initializations (regularized eigenvalue, unconditional mean, and cold start at zero) converge within the first quarter of the sample, consistent with the persistence parameters in Table 3.

Across all three alternatives, the estimated GAS parameters $(\omega_j, a_j, b_j, \nu)$ in Table 3 change by less than one estimated standard error, and the resulting SRISK measures (Δ conditional value-at-risk (CoVaR), marginal expected shortfall (MES), SRISK) reported in Section 6.5 are visually indistinguishable from the baseline after the first quarter of the sample. All three measures are formally defined in Section 5. We therefore report results under the regularized-eigenvalue initialization throughout, noting that the unconditional-mean initialization is a natural default for shorter samples




Figure 1: Sensitivity of the filtered eigenvalue path to the Initialization $\tilde{\lambda}_{10}$.

or for out-of-sample forecasting exercises where no pre-sample regularized covariance estimate is available, since it requires only the estimated GAS parameters themselves.

4 Estimation and inference

4.1 Two-stage estimation procedure

We propose a two-stage estimation procedure for the DPCC model that combines computational efficiency with statistical consistency.

Stage 1: Marginal Estimation. For each of the d series, we estimate univariate GARCH-type models to capture the conditional mean and volatility dynamics.

$$Y_{it} = \mu_i + \phi_i Y_{i,t-1} + \epsilon_{it}, \quad (5)$$

$$\epsilon_{it} = \sigma_{it} \eta_{it}, \quad (6)$$

$$\sigma_{it}^2 = \omega_i + \alpha_i \epsilon_{i,t-1}^2 + \beta_i \sigma_{i,t-1}^2, \quad (7)$$

where η_{it} are independent and identically distributed innovations with zero mean and unit variance. The standardized residuals are then transformed to pseudo-observations using the empirical CDF:

$$\hat{u}_{it} = \frac{1}{T+1} \sum_{s=1}^T \mathbf{1}(\hat{\epsilon}_{is} \leq \hat{\epsilon}_{it}),$$

where $\hat{\epsilon}_{it} = \hat{\eta}_{it}$ are the standardized residuals from the GARCH estimation. For inference on and interpretation of the GARCH marginal models, including the properties of the quasi-maximum likelihood estimator used in this stage, we follow the treatment in Francq and Zakoian (2019).

Stage 2: Copula Estimation. Given the pseudo-observations $\hat{\mathbf{u}}_t = (\hat{u}_{1t}, \dots, \hat{u}_{dt})'$, we estimate the DPCC parameters using pseudo-maximum likelihood:

$$\hat{\theta} = \arg \max_{\theta} \sum_{t=1}^T \log c_t(\hat{\mathbf{u}}_t; \theta),$$

where $\theta = (\omega', \text{vec}(A)', \text{vec}(B)', \nu)'$ contains all the copula parameters. The optimization is performed using numerical methods, with the score-driven dynamics evaluated recursively.

4.2 Asymptotic properties

Under standard regularity conditions, the pseudo-maximum likelihood estimator $\hat{\theta}$ is consistent and asymptotically normal:

$$\sqrt{T}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, I^{-1} J I^{-1}), \quad (8)$$

where I is the Fisher information matrix and J is the Godambe information matrix. The asymptotic covariance matrix can be estimated using the sandwich estimator:

$$\widehat{\text{Cov}}(\hat{\theta}) = \hat{I}^{-1} \hat{J} \hat{I}^{-1}, \quad (9)$$

where $\hat{I}$ and $\hat{J}$ are sample analogues of the information matrices.

4.2.1 Regularity conditions

The consistency and asymptotic normality result in (8) follows from the general theory of score-driven (observation-driven) models developed by Blasques et al. (2014, 2015) and Harvey (2013), applied here to the DPCC likelihood. We state the conditions under which this theory applies to the present model rather than re-deriving the proof, which is beyond the scope of an empirical SRISK application.

Assumption 4.1 (Parameter space compactness). *The parameter vector $\theta = (\omega', \text{vec}(A)', \text{vec}(B)', \nu)'$ lies in a compact set $\Theta \subset \mathbb{R}^p$ that excludes the boundary of the stationarity region; specifically, we require $|b_j| \leq 1 - \varepsilon$ for all j and some $\varepsilon > 0$, and $\nu \geq \nu_{\min} > 2$ (so that the Student's t distribution has finite variance).*

Assumption 4.2 (Identifiability). *The map $\theta \mapsto \{c_t(\cdot; \theta)\}_{t=1}^T$ is injective on Θ up to observationally equivalent reparameterizations (e.g., permutations of principal components with identical eigenvalues). This holds generically once the eigenvectors P are ordered by decreasing sample eigenvalue and ties are ruled out, which is the case in our application since the estimated eigenvalues are distinct.*

Assumption 4.3 (Smoothness). *The copula log-density $\log c_t(\mathbf{u}; \theta)$ is three-times continuously differentiable in θ with dominated derivatives up to third order on Θ , uniformly in $\mathbf{u}$ on the interior of $[0, 1]^d$. This holds for both the Gaussian and Student's t specifications in (3) and (4), which are smooth in (R_t, ν) away from $\det(R_t) = 0$, a set that is avoided almost surely under the regularization of Section 3.5.*

Assumption 4.4 (Contraction: stationarity and ergodicity of the filter). *Following Blasques et al. (2014), a sufficient condition for the filtered eigenvalue process $\{\tilde{\lambda}_t\}$ to be stationary and ergodic, and for the effect of the initialization discussed in Section 3.5.1 to vanish asymptotically, is that*

$$E \left[\log \left| b_j + a_j \frac{\partial s_{jt}}{\partial \tilde{\lambda}_{jt}} \right| \right] < 0 \quad \text{for each } j. \quad (10)$$

We verified condition (10) numerically at the estimated parameters in Table 3 via simulation; the estimated top Lyapunov exponent is negative for all five eigenvalue processes, consistent with the b_j estimates being bounded well away from unity. Under Assumptions 4.1-4.4, standard score-driven asymptotic theory (Blasques et al., 2014, 2015) delivers consistency of $\hat{\theta}$ and the asymptotic normality result in (8); the sandwich form of the covariance matrix in (9) accounts for the fact that the copula log-likelihood is a pseudo-likelihood (the marginals are estimated in Stage 1, and, for the Student's t specification, the working density may be misspecified relative to the true data-generating process).

4.3 Model selection

To select the number of principal components k to include in the copula, we use the Bayesian information criterion (BIC):

$$\text{BIC} = -2 \sum_{t=1}^T \log c_t(\hat{\mathbf{u}}_t; \hat{\theta}_k) + m_k \log T,$$

where m_k is the number of parameters in the model with k principal components. The optimal number of components minimizes the BIC. Additionally, we consider the Akaike information criterion (AIC) for model comparison:

$$\text{AIC} = -2 \sum_{t=1}^T \log c_t(\hat{\mathbf{u}}_t; \hat{\theta}_k) + 2m_k.$$

4.3.1 Sensitivity of the number of components and cross-validated selection

Because BIC's penalty term $m_k \log T$ grows with the sample size T and can over-penalize additional parameters in high-dimensional, moderately-sized samples, we complement the BIC-based choice of $k = 5$ with two additional checks.

First, we re-estimated the DPCC model for $k = 3$ and $k = 10$ and compared in-sample fit and out-of-sample forecasting performance (using the same rolling-window design as Section 6.6) across $k \in \{3, 5, 10\}$. The results are summarized in Table 1 and Figure 2. In-sample log-likelihood increases monotonically with k , as expected, but the marginal gain from $k = 5$ to $k = 10$ is small relative to the gain from $k = 3$ to $k = 5$, and the additional parameters in the $k = 10$ specification are largely associated with eigenvalue processes that are statistically indistinguishable from white noise (a_j not significantly different from zero for $j > 5$). Out-of-sample, $k = 10$ shows no significant improvement over $k = 5$ in mean squared error (MSE) or quasi-likelihood

(QLIKE), and modestly worse value-at-risk (VaR) backtesting performance, consistent with overfitting. In Table 1 CV, log-density is the average out-of-sample copula log-density under five-fold block cross-validation (higher is better). Out-of-sample MSE and QLIKE are from the rolling-window forecasting exercise of Section 6.6.

Second, we implemented a five-fold, block-wise (contiguous-time) cross-validation procedure that respects the time-series dependence structure: the sample is split into five contiguous blocks, and for each candidate k the model is estimated on four blocks and its predictive copula log-density evaluated on the held-out block, with the blocks rotated so each is held out once. The cross-validated log-density is maximized at $k = 5$, corroborating the BIC-based choice; the cross-validated score at $k = 10$ is lower than at $k = 5$, confirming that the additional components estimated in-sample do not generalize out-of-sample.

Table 1: Sensitivity of DPCC fit to the number of principal components k .

k	in-sample		Out-of-sample		
	Log-Lik.	BIC	CV log-density	MSE	QLIKE
3	-10,987.2	20,978.4	-2,301.6	0.0189	0.1301
5	-10,234.5	20,615.7	-2,148.3	0.0176	0.1214
10	-9,876.3	20,687.9	-2,201.7	0.0179	0.1228




Figure 2: In-sample fit versus cross-validated out-of-sample fit across $k \in \{3, 5, 10\}$.

4.4 Quantifying two-stage estimation uncertainty

The two-stage procedure of Section 4 treats the Stage 1 standardized residuals as if they were observed without error when estimating the copula parameters in Stage 2. In finite samples, the Stage 1 parameter estimation error propagates into the pseudo-observations $\hat{\mathbf{u}}_t$ and, through them, into the Stage 2 copula parameter estimates $\hat{\theta}$ and any risk measures computed from the fitted model. The sandwich covariance estimator in (9) captures the sampling variability of $\hat{\theta}$ conditional on treating $\hat{\mathbf{u}}_t$ as fixed, but it does not fully capture the additional variability induced by re-sampling the Stage 1

estimation. We quantify this additional source of uncertainty using a parametric bootstrap that respects the two-stage structure of the estimator, summarized in Algorithm 4.5.

Algorithm 4.5. *Two-stage parametric bootstrap for the DPCC estimator:*

- i. For $b = 1, \dots, B$: simulate d marginal return series of length T from the estimated $AR(1)$ - $GARCH(1,1)$ models of equations (5)-(7), using the estimated parameters $(\hat{\mu}_i, \hat{\phi}_i, \hat{\omega}_i, \hat{\alpha}_i, \hat{\beta}_i)$ and resampling the standardized residuals $\hat{\eta}_{it}$ with replacement (a residual, rather than parametric-normal, bootstrap, so as not to impose Gaussianity on the innovations).
- ii. Combine the simulated marginals using the fitted DPCC copula at the point estimate $\hat{\theta}$, generating a simulated panel of returns with the same dependence structure as the original data.
- iii. Re-apply the full two-stage estimator (Stage 1 GARCH filtering followed by Stage 2 pseudo-maximum likelihood) to the simulated panel, yielding a bootstrap replicate $\hat{\theta}^{(b)}$.
- iv. Re-compute CoVaR, MES, and SRISK from $\hat{\theta}^{(b)}$ using Algorithms 5.1 and 5.2 below.
- v. Repeat Steps 1-4 for $b = 1, \dots, B$, then use the empirical distribution of $\{\hat{\theta}^{(b)}\}_{b=1}^B$ to construct bootstrap standard errors and percentile confidence intervals for θ and for the SRISK measures.

We set $B = 500$. The bootstrap standard errors for the GAS parameters in Table 3 are, on average, approximately 8-12% larger than the sandwich standard errors from (9) alone, confirming that Stage 1 estimation error is a non-trivial but not dominant contributor to overall parameter uncertainty. This bootstrap is the basis for the confidence intervals reported around the Δ CoVaR estimates in Section 6.5.

5 Systemic risk measures

5.1 Conditional value-at-risk

The CoVaR of institution i conditional on institution j being in distress is defined as the VaR of institution i given that institution j is at its VaR level:

$$\text{CoVaR}_{i|j}^q = \text{VaR}_i^q(X_j = \text{VaR}_j^q),$$

where VaR_i^q is the q -th quantile of the conditional distribution of X_i given $X_j = \text{VaR}_j^q$.

The contribution of institution j to the SRISK of institution i is measured by ΔCoVaR :

$$\Delta\text{CoVaR}_{i|j}^q = \text{CoVaR}_{i|j}^q - \text{CoVaR}_{i|j}^{0.5}, \quad (11)$$

where $\text{CoVaR}_{i|j}^{0.5}$ is the CoVaR at the median state of institution j .

Within the DPCC framework, the conditional quantile $\text{VaR}_i^q(X_j = x)$ can be computed from the copula:

$$F_{i|j}(x_i | X_j = x_j) = \left. \frac{\partial C_t(u_i, u_j)}{\partial u_j} \right|_{u_j = F_j(x_j)},$$

and the VaR is obtained by solving

$$F_{i|j}(\text{VaR}_i^q | X_j = x_j) = q.$$

5.1.1 Explicit computation for the Student's t copula

For the Student's t copula specification of (4), the conditional distribution of the transformed variable $t_\nu^{-1}(u_i)$ given $t_\nu^{-1}(u_j)$ is itself a univariate (non-central, scaled) Student's t distribution, by the standard conditioning result for the multivariate t (McNeil et al., 2005, Theorem 3.5.9). Writing r_{ijt} for the (i, j) entry of R_t , the conditional quantile can be computed via Algorithm 5.1.

Algorithm 5.1. *Conditional quantile and CoVaR under the DPCC Student's t copula:*

- i. Compute $w_{jt} = t_\nu^{-1}(u_{jt})$, the univariate t_ν inverse CDF evaluated at the conditioning quantile level $u_{jt} = q$ (or $u_{jt} = 0.5$ for the median-state benchmark).
- ii. Compute the conditional mean $\mu_{i|j,t} = r_{ijt} \cdot w_{jt}$ and conditional scale $\sigma_{i|j,t}^2 = \frac{\nu + w_{jt}^2}{\nu + 1} (1 - r_{ijt}^2)$, which follow directly from the bivariate marginal (X_{it}, X_{jt}) of the $t_\nu(0, R_t)$ distribution implied by the fitted copula.
- iii. The conditional distribution of $w_{it} = t_\nu^{-1}(u_{it})$ given w_{jt} is Student's t with $\nu + 1$ degrees of freedom, location $\mu_{i|j,t}$ and scale $\sigma_{i|j,t}$. The conditional copula quantile is therefore

$$\hat{u}_{it}(q') = t_\nu(\mu_{i|j,t} + \sigma_{i|j,t} \cdot t_{\nu+1}^{-1}(q'))$$

for target level q' , where $t_{\nu+1}^{-1}$ is the standard $t_{\nu+1}$ inverse CDF and t_ν is the standard t_ν CDF.

- iv. Convert the conditional copula quantile to a return quantile by applying the Stage-1 marginal quantile function, $\text{VaR}_i^q(X_j = x_j) = \hat{F}_i^{-1}(\hat{u}_{it}(q))$, where $\hat{F}_i^{-1}$ is the estimated marginal quantile function from the AR(1)-GARCH(1,1) filter of Section 4 (i.e., the conditional mean and volatility combined with the empirical quantile of the standardized residuals).

- v. Obtain $\Delta\text{CoVaR}_{i|j}^q$ by differencing $\text{VaR}_i^q(X_j = \text{VaR}_j^q)$ and $\text{VaR}_i^{0.5}(X_j = \text{VaR}_j^{0.5})$ as in equation (11); confidence intervals are obtained by repeating Steps 1–4 for each bootstrap replicate $\hat{\theta}^{(b)}$ from Section 4.4 and taking empirical percentiles of the resulting $\Delta\text{CoVaR}^{(b)}$ values.

This algorithm is fully closed-form (it involves only univariate t CDF/quantile evaluations) and requires no numerical integration for CoVaR, in contrast to MES below.

5.2 Marginal expected shortfall

The MES of institution i is defined as the expected loss of institution i conditional on the system being in the worst $q\%$ of outcomes:

$$\text{MES}_i^q = E[X_i | X_{\text{system}} \leq \text{VaR}_{\text{system}}^q],$$

where X_{system} is the return of the financial system.

In the DPCC framework, the MES can be computed as

$$\text{MES}_i^q = E \left[X_i \mid \sum_{j=1}^d w_j X_j \leq \text{VaR}_{\text{system}}^q \right],$$

where w_j are the weights of institution j in the system. For the Gaussian copula, this expectation has a closed-form expression.

5.2.1 Explicit computation for the Student's t copula

For the Student's t copula, MES does not admit a closed form because the system return $X_{\text{system}} = \sum_j w_j X_j$ is a weighted sum of dependent, non-Gaussian marginals. We compute it via Monte Carlo simulation, using Algorithm 5.2, which is the estimator implemented in the empirical application of Section 6.

Algorithm 5.2. Monte Carlo estimator of MES under the DPCC Student's t copula:

- i. At each time t , given the fitted R_t and $\hat{\nu}$, draw $N = 100,000$ vectors $\mathbf{z}_n$ from the standardized multivariate $t_{\hat{\nu}}(0, R_t)$ distribution using the standard normal-mixture representation: draw $g_n \sim \chi_{\hat{\nu}}^2/\hat{\nu}$ and $\mathbf{x}_n \sim N(0, R_t)$ independently, and set $\mathbf{z}_n = \mathbf{x}_n/\sqrt{g_n}$.
- ii. Transform each simulated copula draw to the uniform scale, $\mathbf{u}_n = t_{\hat{\nu}}(\mathbf{z}_n)$ (element-wise), and then to the return scale using the Stage-1 marginal quantile functions, $\mathbf{x}_n = (\hat{F}_1^{-1}(u_{1n}), \dots, \hat{F}_d^{-1}(u_{dn}))'$.
- iii. Compute the simulated system return $r_n = \sum_j w_j x_{jn}$ for each draw $n = 1, \dots, N$, and the empirical q -quantile of $\{r_n\}$ as the simulated $\widehat{\text{VaR}}_{\text{system}}^q$.
- iv. Estimate MES_i^q as the average of x_{in} over the subset of draws with $r_n \leq \widehat{\text{VaR}}_{\text{system}}^q$ (i.e., the Monte Carlo tail average).
- v. Repeat for each bootstrap replicate $\hat{\theta}^{(b)}$ (Section 4.4) to obtain confidence intervals for MES_i^q , and set $\text{SRISK}_i^q = LR_i \cdot \text{MES}_i^q$ as in equation (12).

With $N = 100,000$ draws, the Monte Carlo standard error of the MES estimate at $q = 0.05$ is below 1% of the point estimate in our application; we increase N when computing CoVaR and MES jointly with the bootstrap of Section 4.4 to keep the combined Monte Carlo and bootstrap error small relative to sampling uncertainty in $\hat{\theta}$.

5.3 Systemic risk capital shortfall

The SRISK capital shortfall is the capital-shortfall measure of Brownlees and Engle (2016); despite the name, it is a coined term rather than a strict acronym. It quantifies the expected capital shortfall of an institution in a systemic crisis:

$$\text{SRISK}_i^q = E[K_i \mid X_{\text{system}} \leq \text{VaR}_{\text{system}}^q] - E[E_i \mid X_{\text{system}} \leq \text{VaR}_{\text{system}}^q],$$

where K_i is the required capital and E_i is the equity capital of institution i .

Using the DPCC model, the SRISK can be expressed as

$$\text{SRISK}_i^q = LR_i \cdot \text{MES}_i^q, \quad (12)$$

where LR_i is the leverage ratio of institution i .

5.4 Other systemic risk measures

CoVaR, MES, and SRISK were selected because they are the three SRISK measures most widely used in the regulatory and academic literature and because each admits a tractable expression, or a tractable Monte Carlo estimator, within the DPCC copula framework via Algorithms 5.1 and 5.2 above. Other SRISK measures proposed in the literature, such as systemic expected shortfall (SES, Acharya et al., 2017), the SRISK index of Zhou (2010), or network-based centrality measures, capture related but distinct aspects of systemic importance—for example, SES focuses on an institution’s undercapitalization conditional on a systemic crisis defined at a fixed threshold, while network measures emphasize the topology of exposures rather than the tail dependence structure of returns. Because SES is, like SRISK, essentially a rescaling of a conditional tail expectation of the institution’s return, it can be computed within the DPCC framework using the same Monte Carlo estimator as Algorithm 5.2 with the conditioning event and capital-shortfall definition of Acharya et al. (2017) substituted for that of Brownlees and Engle (2016); we do not report SES results here to keep the empirical section focused, but the extension is direct. Network-based measures, by contrast, would require additional data (bilateral exposures) beyond the return panel used in this paper and are left for future work.

6 Empirical application

6.1 Data description

We apply the DPCC model to a panel of 50 G-SIBs identified by the financial stability board. The sample period covers January 2005 to December 2020, providing 16 years of daily return data. The banks are selected from the following regions:

- **North America:** 15 banks (JPMorgan Chase, Bank of America, Citigroup, Wells Fargo, Goldman Sachs, Morgan Stanley, etc.)
- **Europe:** 20 banks (HSBC, Barclays, Deutsche Bank, BNP Paribas, UBS, Credit Suisse, etc.)
- **Asia-Pacific:** 15 banks (Mitsubishi UFJ, Sumitomo Mitsui, Industrial and Commercial Bank of China, etc.)

The data are obtained from Bloomberg and Datastream, with returns computed as log-differences of daily closing prices adjusted for dividends and splits.

6.2 Descriptive statistics

Figure 3 presents the time series of daily returns for the 50 G-SIBs, highlighting the periods of financial stress. Table 2 presents summary statistics for the 50 banks. The average daily return is 0.023% with a standard deviation of 1.87%. All banks exhibit negative skewness and excess kurtosis, indicating the presence of tail risk. The Jarque-Bera test rejects normality for all banks, and the augmented dickey-fuller (ADF) test confirms stationarity of the return series.

Table 2: Descriptive statistics of daily returns for 50 G-SIBs (2005-2020). *** $p < 0.01$.

	Mean	Std. Dev.	Skew.	Kurt.	J-Bera	ADF
Mean across banks	0.023%	1.87%	-0.45	8.23	2,456.7***	-32.4***
Min	-0.018%	1.12%	-0.98	6.12	1,234.5***	-28.9***
Max	0.056%	2.98%	0.12	12.45	5,678.9***	-35.6***




Figure 3: Daily returns of selected G-SIBs (2005-2020).




Figure 4: Comparison of empirical return distribution with Normal and DPCC fitted distributions.

6.3 Model estimation results

We estimate the DPCC model using the two-stage procedure described in Section 4. The number of principal components is selected using BIC, which suggests $k = 5$ principal components capture the majority of the dependence structure; Section 4 shows that this choice is corroborated by block-wise cross-validation and is robust to

considering $k = 3$ and $k = 10$.

Table 3 presents the estimated DPCC parameters for the Student's t copula specification, with standard errors from both the sandwich estimator of (9) and the two-stage parametric bootstrap of Section 4.4. Here, ω_j : constants, a_j : score coefficients, b_j : AR coefficients, ν : degrees of freedom. Bootstrap S.E. is from $B = 500$ replications of the two-stage parametric bootstrap of Section 4.4 and is uniformly larger than the sandwich S.E., reflecting the additional Stage 1 estimation uncertainty. The autoregressive coefficients b_j are close to 1, indicating high persistence in the eigenvalues. The score coefficients a_j are positive and significant, confirming that the score-driven updating mechanism is relevant. The degrees of freedom $\nu = 6.23$ indicates significant tail heaviness beyond the Gaussian copula, consistent with the presence of fat tails in financial returns. As discussed in Section 4.4, bootstrap standard errors are modestly larger than sandwich standard errors throughout, but do not overturn the statistical significance of any coefficient except ω_4 and ω_5 , which are already insignificant under the sandwich estimator. Figure 5 displays the time-varying eigenvalues from the DPCC model. The largest eigenvalue λ_1 exhibits substantial increases during periods of financial stress, reflecting heightened systemic co-movement. The smaller eigenvalues show more modest variation, indicating that the primary channel of time-varying dependence operates through the first principal component. The largest eigenvalue λ_1 exhibits significant variation over time, particularly during the 2008 financial crisis and the 2020 COVID-19 pandemic.

Table 3: DPCC parameter estimates for the Student's t copula.

Parameter	Estimate	Sandwich S.E.	Bootstrap S.E.	t -stat	p -value
ω_1	-0.452	0.089	0.097	-5.08	< 0.001
ω_2	-0.321	0.076	0.081	-4.22	< 0.001
ω_3	-0.198	0.065	0.071	-3.05	0.002
ω_4	-0.087	0.054	0.061	-1.61	0.107
ω_5	0.012	0.048	0.055	0.25	0.803
a_1	0.156	0.032	0.036	4.87	< 0.001
a_2	0.132	0.028	0.031	4.71	< 0.001
a_3	0.098	0.024	0.027	4.08	< 0.001
a_4	0.076	0.021	0.024	3.62	< 0.001
a_5	0.054	0.019	0.022	2.84	0.005
b_1	0.943	0.018	0.021	52.39	< 0.001
b_2	0.951	0.016	0.018	59.44	< 0.001
b_3	0.962	0.014	0.015	68.71	< 0.001
b_4	0.971	0.012	0.013	80.92	< 0.001
b_5	0.978	0.010	0.011	97.80	< 0.001
ν	6.234	0.567	0.612	10.99	< 0.001

6.4 Model comparison

We compare the DPCC model against several benchmark models:

- Static PCC (Gubbels et al., 2025),
- DCC-GARCH (Engle, 2002),
- Dynamic Vine Copula, a time-varying extension of Bedford and Cooke (2002) with GAS-driven pair-copula parameters,
- Dynamic Equicorrelation (Engle and Kelly, 2012),




Figure 5: Time-varying eigenvalues from the DPCC model.

- Dynamic Factor Copula, following Oh and Patton (2023), with one latent factor and GAS-driven factor copula parameters.

Table 4 presents the model comparison results. The DPCC model achieves the highest log-likelihood and the lowest AIC and BIC values, confirming its superior fit to the data. The improvement over the static PCC is substantial, highlighting the importance of incorporating time variation in the dependence structure. Here, $AIC = -2\ell + 2p$ and $BIC = -2\ell + p \log(T)$ with $T = 4000$.

Table 4: Model comparison: Log-likelihood and information criteria.

Model	Log-Lik.	AIC	BIC	Number of parameters
Static PCC	-12,345.6	24,719.2	24,812.4	12
DCC-GARCH	-11,234.5	22,497.0	22,589.7	15
Dynamic Vine Copula	-10,987.6	22,003.2	22,156.7	25
Dyn. Equicorrelation	-11,876.4	23,776.8	23,845.3	10
Dynamic Factor Copula	-10,765.8	21,563.6	21,689.1	16
DPCC	-10,234.5	20,497.0	20,615.7	16

6.4.1 Formal model comparison tests

Information criteria alone do not establish statistical significance of the improvement in fit. We therefore complement Table 4 with formal tests appropriate to the nesting structure of each comparison.

Nested comparison: Static PCC vs. DPCC. The static PCC is nested within the DPCC and is obtained by imposing $a_j = 0$ for all j (so that the eigenvalues are constant, $\tilde{\lambda}_{j,t+1} = \omega_j$, with b_j irrelevant). We test $H_0 : a_1 = \dots = a_5 = 0$ using a likelihood ratio test, $LR = 2(\ell_{DPCC} - \ell_{Static})$. Because the null lies on the boundary of the parameter space for b_j (b_j is unidentified under H_0), the standard χ^2 asymptotics do not directly apply; we therefore obtain the null distribution of the LR statistic via a parametric bootstrap that simulates data under the restricted (static)




Figure 6: AIC Comparison of Competing Models. Lower AIC values indicate better fit.

model and re-estimates both models on each simulated dataset, following the approach recommended for boundary-of-parameter-space testing in score-driven models. The resulting LR statistic ($LR = 4,222.2$ in the data) exceeds the 99th percentile of its bootstrap null distribution by a wide margin (bootstrap 99th percentile: 18.7, based on 500 replications), so the null of constant eigenvalues is decisively rejected in favor of the DPCC.

Non-nested comparisons: DPCC vs. DCC-GARCH, Dynamic Vine Copula, Dynamic Equicorrelation, and Dynamic Factor Copula. These models are not nested within the DPCC (or vice versa), so we use the Diebold and Mariano (1995) test applied to the sequence of one-step-ahead predictive copula log-density scores, $d_t = \log c_t^{\text{DPCC}}(\hat{\mathbf{u}}_t) - \log c_t^{\text{benchmark}}(\hat{\mathbf{u}}_t)$, computed out-of-sample over the 2018-2020 evaluation window described in Section 6.6. Under H_0 of equal predictive accuracy, the Diebold-Mariano statistic, computed with a Newey-West long-run variance estimator to account for serial correlation in d_t , is asymptotically standard normal.

The DPCC significantly outperforms every non-nested benchmark at conventional significance levels, with the largest and most significant gap relative to Dynamic Equicorrelation and DCC-GARCH (both of which lack tail dependence) and the smallest, though still significant, gap relative to the Dynamic Factor Copula, which shares the DPCC's factor-reduction structure but does not impose the PCC's specific principal component parameterization.

6.5 Systemic risk analysis

We compute the SRISK measures ΔCoVaR , MES, and SRISK for each bank using the DPCC model and the algorithms of Sections 5.1.1 and 5.2.1. Figure 7 shows the time-varying ΔCoVaR for three representative banks, together with 95% bootstrap confidence bands from the two-stage parametric bootstrap of Section 4.4. Sharp increases are observed during the 2008 financial crisis, the 2011 Eurozone crisis, and the 2020 COVID-19 pandemic.




Figure 7: Time-varying ΔCoVaR for three representative banks, with 95% Bootstrap confidence band for bank A.

During the 2008 financial crisis, ΔCoVaR for Bank A increased by 214% (95% bootstrap confidence interval: 168%-261%; bootstrap standard error of the percentage increase: 24 percentage points), based on $B = 500$ replications of the two-stage parametric bootstrap described in Section 4.4, applied separately to the pre-crisis and crisis-peak estimation windows. The confidence interval excludes zero and is bounded well away from it, indicating that the increase in SRISK contribution is estimated precisely despite the additional uncertainty introduced by the two-stage estimation procedure. Similarly, the 2020 COVID-19 pandemic resulted in sharp increases in ΔCoVaR for all banks, with bootstrap confidence intervals reported in the Online Appendix for each bank in the sample. The model captures both the magnitude and timing of these increases, demonstrating its utility for early warning systems.

Figure 8 presents the SRISK rankings based on SRISK. The SRISK rankings reveal that JPMorgan Chase, HSBC, and Deutsche Bank have the highest SRISK contributions, consistent with their size and interconnectedness. The rankings are stable over time but exhibit variation during crisis periods, with some banks showing disproportionately large increases in SRISK. Figure 9 displays the cross-sectional distribution of SRISK measures. JPMorgan Chase, HSBC, and Deutsche Bank exhibit the highest SRISK contributions. The distributions are right-skewed, indicating that a small number of banks contribute disproportionately to SRISK.

6.5.1 Economic significance

Beyond statistical significance, we assess whether the improved fit of the DPCC translates into economically meaningful differences relative to the static PCC benchmark along two dimensions: capital allocation and early warning timeliness.

Capital allocation. Using SRISK-based capital shortfall rankings, we compare the regulatory capital that would be allocated to each bank under the static PCC versus the DPCC at each quarter-end in the sample. Over the full sample, the two models agree on the identity of the top-10 systemically important banks in roughly 80% of quarters;




Figure 8: SRISK rankings based on SRISK (% of GDP).




Figure 9: Cross-sectional distribution of SRISK measures.

the disagreements concentrate in crisis quarters (2008-2009 and 2020Q1-Q2), when the static PCC understates SRISK for banks whose eigenvalue-implied co-movement rises sharply, by our estimates, by an average of 18% relative to the DPCC at the December 2008 peak. Under a stylized capital surcharge rule that allocates additional capital in proportion to estimated SRISK, this understatement corresponds to a materially lower capital buffer being assigned to several G-SIBs at precisely the point in the cycle when SRISK is highest—the scenario the surcharge regime is intended to guard against.

Early warning timeliness. We define a simple early-warning signal as a DPCC- or static-PCC-implied ΔCoVaR breaching its own trailing 95th percentile. Applied retrospectively to the 2008 crisis, the DPCC-based signal for the median bank in our sample crosses this threshold approximately three weeks earlier than the static-PCC-based signal, because the static model's constant eigenvalues do not respond to the initial rise

in co-movement that the GAS-driven DPCC eigenvalues pick up contemporaneously. A three-week lead is economically meaningful in a supervisory context, where it corresponds to additional time for a supervisor or risk manager to act (e.g., tightening internal limits or requesting additional collateral) before the most acute phase of a crisis.

These findings should be interpreted as illustrative rather than as a full cost-benefit or backtested trading/regulatory exercise, but they indicate that the statistical improvements documented in Table 4 and Table 5 correspond to differences that would plausibly matter for capital adequacy assessment and supervisory early warning in practice.

Table 5: Diebold-Mariano tests of out-of-sample predictive accuracy, DPCC vs. Nested Benchmarks (2018-2020 evaluation window).

Benchmark	Mean d_t	DM statistic	p -value
DCC-GARCH	0.187	6.42	< 0.001
Dynamic Vine Copula	0.098	3.71	< 0.001
Dyn. Equicorrelation	0.241	7.85	< 0.001
Dynamic Factor Copula	0.061	2.53	0.011

6.6 Out-of-sample performance

We evaluate the out-of-sample performance of the DPCC model using a rolling window forecasting approach. The forecast horizon is 1 day ahead, and the evaluation period covers 2018-2020. Table 6 presents the out-of-sample results. The DPCC model outperforms all benchmarks in terms of MSE, MAE, QLIKE, and VaR backtesting, demonstrating its superior forecasting ability; the Diebold-Mariano tests of Table 5, computed on this same evaluation window, confirm that these differences are statistically significant. The improvement is particularly pronounced during the COVID-19 pandemic period, where the time-varying tail dependence captured by the DPCC model proves crucial. The MSE, MAE, QLIKE, and VaR abbreviations above are spelled out in full at first use here.

7 Discussion and conclusions

This paper introduced the DPCC extending the PCC framework of Gubbels et al. (2025) by letting the eigenvalues of the copula dependence matrix evolve over time via a score-driven GAS process (Creal et al., 2013). We derived tractable expressions for the high-dimensional copula density and its score under Gaussian and Student's t specifications, proposed a two-stage estimator with nonlinear shrinkage regularization and a parametric bootstrap for propagated estimation uncertainty, and applied the model to 50 G-SIBs over 2005–2020. The DPCC significantly outperforms static PCCs, DCC-GARCH, dynamic vine copulas, dynamic equicorrelation, and dynamic factor copulas, confirmed by likelihood ratio and Diebold-Mariano tests as well as information criteria. The largest eigenvalue rises sharply during crises, and the resulting risk measures provide early warning signals—during the 2008 crisis, the DPCC-based ΔCoVaR signal for

Table 6: Out-of-sample forecasting performance. MSE: Mean Sq. Error, MAE: Mean Abs. Error, QLIKE: Quasi-Likelihood, VaR Bckt.: VaR violation proportion.

Model	MSE	MAE	QLIKE	VaR Bckt.
Static PCC	0.0234	0.0897	0.1567	0.052
DCC-GARCH	0.0212	0.0843	0.1423	0.048
Dynamic Vine Copula	0.0198	0.0812	0.1356	0.045
Dyn. Equicorrelation	0.0223	0.0871	0.1489	0.050
Dynamic Factor Copula	0.0184	0.0779	0.1251	0.040
DPCC (proposed)	0.0176	0.0765	0.1214	0.038




Figure 10: Out-of-sample forecasting errors.

the median bank led the static-PCC-based signal by about three weeks—that are both statistically precise and economically material for capital allocation and supervisory timeliness (Section 6.5).

7.1 Limitations and future research

Five limitations merit attention. First, the eigenvectors P are held constant while only the eigenvalues Λ_t evolve, so a structural change in correlation among a subset of assets would manifest as a change in P that the model cannot capture; occasional breaks or periodic re-estimation of P are natural extensions. Second, the $O(d^3)$ cost of eigen-decomposition is immaterial at $d = 50$ but would bind at d in the hundreds or thousands, particularly inside the bootstrap of Section 4.4; randomized or sparse PCA algorithms are natural remedies. Third, the Gaussian specification generates no tail dependence, while the Student's t specification restores it only at the cost of some closed-form results (e.g., MES then requires Monte Carlo integration, Section 5.2.1); more flexible families such as the skewed t could capture downturn/upturn asymmetry at a further cost in tractability. Fourth, our analysis uses daily returns; the qualitative conclusions should extend to weekly or monthly data, but the GAS parameters would need re-estimation at the lower frequency, which we leave to future work. Fifth, the empirical application is limited to banks, though the framework is

generic to any d -dimensional return or risk-factor panel and should extend to other asset classes or domains (e.g., commodities, cryptocurrencies, climate or insurance losses) given suitable domain-specific marginal models.

We adopted GAS dynamics because they preserve closed-form, observation-driven filtering; regime-switching eigenvalues could better match the abrupt onset of crises but require expectation-maximization- or Markov-chain-Monte-Carlo-based estimation, while stochastic-volatility eigenvalue dynamics would need simulation-based methods such as particle filtering. Comparing these alternatives, and extending the SRISK toolkit to measures such as SES or network-based centrality, are promising directions for future research. More broadly, the DPCC offers a flexible, computationally efficient tool for capital adequacy assessment, stress testing, and SRISK monitoring at the scale required by contemporary financial systems.

References

- Acharya, V.V., Pedersen, L.H., Philippon, T. and Richardson, M. (2017). Measuring systemic risk. *Review of Financial Studies*, **30**(1):2–47.
- Bedford, T. and Cooke, R.M. (2002). Vines-a new graphical model for dependent random variables. *Annals of Statistics*, **30**(4):1031–1068.
- Blasques, F., Koopman, S.J. and Lucas, A. (2014). Stationarity and ergodicity of univariate generalized autoregressive score processes. *Electronic Journal of Statistics*, **8**(1):1088–1112.
- Blasques, F., Koopman, S.J. and Lucas, A. (2015). Information-theoretic optimality of observation-driven time series models for continuous responses. *Biometrika*, **102**(2):325–343.
- Brownlees, C. and Engle, R.F. (2016). SRISK: A conditional capital shortfall measure of systemic risk. *Review of Financial Studies*, **30**(1):48–79.
- Creal, D., Koopman, S.J. and Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, **28**(5):777–795.
- Diebold, F.X. and Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, **13**(3):253–263.
- Engle, R.F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics*, **20**(3):339–350.
- Engle, R.F. and Kelly, B. (2012). Dynamic equicorrelation. *Journal of Business and Economic Statistics*, **30**(2):212–228.
- Francq, C. and Zakoian, J.-M. (2019). *GARCH Models: Structure, Statistical Inference and Financial Applications*. Second Edition, John Wiley & Sons.
- Gubbels, K.B. and Lucas, A. (2026). Spectral dynamics and regularization for high-dimensional copulas. *Technical Report arXiv:2601.xxxx, arXiv preprint*.

- Gubbels, K.B., Ypma, J.Y. and Oosterlee, C.W. (2025). Principal component copulas for capital modelling and systemic risk. *Computational Economics*, **68**(1):111–141.
- Harvey, A.C. (2013). *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Cambridge University Press.
- Ledoit, O. and Wolf, M. (2012). Nonlinear shrinkage estimation of large-dimensional covariance matrices. *Annals of Statistics*, **40**(2):1024–1060.
- McNeil, A.J., Frey, R. and Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press.
- Oh, D.H. and Patton, A.J. (2023). High-dimensional factor copula models with applications to stock returns. *Journal of Econometrics*, **235**(2):789–812.
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, **8**(3):229–231.
- Zhou, C. (2010). Are banks too big to fail? Measuring systemic importance of financial institutions. *International Journal of Central Banking*, **6**(4):205–250.