

Structured factor copulas for modeling the systemic risk of European and United States banks

Hoang Nguyen^(a), Audronė Virbickaitė^(b), M. Concepción Ausín^(c),

and Pedro Galeano^(c)

^(a)Department of Management and Engineering, Linköping University, Sweden

^(b)Department of Quantitative Methods, CUNEF Universidad, Madrid, Spain

^(c)Department of Statistics and UC3M-Santander Big Data Institute, Universidad Carlos III de
Madrid, Spain

March 28, 2024

Abstract

In this paper, we employ Credit Default Swaps (CDS) to model the joint and conditional distress probabilities of banks in Europe and the U.S. using factor copulas. We propose multi-factor, structured factor, and factor-vine models where the banks in the sample are clustered according to their geographic location. We find that within each region, the co-dependence between banks is best described using both, systematic and idiosyncratic, financial contagion channels. However, if we consider the banking system as a whole, then the systematic contagion channel prevails, meaning that the distress probabilities are driven by a latent global factor and region-specific factors. In all cases, the co-dependence structure of bank CDS spreads is highly correlated in the tail. The out-of-sample forecasts of several measures of systemic risk allow us to identify the periods of distress in the banking sector over the recent years including the COVID-19 pandemic, the interest rate hikes in 2022, and the banking crisis in 2023.

Keywords: Bank risk; Contagion; Credit Default Swaps; Crisis; Default; Distress; Factor vine copulas.

JEL Classification: G01; G21; C01; C11; C58.

1 Introduction

Gaining insight into the intricate patterns of co-dependence among financial institutions is crucial. As highlighted by Coval et al. (2009); Zimmer (2012); Oh and Patton (2017), the 2008 financial crisis revealed the limitations of existing financial models, which failed to accurately capture the complex dependencies among assets. Their shortcomings stemmed from an overreliance on Gaussianity, simplistic correlation-based assumptions, and a neglect of the risks posed by extreme events (tail dependencies). Recently, the 2023 collapse of three U.S. regional banks - Silicon Valley, Signature, and First Republic - alongside the European Credit Suisse bank serves as a compelling reminder of the critical need to model the interdependence of financial institutions using complex non-linear, non-Gaussian models. This event, referred to by Akhtaruzzaman et al. (2023) as an instance of financial contagion, emphasizes the limitations of traditional models in predicting and mitigating the risks of such interconnected financial crises, further justifying the importance of advancing beyond simplistic correlation-based assumptions to more accurately capture the intricate dynamics of financial markets. Financial contagion refers to the cascading effects of some initial shock to a small portion of a financial system can have on the entire system (Dornbusch et al., 2000; Forbes and Rigobon, 2002; Constancio, 2012). Modeling and forecasting contagion is of major importance: for instance, given one bank has defaulted, what is the probability that the others will default? This issue has received a great deal of attention, as evidenced by Acemoglu et al. (2015); Dungey and Gajurel (2015); Martins (2023).

Modeling banking contagion is not a simple task, since it can arise from several sources, as widely documented by the finance-related literature (Ballester et al., 2016; Kleinow and Moreira, 2016; Agosto et al., 2020). Most of the articles differentiate between two channels of contagion: systematic and idiosyncratic, also characterized in Agosto et al. (2020) as global and local, respectively. The systematic or global channel is associated with certain global (i.e., inter-country) factors, as seen in the 2008 financial crisis, while the idiosyncratic or local channel is related to some bank-specific (i.e., inter-institutional) factors, as witnessed in 2023. Understanding how contagion works and the causes of systemic risk has been proven to be essential in the context of global financial markets. For example, Agosto et al. (2020) have found evidence of significant vulnerability across countries and institutions at the start of the 2008 global financial crisis and during the 2011 sovereign debt crisis. Similarly, Kleinow and Moreira (2016) identify critical factors driving systemic risk in the European banking sector, and demonstrate the importance of these factors for bank regulatory practices. Lastly, Ballester et al. (2016) emphasize the significance of common factors in the spread of risk spillovers, suggesting that global connections among banking markets play a key role in the

transmission of shocks. However, the question persists as to whether banking contagion results from shared underlying factors, indicating systematic contagion, or if it transmits from one bank to another, demonstrating an idiosyncratic nature. Understanding the source of contagion is crucial as it dictates divergent strategies for policymakers.

The first contribution of this article lies in identifying both systematic and idiosyncratic sources of financial contagion—global and local, respectively—through a novel statistical-based approach that captures complex co-dependence patterns (fat-tailed, asymmetric, non-Gaussian) among financial institutions. Figure 1 presents a simplified diagram of the two contagion transmission sources: the panel on the left illustrates the case of systematic contagion caused by some common factor F that affects all banks B_i in the system, meanwhile the panel on the right shows the case of idiosyncratic contagion, which is caused by a single bank B_1 . In practice, it is highly probable that both channels operate simultaneously, with the intensity of each source fluctuating over time. Additionally, the presence of multiple common factors, an ambiguous factor structure, and uncertainty regarding the quantity and nature of connections between the nodes are additional considerations. In other words, we are facing a complex co-dependence structure, that is high-dimensional, (potentially) non-linear, and non-Gaussian. Therefore, in this article we propose to model the complex co-dependence structure between financial institutions using the factor copulas of Krupskii and Joe (2013, 2015) and the factor-vine copulas of Brechmann and Joe (2014) and Nguyen (2019). In these models, the dependence structure is defined through latent variables using bivariate copulas. More precisely, we employ five model designs, starting with one- and two-factor copulas (Krupskii and Joe, 2013), that permit the units (banks) in the system to be affected by one or several latent factors. Then, we also consider alternative factor structures that do not suffer from the curse of dimensionality, namely bi-factor and nested-factor models (Krupskii and Joe, 2015). Nonetheless, the aforementioned setups solely accommodate unit-factor interaction, i.e. systematic contagion. To relax this assumption and enable inter-institutional interaction, we adopt the factor-vine copula model (Brechmann and Joe, 2014; Nguyen, 2019) that effectively captures both forms of financial contagion, systematic and idiosyncratic.

The second contribution of this article is the use of the Variational Bayes (VB) estimation method for the proposed factor copulas. Variational Inference is a computational technique used in Bayesian statistics and machine learning to estimate complex probability distributions that allows for fast inference and prediction even with very large datasets, see Blei et al. (2017) for a review. In particular, we follow Kucukelbir et al. (2017) to approximate the posterior distribution of the copula model parameters. A very relevant feature of our proposed inference procedure is that it can also facilitate an automated model selection




Figure 1: Systematic and idiosyncratic contagion.

The plot on the left shows the systematic contagion caused by a common factor F that affects all banks B_i (inter-country or global), while the plot on the right shows the idiosyncratic contagion caused by a single bank B_1 (inter-institutional or local).

process for determining the type of bivariate relationships between the variables, something that cannot be easily done using e.g. maximum likelihood, for which these relationships must be established in advance (Krupskii and Joe, 2013, 2015).

The third contribution of this article is an empirical application using the most recent CDS (Credit Default Swap) data of European and U.S. banks. A CDS is a financial contract in which a buyer of corporate or sovereign debt reduces possible loss arising from the default of the issuer of such debt. The buyer pays the CDS premium regularly over the life of the CDS contract in exchange for protection against the default of a specific borrower. The interest in modeling and forecasting CDS spreads stems from their efficacy as indicators of market perceptions on the likelihood of default for the reference entity, be it a company or a nation (Creal and Tsay, 2015; Ye et al., 2022). As noted in Oh and Patton (2018), the main factor responsible for CDS spread fluctuations is changes in the probability of default. In other words, changes in the CDS spreads signal changes in the credit risk, and the co-movements in the CDS spreads are directly attributable to financial contagion. Using CDS data is crucial as it accurately reflects market views on systemic risk, providing a clear, immediate measure of financial stability and contagion impacts, thus greatly enhancing the understanding of risk management. Indeed, CDS spreads have been widely employed to measure and model financial and sovereign contagion, see for example Beirne and Fratzscher (2013); Alter and Beyer (2014); Chen et al. (2020); Agosto et al. (2020); Ye et al. (2022). Furthermore, as demonstrated by Alter and Schöler (2012); Bruyckere et al. (2013); Alter and Beyer (2014), there exist risk spillovers between banks and sovereigns. Therefore, successful forecasting of default probabilities of banks can help regulators and policymakers to act accordingly in order to reduce the possible risk spillover toward sovereigns.

In the empirical illustration, we employ the daily CDS spread data for eight European and seven U.S. banks for a period of more than 16 years. The chosen period encompasses the 2008 global financial crisis, the 2009-2014 European sovereign debt crisis, the COVID-19 pandemic of 2020, the 2022 interest rate spikes, and the collapse of several financial institutions in

2023. To our knowledge, this is the longest (time-wise) CDS series considered in the default probability modeling context in the banking sector. Our findings indicate that the financial contagion between Europe and U.S. can be attributed to systematic channels only, where there is a single global factor affecting both markets and two regional factors affecting each market individually. Moreover, if we consider European and U.S. markets individually, the financial contagion sources change. Regional financial contagion is transmitted not only from some latent common factor (systematic) but also from one bank to another (idiosyncratic). By using the best-fitting bi-factor and factor-vine copula models we obtain 20-step-ahead forecasts for various measures of systemic risk. We show that our proposed models can accurately capture the distress periods actually observed in historical CDS spread series.

The rest of the paper is organized as follows. Section 2 contains a brief review of the related literature on the use of copula in modeling financial interconnectedness. Section 3 introduces the marginal model for bank default risk and the factor copulas used for modeling the joint distress probabilities. We present the Variational Bayesian inference algorithm in Section 4. Section 5 contains the application of the proposed factor copula for modeling the joint and conditional distress probabilities of 15 European and U.S. banks. Finally, conclusions are drawn in Section 6.

2 Related literature

Rodriguez (2007) was one of the first authors to study financial contagion between countries/regions using copulas. Using East Asian and Latin American stock indices he found evidence of changing dependence during the periods of turmoil and increased tail dependence in some markets but not the others. Aloui et al. (2011) and Kenourgios et al. (2011) study financial contagion using copulas by analyzing the co-dependence between the developed markets, such as the U.S. and U.K., and emerging markets, in particular, the BRICs. Arakelian and Dellaportas (2012) rely on copulas to capture instances of financial contagion by employing a change-point approach. The authors propose to test the existence of financial contagion by examining whether the number of breaks in copula parameters and individual marginal volatilities is different from zero. This approach provides evidence of financial contagion in equity markets and between sovereign entities during both the Asian and Mexican crises.

Other authors, instead of focusing on country/region contagion, consider interconnectedness at the firm-level. Brechmann et al. (2013) study contagion in financial markets by using vine copulas. More precisely, the authors find evidence of asymmetric tail behavior and that U.S. banks have a stronger impact on the international financial markets as com-

pared to Europe and Asia-Pacific using CDS spread data for 38 large international banks and insurers for the period of 2006-2011. Similarly, Kleinow and Moreira (2016) investigate the drivers of systemic risk and contagion among European banks. The authors employ copulas to quantify the systemic risk contribution and systemic risk sensitivity by using the CDS spread data of 36 European financial institutions from 2005 to 2013. Cerrato et al. (2017) investigate joint credit risk in the U.K. banking sector by analyzing the weekly CDS spreads from 2007 to 2015. The study reveals a strong connection between the time-varying and asymmetric dependence structure of CDS log-differences and the joint probability of default. The authors employ a Generalized Autoregressive Score-based dynamic asymmetric copula model. The results confirm that a significant part of the dependence structure of CDS spread log-differences is influenced by market factors.

An alternative method for analyzing the spread of systemic risk involves the use of network-based models. For instance, Glasserman and Young (2015) investigate how the interconnectedness of financial institutions can increase expected losses and defaults by assuming only minimal information about network structure. Paltalidis et al. (2015) explore the spread of financial contagion within the Eurozone by developing a dynamic financial network. The authors found that sovereign credit risk channel is the primary source of financial contagion in the banking network. Nonetheless, Torri et al. (2018) highlight a significant limitation of network-based models: the network’s structure often has to be inferred from data that is both noisy and aggregated, or by making simplifications that may not hold true, such as Gaussianity, for example.

Closest to our proposed approach, Oh and Patton (2018) employ a new class of one-factor dynamic copula to model the co-dependence of CDS spreads among 100 U.S. firms over the period 2006 to 2012. The authors find that while the probability of distress for individual firms has decreased since the height of the financial crisis, the current levels of joint probability of distress are significantly higher than those observed before the financial crisis. Krupskii and Joe (2020) propose a time-varying factor copula model to capture the non-Gaussian co-dependence between the variables of interest. The authors apply their proposed model to the CDS spread data of 24 U.S. firms and provide estimates of the joint probability of distress, similar to the approach in Lucas et al. (2014) and Oh and Patton (2018).

3 Model specification

This section describes in detail the models used to capture the interconnectedness of the default probabilities among financial institutions (banks). To this end, we first show how

the CDS spreads are related to the default probabilities. We then use existing models to estimate the marginal distributions for the log-differences of the CDS spread data. Finally, we describe in detail the factor copulas used for modeling the joint behavior of the default probabilities. We define three classes of copula models: the multi-factor copulas of Krupskii and Joe (2013), which include the one and two-factor copulas; the structured factor copulas of Krupskii and Joe (2015), which include the bi-factor and nested copulas; and the novel factor-vine copulas of Nguyen (2019), resulting in five distinct model architectures.

3.1 Default probability and marginals

To assess default risk, one standard approach involves calculating the implied default probability by valuing the present value of both the premium payment and the protection payment. Following the simplified model in Lange et al. (2016), let S be the CDS spread, usually measured in basis points (bps), at time $t = 0$ where the risk-free interest is r and the risk neutral loss given default is L . The implied risk neutral default probability p is the risk-neutral solution of premium payment and protection payment discounted for τ period of payment, i.e.

$$\frac{S}{1+r} + \frac{S}{(1+r)^2} + \dots + \frac{S}{(1+r)^\tau} = \frac{pL}{1+r} + \frac{(1-p)pL}{(1+r)^2} + \dots + \frac{(1-p)^{\tau-1}pL}{(1+r)^\tau} \quad (1)$$

or equivalently,

$$\frac{(1 - (1+r)^{-\tau})S}{r} = pL \frac{1 - (1-p)^\tau(1+r)^{-\tau}}{r+p}. \quad (2)$$

Here we assume that the default event only happens at the end of each payment period. The interest rate r as well as the loss given default L are constant. We choose the EURIBOR one year as the proxy of the risk-free interest rate and set $L = 75\%$ as in Lucas et al. (2014). Note that, when r , p and τ are small enough, we could approximate $S \approx pL$.

The CDS spread series usually exhibit very high levels of autocorrelations, sometimes being very close to the non-stationary region, therefore, we consider the log-differences of daily CDS spreads as $s_{i,t} = \log(S_{i,t}) - \log(S_{i,t-1})$ for each bank $i = 1, \dots, d$ and time unit $t = 1, \dots, T$. Additionally, we assume that CDS log-differences, $s_{i,t}$, follow a stationary, asymmetric autoregressive GJR-GARCH model with skewed and fat tail innovations. More precisely, we assume a AR(p)-GJR-GARCH(1,1) model with standardized skew- t errors given by:

$$\begin{aligned} s_{i,t} &= \mu_i + \sum_{j=1}^p \phi_{j,i} s_{i,t-j} + a_{i,t}, \\ a_{i,t} &= \sigma_{i,t} \eta_{i,t}, \quad \eta_{i,t} \sim \text{skew-}t(0, 1, \xi_i, \nu_i), \\ \sigma_{i,t}^2 &= \omega_i + (\alpha_i + \gamma_i J_{a_{i,t-1} < 0}) a_{i,t-1}^2 + \beta_i \sigma_{i,t-1}^2, \end{aligned} \quad (3)$$

where μ_i is a constant mean, $\phi_{j,i}$ are the autoregressive parameters, $(\omega_i, \alpha_i, \beta_i, \gamma_i)$ are the GJR-GARCH parameters and $I_{a_{i,t-1} < 0}$ indicates whether $a_{i,t-1}$ is negative, in which case it is 1, and when positive, in which case it is 0. To guarantee stationarity and positive variance, the root of the AR polynomial equation has to be less than one, $(\alpha_i + \gamma_i) > 0$ and $0 < (\alpha_i + \beta_i + \gamma_i/2) < 1$. As a consequence, $a_{i,t}$ is a sequence of innovations or shocks, and $\eta_{i,t}$ is a sequence of independent standardized innovations which is assumed to follow a standardized skew- t distribution with skewness ξ_i and degrees of freedom ν_i , (Fernández and Steel, 1998).

We note that any other alternative model can be specified for the marginals. Krupskii and Joe (2013, 2015) employ an AR(1)-GARCH(1,1)- t model. Creal and Tsay (2015) consider a stochastic volatility model with skew- t errors, Oh and Patton (2017) use an AR(1)-GJR-GARCH(1,1) model with a nonparametric error distribution, and Oh and Patton (2018) employ an AR(5)-GJR-GARCH(1,1) model with skew- t errors. These models can be different for each of the individual series. If the model is specified correctly, the probability integral transform results in uniformly distributed data, i.e., $F_{s_{i,t}}(s_{i,t}) = u_{i,t} \sim U(0, 1)$, where $F_{s_{i,t}}$ denotes the model distribution function for the CDS log-differences. Next, using the transformed variables $\mathcal{U} = \{u_{i,t} : i = 1, \dots, d; t = 1, \dots, T\}$, the co-dependence structure of the CDS spreads is modeled via different factor copula models.

3.2 Co-dependence modeling via copulas

The construction of flexible multivariate distributions using copulas started with the seminal work of Sklar (1959). Since then, many studies have focused on utilizing these tools for various issues. For more information, please refer to Joe (2014). A copula is a distribution on the multidimensional unit cube whose marginals are uniform. The use of copulas allows for the separation of the modeling of the individual marginals from that of the co-dependence structure as follows. Consider a collection of random variables Y_1, \dots, Y_d with the corresponding marginal cumulative distribution functions $F_{Y_i}(y_i) = P[Y_i \leq y_i]$ for $i = 1, \dots, d$ and a joint distribution function $H(y_1, \dots, y_d) = P[Y_1 \leq y_1, \dots, Y_d \leq y_d]$. According to Sklar (1959), there exists a copula C such that $H(y_1, \dots, y_d) = C(u_1, \dots, u_d)$, where $u_i = F_{Y_i}(y_i)$, for $i = 1, \dots, d$; the copula C is unique if the marginals are continuous. It is straightforward that u_i are realizations of a uniform random variable $U_i = F_i(Y_i) \sim U(0, 1)$, for $i = 1, \dots, d$. As noted in Oh and Patton (2017), the copula decomposition property into margins, $F_i(\cdot)$, and dependence structure, $C(\cdot)$, has two important advantages over considering the joint distribution, $H(\cdot)$, directly. Firstly, estimation can be performed independently, greatly reducing the computational burden. Secondly, the choice of the marginal models is independent across

series as well as from the co-dependence structure, as long as the models fit the data well. This two-stage approach is called inference functions for the margins (IFM) method, and, as shown by (Joe, 2005), under weak regularity conditions, the estimated parameters of the copula model are consistent and asymptotically efficient.

Unfortunately, finding adequate copula functions even in moderate dimensions may become complicated. One popular approach is the use of vine copulas where the dependence structure is defined as a product of bivariate copulas using tree representations, see Czado and Nagler (2022) for a review. However, the requirement of a significant number of parameters becomes an important problem when the dimension becomes high. Thus, truncated vine copulas and factor copulas are effective tools for capturing the dependence in high dimensional settings without requiring a large number of parameters. For example, Brechmann et al. (2013) assume that there is no dependence in higher tree level of vine copulas. Alternatively, Krupskii and Joe (2013, 2015) propose factor copula models where observable variables link to unobservable variables via bivariate copulas. On the other hand, Creal and Tsay (2015) and Oh and Patton (2017) extend the classical factor analysis by inverting the dependence structure from latent elliptical or skew-elliptical distributions to the constrained copula domain. In this paper, we concentrate on the former approach as it allows for a more flexible choice of copula functions.

3.2.1 Multi-factor copulas

Krupskii and Joe (2013) proposed a multi-factor copula model as truncated C-vines rooted at latent variables. Note that C-vines are vines whose trees have a main root node. They assume that the observables U_1, \dots, U_d are rooted at the latent variables V_0, \dots, V_{p-1} such that U_i , for $i = 1, \dots, p$, are independent given the latent factors V_j , for $j = 1, \dots, p$, which follow uniform distributions $V_j \sim U(0, 1)$. Therefore, in each tree level of the factor copula model, a common factor stays at the root of the C-vine, and observable copula variables are linked to the common latent variable through bivariate copula functions. Table 1 in the Online Appendix shows some common bivariate copula functions as the building block for the one-factor and two-factor copula architectures which are illustrated in Figure 2. For example, let us denote each bivariate copula density as $c_{U_i, V_0}(u_i, v_0)$, each bivariate copula distribution function as $C_{U_i, V_0}(u_i, v_0)$ and the conditional distribution of U_i given V_0 as $F_{U_i|V_0}(u_i|v_0)$ in the first tree level of the two-factor copula model. The copula data in the second level of the two-factor copula are obtained as pseudo observables $u_{i|v_0} = F_{U_i|V_0}(u_i|v_0) \in (0, 1)$, where $u_{i|v_0}$ are the realizations of a random variable $U_{i|V_0}$. The two-factor copula assumes that the variables in the second level $U_{i|V_0}$ are conditionally independent given the latent factor V_1 . Hence, it is trivial to calculate the conditional copula density $p(u_1, \dots, u_d|v_0, \dots, v_p)$ and the

integrated copula density $p(u_1, \dots, u_d)$, see details in Appendix A.




Figure 2: Multi-factor copulas.

The plot on the left illustrates the one-factor copula, meanwhile the plot on the right illustrates the two-factor copula for $d = 5$.

3.2.2 Structured factor copula

To account for a more complex dependence structure and address the curse of dimensionality problem, Krupskii and Joe (2015) propose an alternative approach by dividing d variables into G groups, so in total, there are $G + 1$ latent variables $V_0, V_1, \dots, V_G \sim U(0, 1)$. This gives rise to bi-factor and nested-factor copulas, see the illustration in Figure 3. The group division is arbitrary and is usually based on some external information, such as geographic location or firm sectors, for example. Alternatively, Oh and Patton (2023) adapt a k -means clustering algorithm for the group division.

In the bi-factor copula model, all i_g -th variables in group g are conditionally independent given the latent variables, V_g, V_0 , for $g = 1, \dots, G$, and U_{i_g} does not depend on g' , for any other $g' \neq g$, given V_0 . Note that the bi-factor copula model is an extension of the two-factor copula, which is a bi-factor copula with $G = 1$. The nested-factor copula model assumes that G latent variables, V_1, \dots, V_G , are generated from a common latent variable V_0 and all observables U_{i_g} in group g are independent given the latent group factor, V_g , for $g = 1, \dots, G$. Krupskii and Joe (2015) demonstrate the flexibility of the structured factor copula over the multi-factor copula models and derive the explicit form of copula density, see details in Appendix A.

3.2.3 Factor-vine copula

Finally, Nguyen (2019) extends the factor copula model proposed by Krupskii and Joe (2013) to include a factor copula at the first tree level and truncated vine copulas at higher tree levels. The new proposed model aims to capture different behaviors at the tail of the joint distribution and provides interpretable economic meaning. This idea is similar to Brechmann and Joe (2014) who construct Gaussian copulas using one-factor and Gaussian truncated




Figure 3: Structured factor copulas.

The plot on the left illustrates the bi-factor copula, and the plot on the right illustrates the nested-factor copula for $d = 5$ and $G = 2$.

vine, however, their proposal ignores the issues of fat tails and asymmetric dependence among variables.

Following the notations in Bedford and Cooke (2002), Kurowicka and Cooke (2006), and Brechmann et al. (2012), we specify the truncated factor-vine copula as a sequence of trees, and the nodes in each tree are linked by bivariate copula functions. Let T_0, T_1, \dots, T_K be the trees in the truncated factor-vine copula model where $K \leq d - 1$, and let N_i and E_i be the set of nodes and edges in tree T_i , for $i = 0, \dots, K$, such that the following requirements are satisfied:

1. Tree T_0 has nodes $N_0 = \{0, 1, \dots, d\}$ and edges $E_0 = \{\{0, 1\}, \{0, 2\}, \dots, \{0, d\}\}$.
2. For $i = 1, \dots, K$, the nodes in tree T_i are the edges in tree T_{i-1} , that is, $N_i = E_{i-1}$.
3. If two edges in tree T_i are joined by an edge in tree T_{i+1} , they must share a common node.

The truncated factor-vine copula is a specification of truncated regular vine copulas (R-vine). Note that R-vines are vines only verifying requirements 2 and 3. Thus, truncated factor-vines are truncated R-vines where at tree T_0 , there is a common root latent variable, V_0 , that links to all observable variables by bivariate copula functions. Let the node set be $N = \{N_0, \dots, N_K\}$ and let the edge set be $E = \{E_0, \dots, E_K\}$. One associates each edge $e(j, k)$ in E_i with a bivariate copula distribution function, $C_{j,k|D(e)}$, and a bivariate copula density, $c_{j,k|D(e)}$, where j and k are conditioned nodes and $D(e)$ is the conditioning set. The bivariate copulas in tree T_0 have an empty conditioning set.

Figure 4 shows an example of a truncated factor-vine copula with two levels. Level 1 (on the left) contains a one-factor copula, and the remaining co-dependence among the variables, conditional on V_0 , is captured by a regular vine copula on level 2 (on the right). Appendix A derives the conditional factor-vine copula density.




Figure 4: Truncated factor-vine copula.

The figure shows a truncated factor-vine copula with two levels for $d = 5$. Level 1 on the left contains a one-factor copula while the remaining dependence of variables, conditional on v_0 , is a regular vine copula on level 2 (on the right).

4 Bayesian inference

Krupskii and Joe (2013, 2015) use maximum likelihood estimation for structured factor copula models. However, the parametric families of the bivariate copula links are chosen at the beginning which limits the flexibility of the model. In this section, we start with a factor copula with a known structure and connection links. We apply Variational Bayesian inference, similar to Nguyen et al. (2020), to estimate the copula parameters. Then, we discuss an extension of an automated process to identify the dependence structure and to select the appropriate bivariate copula functions.

4.1 Prior and posterior distributions

Assuming that we have a pre-specified factor copula structure together with bivariate linking copulas, the goal is to make inference on the latent variables $\mathbf{v}_{0:G,1:T}$ and copula parameters $\boldsymbol{\theta}$ given the copula data $\mathcal{U} = \{u_{i,t} : i = 1, \dots, d; t = 1, \dots, T\}$. For example, in the truncated two-level factor-vine copula model, let $\Theta = \{\mathbf{v}_{0,1:T}, \boldsymbol{\theta}\}$ be the set of parameters of interest, where $\mathbf{v}_{0,1:T} = \{v_{0,1}, \dots, v_{0,t}\}$ and $\boldsymbol{\theta} = \{\boldsymbol{\theta}_{01}, \dots, \boldsymbol{\theta}_{0d} \text{ and } \boldsymbol{\theta}_{jk}, \text{ for } e(j,k) \in E_1\}$. We assume a uniform distribution $U(0, 1)$ for the latent variables, $v_{0,t}$, and an uninformative but proper prior distribution for any set of bivariate copula parameters in $\boldsymbol{\theta}$ so that the Kendall- τ is in the range of $[-0.9; 0.9]$. See the details of prior distributions for copula parameters in Table 1 of the Online Appendix.

The prior density is a product of independent prior distributions of the model parameters, $\pi(\Theta) = \prod_{t=1}^T \pi(v_{0,t}) \prod_{i=1}^d \pi(\boldsymbol{\theta}_{0i}) \prod_{e\{j,k\} \in E_1} \pi(\boldsymbol{\theta}_{jk})$. Using the conditional copula density in Appendix A, the posterior distribution up to a normalizing constant of the truncated two-level

factor-vine copula can be written as,

$$\begin{aligned}
\pi(\Theta|\mathcal{U}) &\propto \prod_{t=1}^T p(u_{1,t}, \dots, u_{d,t}|v_{0,t}, \Theta)\pi(\Theta) \\
&\propto \prod_{t=1}^T \left[\prod_{e\{j,k\} \in E_1} c_{U_j, U_k|V_0}(u_{j,t}|v_{0,t}, u_{k,t}|v_{0,t}; \theta_{jk}) \prod_{i=1}^d c_{U_i, V_0}(u_{i,t}, v_{0,t}; \theta_{0i}) \right] \times \quad (4) \\
&\prod_{t=1}^T \pi(v_{0,t}) \prod_{i=1}^d \pi(\theta_{0i}) \prod_{e\{j,k\} \in E_1} \pi(\theta_{jk}).
\end{aligned}$$

The truncated factor-vine copula requires obtaining a new set of pseudo data, $u_{i,t|v_{0,t}} = F_{U_{i,t}|V_0}(u_{i,t}|v_{0,t})$, whenever we generate a new sample of any element in Θ which makes the implementation of the Markov Chain Monte Carlo (MCMC) computationally expensive. Hence, in the next section, we employ an approximate approach to obtain samples from the posterior distribution, namely, Variational Bayes.

4.2 Variational inference

Variational inference is a computational technique used in Bayesian statistics and machine learning for approximating complex probability distributions. For the factor copula models, described above, we are interested in making inferences on the factor copula parameters Θ based on the posterior distribution $\pi(\Theta|\mathcal{U})$. Variational Bayes proposes a family of distributions \mathcal{Q} and finds a proposal distribution $q(\cdot; \lambda) \in \mathcal{Q}$, parameterized by a vector λ , to minimize the Kullback-Leibler divergence between $q(\Theta; \lambda)$ and $\pi(\Theta|\mathcal{U})$. However, as $\pi(\Theta|\mathcal{U})$ is defined up to a normalizing constant, VB instead maximizes the evidence lower bound (ELBO):

$$\begin{aligned}
\text{ELBO}(q) &= \mathbb{E}_q[\log p(\mathcal{U}, \Theta)] - \mathbb{E}_q[\log q(\Theta; \lambda)] \\
&= \log p(\mathcal{U}) - KL(q(\Theta; \lambda) || \pi(\Theta|\mathcal{U})) \leq \log p(\mathcal{U})
\end{aligned}$$

such that when $q(\Theta; \lambda) = \pi(\Theta|\mathcal{U})$, we obtain $\text{ELBO} = \log p(\mathcal{U})$. Therefore, VB turns Bayesian inference into an optimization problem of finding λ to maximize the ELBO.

Blei et al. (2017) review different approaches in choosing the variational distributions and optimization techniques. As the posterior distribution of the structured factor copula is complex, we follow Kucukelbir et al. (2017) and use a black box variational inference based on stochastic approximation of ELBO. The general framework is to sample from the distribution $q(\Theta; \lambda)$ and use these values to approximate the expectations in ELBO using Monte Carlo integration. In order to find the optimal ELBO, we also sample from $q(\Theta; \lambda)$ and take the average of the derivatives of ELBO with respect to the parameters and update

the variational parameters according to the gradients. For example, in the truncated factor-vine copula presented in the previous section, the mean-field approximation assumes that \mathcal{Q} is a Gaussian distribution and

$$q(\Theta; \lambda) = \prod_{t=1}^T \pi(v_{0,t}; \lambda) \prod_{i=1}^d \pi(\theta_{0i}; \lambda) \prod_{e\{j,k\} \in E_1} \pi(\theta_{jk}; \lambda), \quad (5)$$

where λ contains the posterior mean and posterior variance of the Gaussian distribution. We illustrate the schematic representation of the algorithm in Appendix B. Nguyen (2019) presents multiple simulation studies regarding the trade-off between the computational speed and estimation accuracy using VB as compared to MCMC for truncated factor-vine copula models. The author shows that VB approximates quite well the posterior distribution of the latent variable and slightly underestimates the standard deviations of the bivariate copula parameters.

4.3 Model selection

Given a pre-specified factor copula structure, it is fast to estimate the model parameters using VB. However, the parametric families of the bivariate copula links are usually unknown due to the nature of the latent variables. For the structured factor copula model, Krupskii and Joe (2015) use graphical tools and Fan and Joe (2023) use a proxy estimation to define the bivariate copula links between variables. On the other hand, Nguyen et al. (2020) derive an automatic procedure to recover the hidden dependence structure. This procedure iterates between the estimation of the latent common factors and the copula model selection between the latent factors and the observables. For example, we first assign all Gaussian bivariate copulas to the structured factor copula and estimate the common latent factors using VB. Given the posterior medians of the latent factors, we consider other copula models and select the best family for the bivariate links between the observables and the latent factors using the Bayesian information criterion (BIC). We reassign the chosen copula families for the links and reestimate the model until the bivariate links remain unchanged or until a certain number of iterations is reached. Nguyen et al. (2020) show that this procedure increases the BIC of the factor copula model. Also, the accuracy rate of the correct links is above 70% in high dimensional simulations with very similar log-likelihoods as compared to true simulated models. For the truncated factor-vine copula models, we can select the bivariate copulas jointly among levels or level-by-level, which is similar to Dissmann et al. (2013). We use a one-factor copula model to choose the bivariate links in the first level and apply the algorithm in Dissmann et al. (2013) for higher levels given the pre-specified copulas in the

first level. Nguyen (2019) shows that the accuracy rates are also high in all levels of the factor-vine copulas.

5 Empirical illustration

5.1 Data

In this section, we employ the proposed factor copulas to analyze the joint distress probability of European and U.S. banks. We select only the banks that have actively traded in CDS during the designated period between January 1, 2007 and May 1, 2023 (16 years and 4 months). Finally, we extract daily data for *5-year maturity CDS of senior unsecured debt* for 8 European and 7 U.S. banks from Bloomberg. The ticker symbols and full names are summarized in Table 1. The selection of specific U.S. and European banks for this analysis was primarily influenced by the availability and depth of CDS data within the Bloomberg database. The limitation in the span of CDS data for some banks, or the lack of such data for others, imposed a focus on those institutions that offered a sufficient historical record to ensure the reliability of our analysis. Furthermore, the selected banks have high liquidity (as evidenced by individual CDS spread plots in the Online Appendix), which is important for correctly predicting systemic risk and the spread of financial contagion. The proposed approach can be easily scaled up to higher dimensions, depending on data availability. Nguyen (2019) and Nguyen et al. (2020) provide simulation studies containing 100 series, and show that VB estimation and automatic model selection procedure can be executed quickly and accurately.

Table 1: Bank data.

European banks		U.S. banks	
Ticker symbol	Full name	Ticker symbol	Full name
BCS	Barclays	AXP	American Express Company
BNP	BNP Paribas	BAC	Bank of America Corporation
CS	Credit Suisse	C	Citigroup
GLE	Societe Generale	GS	The Goldman Sachs Group
ISP	Intesa Sanpaolo	JPM	JPMorgan Chase
SAN	Santander	MS	Morgan Stanley
UBS	UBS	WFC	Wells Fargo
UCG	UniCredit		

Notes: The table reports main information for the European and U.S. banks: ticker symbol, full name and the main operating market.

Before proceeding, we impute some missing values for some CDS spread series, see details in Online Appendix. After the treatment, we have a sample of $T = 4110$ observations. Figure 5 draws the dynamics of the CDS spreads, measured in basis points, and the implied

marginal default probabilities, obtained using Equation 2. For both regions, various significant increases in CDS spreads indicate critical financial stress periods within the banking sector. The first increase is during the global financial crisis, where the U.S. financial sector suffered a greater impact as compared to the European banking sector. The second visible increase is during the turbulent developments in the European sovereign debt market during 2009-2014 (Kleinow and Moreira, 2016), which affected both, Europe and U.S. banking sectors. As noted in the Financial Times, “Credit default swaps on major U.S. banks (...) hit fresh highs for the year on Tuesday as problems in Greece intensified.”¹. There is another notable increase at the year 2020 mark, which corresponds to the COVID-19 pandemic outbreak. Finally, the last peak in the CDS spreads is at the very end of the sample, which reflects the effect of the 2022 interest rate spikes on the banking sector and the early 2023 banking panic. The European region has a more pronounced increase because Credit Suisse bank is included in the sample, meanwhile, the CDS spread data for the three bankrupt U.S. banks is not available and, therefore, not reflected in the plot. Major descriptive statistics for the crude CDS spread data are reported in Online Appendix. The minimum values of CDS spreads vary between 5 and 20 bps, meanwhile, the maximum values go as high as 1360 bps for Morgan Stanley during the global financial crisis and 1082 bps for Credit Suisse in 2023. The bottom row of Figure 5 draws the corresponding implied default probabilities for both regions. We can observe that the implied probabilities closely resemble the CDS spread dynamics. For both regions, the highest average implied default probabilities were reached during the worldwide financial crisis, and they have not reached those levels ever since (at least for the banks in the sample).

Next, the crude CDS spread data is converted into the log-differences $s_{i,t}$, as described in Section 3. Table 2 reports the summary statistics for the log-differences (in %) of the CDS spreads of all 15 banks, and Figure 6 plots the evolution of the log-differences. Most of the variables exhibit positive skewness and excess kurtosis, similar to the results in Oh and Patton (2018). Note that in some cases the log-differences are as low as negative 100%. This does not pose any problems in the analysis, as the log transformation is used for mathematical convenience only. In the modeling and forecasting of systemic risk in Section 5.5 we work with the CDS spread data and not the log-differences.

Finally, Figure 7 plots the correlations between the log-differences. We can clearly see two groups: European banks are in the bottom left corner, meanwhile U.S. banks are in the right top corner. American Express Company (AXP) has the lowest correlation with any of the banks. The overall correlations are always positive and range from medium to very high, indicating a very high level of interconnectedness among the banks; more regional

¹<https://www.ft.com/content/515508a2-9edc-11e1-9cc8-00144feabdc0>




Figure 5: CDS spreads and implied default probabilities by region.

The figures show the evolution of the CDS spreads measured in basis points (top row) and the evolution of the implied default probabilities (bottom row) for January 2007 to May 2023. The European region contains 8 banks and the U.S. has 7 banks in total. The dynamics of each individual bank are in gray, and the average is in black.

Table 2: Descriptive statistics for the log differences (in %) of CDS spreads.

	mean	Q2	sd	skew	kurt	min	max	1stACF	ADF(10)	LB(10)	ARCH(10)
BCS	0.07	-0.05	4.26	-0.22	17.23	-47.30	35.74	0.16	0.01	0.00	0.00
BNP	0.06	-0.04	4.53	0.17	11.03	-31.51	34.86	0.12	0.01	0.00	0.00
CS	0.07	-0.02	4.45	-2.35	85.99	-101.62	57.53	0.18	0.01	0.00	0.00
GLE	0.06	-0.04	4.45	0.40	13.13	-39.20	40.96	0.12	0.01	0.00	0.00
ISP	0.07	-0.03	5.01	1.26	31.68	-46.07	80.52	0.06	0.01	0.00	0.00
SAN	0.05	-0.05	4.38	0.07	12.09	-42.76	35.56	0.18	0.01	0.00	0.00
UBS	0.07	-0.04	4.15	0.18	15.79	-47.25	35.19	0.12	0.01	0.00	0.00
UCG	0.06	-0.06	4.45	0.37	14.58	-40.36	41.97	0.18	0.01	0.00	0.00
AXP	0.04	-0.03	3.83	0.47	13.39	-31.41	30.27	0.00	0.01	0.00	0.00
BAC	0.07	-0.07	4.67	0.65	44.14	-60.07	68.96	0.06	0.01	0.00	0.00
C	0.06	-0.09	4.33	-0.16	29.03	-52.11	40.12	0.12	0.01	0.00	0.00
GS	0.04	-0.09	4.31	0.39	43.43	-57.55	64.15	0.06	0.01	0.00	0.00
JPM	0.04	-0.07	3.94	0.02	18.77	-49.07	35.98	0.12	0.01	0.00	0.00
MS	0.03	-0.10	4.82	-3.33	135.57	-118.53	79.94	0.01	0.01	0.00	0.00
WFC	0.07	-0.05	4.03	1.15	25.13	-41.46	54.58	0.15	0.01	0.00	0.00

Notes: The table reports the descriptive statistics for the log differences (in %) of CDS spreads. 1stACF is the first-order autocorrelation, $ADF(m)$ is the p -value for the Augmented Dickey-Fuller test with m lags, the $LB(k)$ is the p -value for the Ljung-Box test for autocorrelation of k lags, and $ARCH(l)$ is the p -value for the test for ARCH effects with l lags. The data is from January 1st, 2007 to May 1st, 2023.




Figure 6: CDS log-differences by region.

The figures show the evolution of the CDS spread log-differences (in %) from January 1st, 2007 to May 1st, 2023. European region contains 8 banks and the U.S. has 7 banks in total. The dynamics of each individual bank are in gray, and the average is in black.

than global. European banks seem to show a stronger overall correlation than U.S. banks, indicating a more highly interconnected system.




Figure 7: Correlation matrix of CDS log-differences.

The figure shows the sample Pearson correlation matrix between the log-differences of the CDS spreads. European banks are in the lower left quadrant, and U.S. banks are in the upper right quadrant.

5.2 Marginals

Next, we model the log-differences of CDS spreads $s_{i,t}$ using an AR(4)-GJR-GARCH(1,1) with the skew- t errors, as described in Section 3.1. All details related to the marginal modeling, such as the parameter estimation results, residual descriptive statistics, and plots are in the Online Appendix. Most of the banks have estimated positive AR(1) and AR(2) and negative AR(3) and AR(4) coefficients, similarly as in Oh and Patton (2017). The persistence parameter of the GJR-GARCH(1,1) model indicates strong volatility persistence

for all assets, and the asymmetric volatility response parameter γ is negative in almost all cases, except for AXP, where it is statistically insignificant. In contrast to the “conventional” financial assets, where positive log returns typically signal positive developments, positive values of CDS spread log-differences mean an increase in CDS spreads, which conveys adverse developments or negative news. Therefore, the estimated leverage parameter γ has a negative sign: past negative log-differences are good news that decrease the current level of volatility. Finally, the estimated skewness and degrees of freedom parameters indicate that the error distribution is asymmetric and fat-tailed for all banks. The specified model is appropriate for the data, as the model residuals exhibit no autocorrelation or heteroscedasticity, and the probability integral transforms of the log-differences result in uniformly distributed data, i.e. $F_{s_{i,t}}(s_{i,t}) = u_{i,t} \stackrel{\text{iid}}{\sim} U(0, 1)$, as confirmed by the Kolmogorov–Smirnov, Anderson-Darling, and data-driven smooth tests for uniformity. All details are available in the Online Appendix.

5.3 In-sample results

Next, we model the uniformly distributed data $\mathcal{U} = \{u_{i,t} : i = 1, \dots, d; t = 1, \dots, T\}$ using the one-factor, two-factor, nested-factor, bi-factor, and two-level truncated factor-vine copula models proposed in Section 3. Apart from the proposed specifications, we also consider three benchmark models: a two-factor copula with only Gaussian links (Two-factor G), a Constant Conditional Correlation model with a full correlation matrix (Full G (CCC), Bollerslev, 1990), and a truncated-vine copula (Trunc Vine, Brechmann et al., 2012; Czado and Nagler, 2022). The two-factor Gaussian model does not accommodate asymmetric or fat-tailed co-dependence. The CCC model, while fully Gaussian and based on correlations, does not incorporate a factor structure. Finally, the truncated vine copula model allows only for unit-to-unit interaction (without common factors), which can also be non-Gaussian. The *ad hoc* grouping of the variables, necessary in the nested-factor and bi-factor copulas is based on the geographic location of the banks (Europe and U.S.). We have fitted the proposed models on the full sample as well as four approximately equal-length subsamples to account for a possible change of co-dependence structure over time. Table 3 reports the in-sample Bayesian Information Criteria, while Online Appendix contains the results for the log likelihoods. Similarly, Oh and Patton (2018) have also employed log likelihoods, AIC, and BIC criteria for the in-sample model evaluation.

The first four columns of Table 3 contain the BICs for the four subperiods, meanwhile, the last column reports the overall BIC for the full sample. For both regions individually, the factor-vine copula is decidedly the best, indicating that both contagion channels are at work: systematic and idiosyncratic. Nonetheless, there is one exception. For the U.S.

Table 3: In-sample Bayesian Information Criteria.

	Subsample 1 Jan 2007 - Dec 2010	Subsample 2 Jan 2011 - Dec 2014	Subsample 3 Jan 2015 - Dec 2019	Subsample 4 Jan 2020 - May 2023	Full sample Jan 2007 - May 2023
U.S. banks					
Two-factor G	-3768.02	-5515.66	-6442.52	-5674.15	-20266.85
Full G(CCC)	-3809.97	-5530.34	-6503.93	-5640.27	-20260.08
Trunc Vine	-4178.36	-5529.05	-6854.12	-6061.82	-21793.50
One-factor	-4200.99	-5830.27	-7006.16	-6170.19	-22622.88
Two-factor	-4182.95	-5803.62	-6979.97	-6098.57	-22610.04
Factor-Vine	-4455.18	-5834.02	-7017.22	-6161.97	-22698.24
European banks					
Two-factor G	-7248.53	-8956.40	-8413.19	-6473.14	-30475.41
Full G(CCC)	-7342.71	-9070.69	-8477.18	-6507.00	-30939.64
Trunc Vine	-7348.53	-8960.56	-8608.54	-6972.59	-31828.56
One-factor	-7359.31	-8951.08	-7829.39	-6725.13	-30878.72
Two-factor	-7407.53	-8884.31	-7873.32	-6980.22	-32215.13
Factor-Vine	-7446.67	-9125.10	-8646.63	-7096.86	-32670.12
U.S. and European banks					
Two-factor G	-11135.07	-14671.44	-14361.85	-12402.50	-51236.18
Full G(CCC)	-11237.97	-14747.83	-15071.31	-12427.77	-52604.30
Trunc Vine	-11803.37	-14851.49	-15779.86	-13519.86	-55033.73
One-factor	-9175.17	-11433.56	-10301.33	-10392.73	-41131.21
Two-factor	-11744.15	-14924.44	-14964.80	-13293.54	-54773.96
Bi-factor	-12016.55	-15322.08	-15809.10	-13429.70	-56212.63
Nested-factor	-11761.48	-15202.06	-15191.11	-13472.43	-54336.74
Factor-Vine	-11563.51	-14732.43	-15498.63	-13346.46	-53731.25

Notes: The table reports the in-sample Bayesian Information Criteria (BIC) for four subperiods and the full sample. We compare our proposed factor copula model with benchmarks: a two-factor copula with only Gaussian links (Two-factor G), a Constant Conditional Correlation model with a full correlation matrix (Full G (CCC), Bollerslev, 1990), and a truncated-vine copula (Trunc Vine, Brechmann et al., 2012; Czado and Nagler, 2022).

in the subsample 4 (January 2020 - May 2023), a period that includes COVID-19, 2022 interest rate spikes, early 2023 banking panic, the best is the one-factor copula. From the economic perspective, this means that during this specific period, the interconnectedness of the U.S. banking sector can be explained due to some common latent factor rather than the interactions between banks. Finally, in considering both regions together, the structure that describes the co-dependence between U.S. and European banks the best is the bi-factor model (with the exception in subsample 4). Such specification permits for a common latent global factor, and two regional latent factors to govern the co-dependence in the banking sector. The above results provide some interesting insights: on a regional level, banking sector contagion has two channels, systematic and idiosyncratic (inter-institutional). However, on the global level, the banking sector is mostly affected by global and regional common factors, and not by individual banks, i.e. contagion is mostly due to systematic factors.

For all models, the vast majority of factor-unit links are t copulas, indicating fat-tailed but symmetric dependence between the latent factor(s) and individual banks. In other words, the systematic contagion channel is best described using fat-tailed specifications. This is also confirmed by the overwhelming underperformance of both Gaussian distribution-based models, the two-factor Gaussian and the CCC model. The degrees of freedom parameter can vary from 2.2 to 41.1 for different models, different banks, and different subsamples. For

the factor-vine copula, even though the factor-unit links are mostly t , the unit-unit links at the second tree level are more heterogeneous, and depend on the units (banks) as well as the time period. As an example, Figure 8 draws the second tree level for the factor-vine copula for the last subperiod and the full sample. For the last subperiod (top row), in the U.S., the selected C-vine structure assigns three links to Wells-Fargo bank with t , Frank, and Gumbel copulas. In Europe, the Italian UniCredit (UCG) bank has as many as five links, with many different types of links. Bottom row presents a similar plot, but for the entire sample. For the U.S., the Bank of America (BAC) now has the most links, all Frank copulas (symmetric and fat-tailed). In Europe, the UniCredit bank remains with the highest number of links. Interestingly, the link structure is more stable for Europe than for the U.S. across both subsamples. In other words, for the U.S, most of the co-dependence is captured in the first level via the latent factor, and the remaining co-dependence is weak and more difficult to identify, resulting in different vine copula structures for each subperiod. On the other hand, in the European region, the latent factor at the first tree level captures only a part of the co-dependence, and the remaining links at the second level are relatively strong and retain similar structures across different subsamples.




Figure 8: Second level of the truncated factor-vine copula.

The figure shows the detailed second tree level of factor-vine copula for the U.S. (on the left) and European (on the right) banks. Top row is for the last subsample: from January 2020 to May 2023, bottom row is for the entire period: from January 2007 to May 2023.

5.4 Out-of-sample results

For the out-of-sample forecasting exercise, we retain the last 1000 observations, from February 28th 2019 till May 1st, 2023. We then roll the dataset in 20-day increments (approximately one month), re-estimate all the models, and predict the co-dependence structure for the next 20-day period. In order to evaluate the out-of-sample forecasting performance of individual models we calculate the negative log predictive score (LPS), the negative conditional likelihood score (CdL), and the variogram score (VarS). The LPS takes into account the goodness of fit of the entire predictive distribution and is one of the most commonly used measures to evaluate the predictive model performance. The CdL, on the other hand, allows to evaluate a certain region of the predictive distribution, for example, the upper tail. The upper tail is defined as a d -dimensional quadrant where the CDS log-differences for all the banks fall within the top 50th percentile. This typically covers approximately 10% of the data, depending on the time period. Note that contrary to the “conventional” assets, the extreme negative events of the CDS spread log-differences happen in the upper tail, and not in the lower tail. Finally, the VarS is based on pairwise differences of the components of the multivariate forecast. This way we can evaluate point-wise forecast accuracy. For details concerning the calculation of each of the scores and their properties, refer to Bjerregård et al. (2021). The results for LPS, CdL, and VarS for the proposed models are reported in Table 4.

Table 4: Out-of-sample forecast results.

	One-factor	Two-factor	Bi-factor	Nested-factor	Factor-Vine
(a) Negative log predictive scores					
US banks	-2874.50	-2884.82	-	-	-2892.41
Euro banks	-3165.57	-3254.41	-	-	-3307.69
US and Euro banks	-4932.64	-6196.65	-6365.17	-6313.78	-6127.93
(b) Negative conditional likelihood predicted scores					
US banks	-1532.82	-1537.80	-	-	-1533.43
Euro banks	-1755.38	-1799.65	-	-	-1788.37
US and Euro banks	-1910.98	-2036.85	-2032.17	-1977.84	-2037.89
(c) Variogram score ($p = 0.5$)					
US banks	811.76	811.84	-	-	809.90
Euro banks	1020.67	1017.73	-	-	1016.08
US and Euro banks	4324.11	4323.60	4264.03	4430.45	4272.60

Notes: The table reports the sum of the negative log predictive scores (LPS), negative conditional likelihood scores (CdL), and variogram scores (VarS) for the last 1000 observations (from February 28th, 2019 to May 1st, 2023, both inclusive). The scores are calculated by rolling the sample every 20 days, re-estimating the models, and predicting the co-dependence structure for the next 20 days. Lower numbers mean better. The parameter $p = 0.5$ next to the VarS indicates the order of the score, see the discussion in Bjerregård et al. (2021).

In terms of LPS and VarS, the out-of-sample forecasting results are consistent across the two scoring rules. Similar to the in-sample results, the banking sector co-dependence is best

described via factor-vine copula if we consider both regions individually. In other words, the financial contagion has systematic and idiosyncratic sources in regional financial systems. On the other hand, if we consider both banking systems jointly, then the bi-factor copula provides the best out-of-sample forecasting results. This means that both financial systems are governed by some common global factor and two individual regional factors. In this case, the bank-to-bank (idiosyncratic) contagion is not as important. Interestingly, the CdL metric provides alternative model ordering: if we consider the upper tail of the distribution only, then for both regions the two-factor copula is the best, meanwhile for both regions jointly the factor-vine copula provides the smallest value of the CdL. In other words, to account for the behavior in the upper tail, a multi-source systematic contagion channel is prevalent within each region individually. Furthermore, when considering the upper tail of the European and U.S. banking systems collectively, both global and local contagion channels are in operation.

Even though both contagion channels are in operation, we further investigate which contagion channel is dominant, if any. Following Diebold and Yilmaz (2014), we decompose the K -step ahead variance forecast error of the reduced-form vector autoregression (VAR) to assess the bank-wise and overall system connectedness. For illustration, we select March 16th, 2023, around the recent collapse of several U.S. banks, and compare connectedness in two models: the one-factor copula, allowing for the systematic contagion channel only, and the factor-vine copula, permitting both systematic and idiosyncratic channels.

We find that the bank-wise and global connectedness measures are relatively similar in both models, suggesting the dominance of the systematic contagion channel. The results are consistent with the findings in the related literature. For example, Ballester et al. (2016) has studied financial contagion between European and U.S. banks using the CDS spread data. By employing a generalized VAR approach, they show that the systematic contagion component is always greater than the idiosyncratic component, emphasizing the importance of common factors in the risk spillover transmission. These insights have important implications for policymakers and practitioners in the financial sector: systematic shocks to the banking system have a much greater effect than the idiosyncratic ones, and should be handled differently.

All details regarding connectedness estimation and the related tables are available in Appendix C at the end of the manuscript.

5.5 Measuring systemic risk

Modeling and forecasting CDS series allows to model and forecast instances of financial distress, which are defined as events where the implied probability of default exceeds a certain

threshold. Instead of relying solely on the implied probability of default, we focus on the directly observable CDS series to identify periods of distress. A similar approach was employed by Oh and Patton (2018), for example. The CDS prediction is carried out by simulating 10,000 samples of the d -variate variable u_t and applying the inverse probability integral transform (the quantile function) resulting from the forecasted AR(p)-GJR-GARCH(1,1)-skew- t model. This procedure is repeated for each forecast instance t and for each asset i .

This enables us to calculate four measures of systemic risk, namely the individual Probability of Distress (PD) (Oh and Patton, 2018), the Joint Probability of Distress (JPD) (Segoviano and Goodhart, 2009; Lucas et al., 2014; Oh and Patton, 2018; Krupskii and Joe, 2020), the Expected Proportion in Distress (EPD) (Hartmann et al., 2005; Oh and Patton, 2018) and the Expected Shortfall (ES) (Oh and Patton, 2017).

Distress happens when the observed CDS spread is larger than some threshold. The distress indicator for bank i , for $i = 1, \dots, d$ can be defined as $D_{i,t+20} = \mathbf{1}\{S_{i,t+20} > C_{i,t+20}\}$. Hence, the distress indicator D will be equal to one if the 20-step-ahead predicted CDS spread $S_{i,t+20}$ is larger than some threshold C . The threshold can be calculated as the upper 95 (or 99) percentile of the historical CDS spreads. Since in the forecasting exercise we have a rolling sample, the threshold C is time-varying. Therefore, C is recalculated as the upper 95 percentile of the last 1000 observations of the CDS spreads. Then the systemic risk measures are calculated as follows:

$$\begin{aligned}
 PD_{i,t} &= \Pr_t [D_{i,t+20} = 1], \\
 JPD_{k,t} &= \Pr_t \left[\left(\frac{1}{d} \sum_{i=1}^d D_{i,t+20} \right) \geq \frac{k}{d} \right], \\
 EPD_{i,t} &= \mathbb{E}_t \left[\frac{1}{d} \sum_{j=1}^d D_{j,t+20} \mid D_{i,t+20} = 1 \right], \\
 ES_{i,t} &= \mathbb{E}_t \left[S_{i,t+20} \mid \left(\sum_{j=1}^d \mathbf{1}\{S_{j,t+20} > C_{j,t+20}\} \right) > 1 \right].
 \end{aligned}$$

First, PD measures the probability of each individual bank being in distress in the next 20 days. Since at each instance t for each asset i we simulate 10,000 scenarios for the spreads $S_{i,t+20}$, this probability is estimated as a sample proportion of the spreads that reach beyond the threshold. Second, the JPD measures the probability that a certain proportion of banks will be in distress. This probability should be decreasing as k increases since it is less and less likely for more and more banks to be in distress at the same time. Third, the EPD presents an expected proportion of banks in distress when bank i has suffered distress. It might provide some insights into individual banks' relative importance in the system and its

effect on the rest of the banks. Finally, the ES metric calculates the expected CDS spread of bank i given that 2 or more banks have CDS spreads above some threshold C .

Figure 9 summarizes the results for all banks for the out-of-sample period. Please note that the presented plots offer a general overview of the trends in systemic risk measures. For precise values of a specific risk measure for a particular bank on a given day, refer to the detailed tables in the Online Appendix. Also, the presented results are for the entire system as a whole, i.e. using the bi-factor copula model. All detailed tables for the region-wise (Europe and the U.S.) systemic risk measures are available in the Online Appendix as well.

First, from sub-figures (a1) and (a2) we can observe an overall increase in the individual distress probabilities (PD) in March - April 2020. The increase was universal for all banks in the system, to as high as 0.83 for WFC. This dramatic increase in the distress probability can be attributed to the COVID-19 pandemic. Stepankova and Teply (2023) have documented the increase in the probability of default in the banking sector during the COVID-19 pandemic, and Duan et al. (2021) have found that the pandemic has increased systemic risk in the banking sector across as many as 64 countries. Another universal peak can be seen in the year 2022 and the beginning of 2023, with a sharp drop in August 2022. This persistent high probability of distress throughout the year can be attributed to the hike in the interest rates: only in 2022 the Fed raised interest rates five times and ECB - three times. As noted in a report by the Financial Stability Oversight Council (2013), a fast spike in yields can affect financial stability. For example, if a bank has a large portion of its assets locked in long-term bonds, the market value of those bonds goes down when interest rates start going up. If at the same time, a large number of depositors withdraw their funds due to concerns about the bank's solvency or liquidity (a phenomenon called a bank run), this might lead to a liquidity crunch and, in an extreme case, a bank's failure (Lyócsa et al., 2023). This indeed played an important part in the collapse of three regional U.S. banks (the Silicon Valley, Signature, and First Republic banks) in early 2023 (Lyócsa et al., 2023), which also triggered the downfall of the already vulnerable Credit Suisse. These events are visible in the plot as the peak of the probability of distress at the end of the sample.

Next, sub-figure (b) draws the Joint Probability of Distress (JPD) series that represents the probability that at a given moment in time, at least d banks will be in distress, for $d = 1, \dots, 15$. Similarly to the previous plot, there are two major turbulent periods, namely the COVID-19 pandemic and the interest rate hikes in 2022. During those periods, the probability that at least one bank will be in distress reaches almost one. Same as in the individual distress probabilities series, there is a sharp drop in the JPD in August 2022, which has no apparent explanation. Finally, the 2022 - March 2023 peak is also visible at the end of the period.

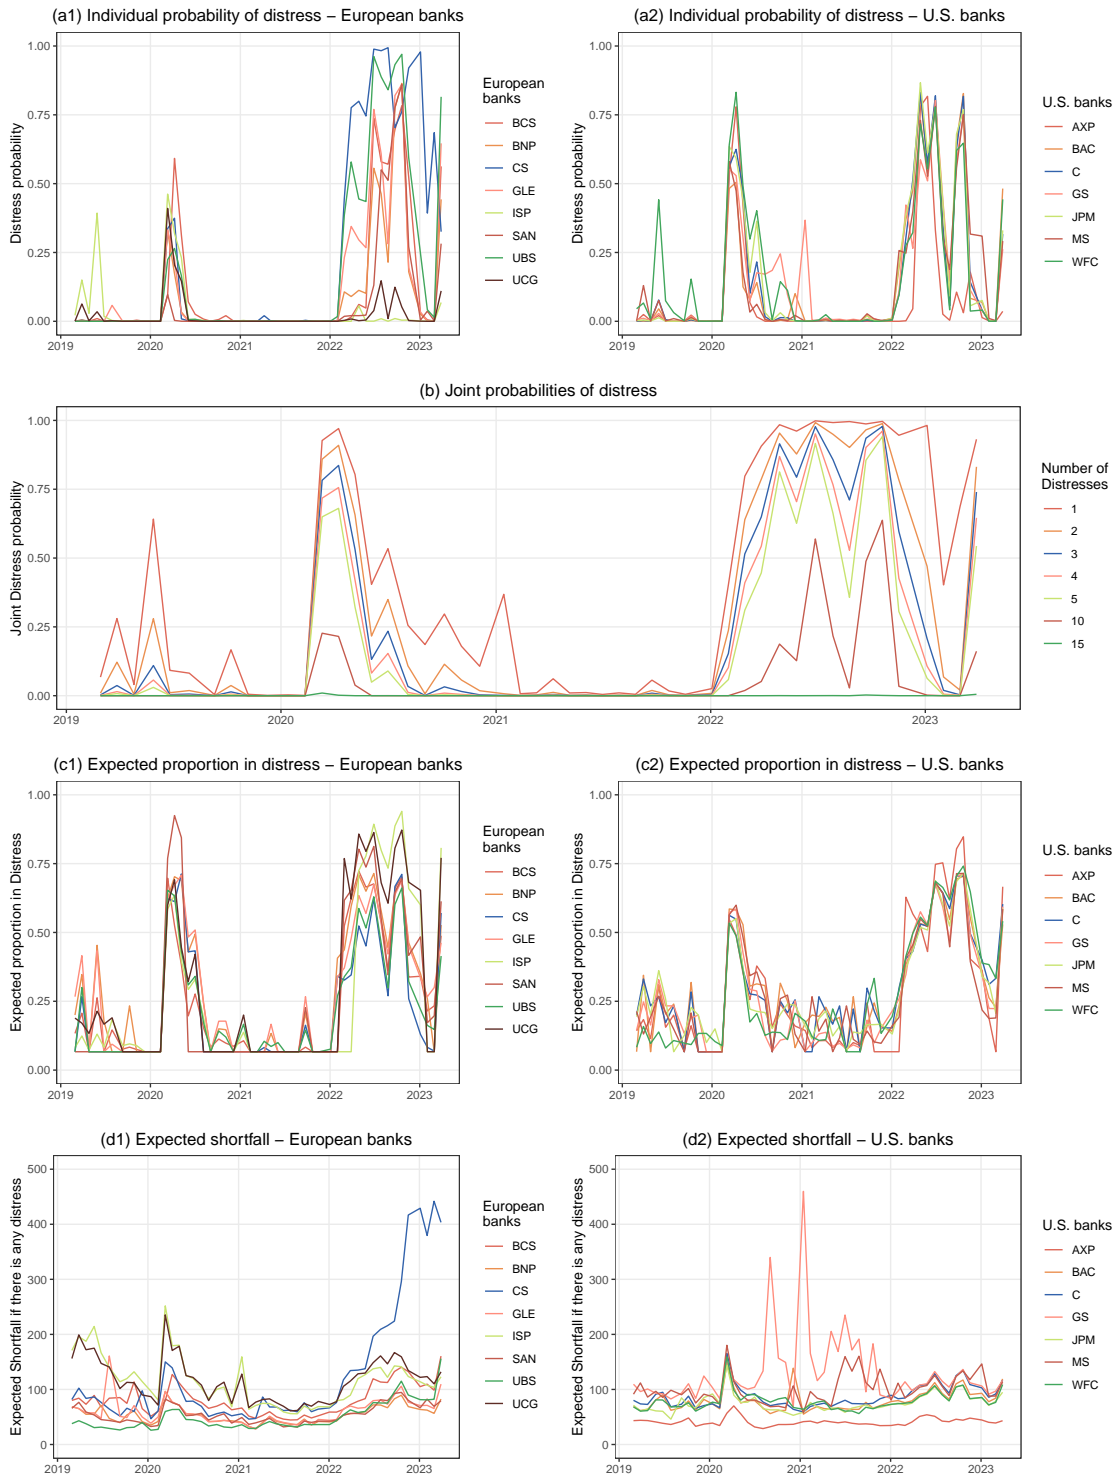


Figure 9: Forecasts of systemic risk.

The figure shows the forecast measures of systemic risk - PD , JPD , EPD , and ES - in the next 20 days for the out-of-sample period from February 28th, 2019 to March 29th, 2023.

Sub-figures (c1) and (c2) present the Expected Proportion in Distress (EPD), which permits to evaluate the individual banks’ relative importance in the system. For example, on February 28th, 2019 the *EPD* value for Barclay’s bank (BCS) is 0.13, which means that if in 20-days-ahead Barclay’s bank suffers distress, then 13% of all banks in the system will also suffer distress. Overall, it seems that U.S. banks’ relative importance in the overall system is of slightly smaller magnitude as compared to the European banks, at least for the selected sample. Interestingly, during the beginning of the COVID-19 pandemic, the Intesa Sanpaolo bank’s *EPD* skyrocketed from 0.07 to as high as 0.89. Italy was the European country where COVID-19 struck the strongest: it had the highest death toll and it was the first one to go into lockdown. Therefore, if Italy’s largest bank, Intesa Sanpaolo, were to suffer distress during that period, that would have sent a very strong signal to the rest of the countries about the possible adverse effects of the COVID-19 pandemic on the banking sector. The figure also gives a good approximation of the periods for “heightened sensitivity”. For example, we can observe a universal increase in the *EPD* values during periods of financial distress: the COVID-19 pandemic, the 2022 interest rate hikes, and the 2023 collapse of some European and U.S. banks.

Finally, sub-figures (d1) and (d2) summarize the results for the Expected Shortfall (ES) of CDS spreads when 2 or more banks are in distress. For example, during the COVID-19 pandemic (the very beginning of 2021), Goldman Sachs bank had the predicted *ES* value of 459 bps, given that two or more other banks in the system were to suffer distress. This indicates a high vulnerability of the Goldman Sachs bank during this specific period. In particular, for that specific time, Goldman Sachs’s probability of distress is estimated to be 0.37, which can be attributed to a sudden 30% daily increase in the CDS spreads (from 52 to 71 bps). Finally, at the end of the sample, Credit Suisse exhibits expected CDS spreads as high as 400, which closely resembles the actual observed CDS data for this specific bank.

6 Conclusion

In this paper, we employ five different models using multi-factor, structured factor, and factor-vine copulas to forecast both the joint and conditional probabilities of distress for a sample of banks. We employ the Variational Bayes estimation method to approximate the posterior distributions of both latent factors and copula model parameters. We also incorporate an automatic procedure to select the best bivariate links among the nodes. We find that in-sample, for U.S. and Europe individually, the best fitting model is the factor-vine copula. In other words, there are two financial contagion channels at work: systematic (governed by some latent factor at the first level) and idiosyncratic (governed by vine copulas

at the second level). Most of the factor-unit links at the first level are t copulas indicating fat-tailed and symmetric dependence structure, meanwhile, the unit-unit links at the second level are of all types: asymmetric and/or fat-tailed, depending on the banks and the time period. Moreover, if we consider U.S. and European banking systems jointly, the interconnectedness between the CDS spread log-differences is best captured using the bi-factor copula. This shows the existence of one common global factor and two regional factors, pointing to the dominance of the systematic contagion channel. Using the best-fitting copula models we then forecast a variety of risk measures that capture the individual, joint, and conditional probabilities of distress. The findings closely coincide with the observed historical data: there are clear peaks in probabilities of distress during turbulent periods, such as the COVID-19 pandemic, 2022 interest rate hikes, and the 2023 banking panic. Also, the model captures the increase in the probability of distress for certain banks in the system, such as Credit Suisse for example.

Acknowledgments

The second author is partially supported by the grants PID2022-138289NB-I00 (*Modeling and forecasting financial returns and macroeconomic variables in data-rich environments: new methods and applications*) and PID2020-113192GB-I00/AEI/10.13039/501100011033 (*Mathematical Visualization: Foundations, Algorithms and Applications*) from the Spanish State Research Agency (Ministerio de Ciencia e Innovación). The third and fourth authors are supported by the grant PID2022-138114NB-I00/AEI/10.13039/501100011033 (*High dimensional dependence modeling*) from the Spanish State Research Agency (Ministerio de Ciencia e Innovación).

Declaration of interest: None

Appendix

A Factor copulas and factor-vine copulas

A.1 One-factor copula models

In the one-factor copula model, the variables U_1, \dots, U_d are conditionally independent given the latent variable, V_0 , as shown in Krupskii and Joe (2013). We have that the conditional distribution function of U_i given V_0 is given by,

$$F_{U_i|V_0}(u_i|v_0; \boldsymbol{\theta}_{0i}) = \frac{\partial F_{U_i, V_0}(u_i, v_0; \boldsymbol{\theta}_{0i})}{\partial v_0} = \frac{\partial C_{U_i, V_0}(u_i, v_0; \boldsymbol{\theta}_{0i})}{\partial v_0},$$

where F_{U_i, V_0} is the joint CDF of U_i and V_0 , and $C_{U_i, V_0}(u_i, v_0; \boldsymbol{\theta}_{0i})$ is the bivariate copula distribution function parameterized by $\boldsymbol{\theta}_{0i}$ as the vector of parameters of the bivariate copula. Then, the conditional copula density and the integrated copula density can be derived respectively as,

$$\begin{aligned} p(u_1, \dots, u_d|v_0; \boldsymbol{\theta}) &= \frac{\partial F(u_1, \dots, u_d|v_0; \boldsymbol{\theta})}{\partial u_1 \dots \partial u_d} = \prod_{i=1}^d \frac{\partial F_{U_i|V_0}(u_i|v_0; \boldsymbol{\theta})}{\partial u_i} = \prod_{i=1}^d \frac{\partial C_{U_i, V_0}(u_i, v_0; \boldsymbol{\theta}_{0i})}{\partial u_i \partial v_0} \\ &= \prod_{i=1}^d c_{U_i, V_0}(u_i, v_0; \boldsymbol{\theta}_{0i}), \end{aligned} \tag{6}$$

and,

$$p(u_1, \dots, u_d; \boldsymbol{\theta}) = \int_0^1 p(u_1, \dots, u_d|v_0; \boldsymbol{\theta}) dv_0 = \int_0^1 \prod_{i=1}^d c_{U_i, V_0}(u_i, v_0; \boldsymbol{\theta}_{0i}) dv_0, \tag{7}$$

where $\boldsymbol{\theta} = [\boldsymbol{\theta}'_{01}, \dots, \boldsymbol{\theta}'_{0d}]'$ is the vector of copula parameters.

A.2 Two-factor copula models

Similarly, Krupskii and Joe (2013) compute the conditional copula density and the integrated copula density of the two-factor copula model respectively as,

$$\begin{aligned} p(u_1, \dots, u_d | v_0, v_1; \boldsymbol{\theta}) &= \prod_{i=1}^d p(u_i | v_0, v_1; \boldsymbol{\theta}) \\ &= \prod_{i=1}^d c_{U_i, V_1|V_0}(u_i|v_0, v_1; \boldsymbol{\theta}_{1i}) c_{U_i, V_0}(u_i, v_0; \boldsymbol{\theta}_{0i}), \end{aligned} \tag{8}$$

and

$$p(u_1, \dots, u_d; \boldsymbol{\theta}) = \int_0^1 \int_0^1 p(u_1, \dots, u_d | v_0, v_1; \boldsymbol{\theta}) p(v_1) p(v_0) dv_1 dv_0, \quad (9)$$

where $u_i|v_0 = F_{U_i|V_0}(u_i|v_0)$, and $\boldsymbol{\theta} = [\boldsymbol{\theta}'_{01}, \dots, \boldsymbol{\theta}'_{0d}, \boldsymbol{\theta}'_{11}, \dots, \boldsymbol{\theta}'_{1d}]'$ is the set of copula parameters. The joint copula distribution and density require a high dimensional integral which makes it difficult to consider more than three latent variables, see Joe (2014).

A.3 Structured factor copula models

The nested factor copula model defines a hierarchical dependence structure from a common root variable. The conditional density function for the nested factor copulas is derived following Krupskii and Joe (2015),

$$p(u_1, \dots, u_d, v_1, \dots, v_G | v_0; \boldsymbol{\theta}) = p(u_1, \dots, u_d | v_1, \dots, v_G; \boldsymbol{\theta}) p(v_1, \dots, v_G | v_0; \boldsymbol{\theta}), \quad (10)$$

where, similar to the one-factor copula, the conditional copula density for the observables is,

$$p(u_1, \dots, u_d | v_1, \dots, v_G; \boldsymbol{\theta}) = \prod_{g=1}^G \prod_{i_g=1}^{d_g} c_{U_{i_g}, V_g}(u_{i_g}, v_g; \boldsymbol{\theta}_{g i_g}),$$

and for the latent factor is,

$$p(v_1, \dots, v_G | v_0; \boldsymbol{\theta}) = \prod_{g=1}^G c_{V_g, V_0}(v_g, v_0; \boldsymbol{\theta}_{0g}).$$

where $c_{U_{i_g}, V_g}(u_{i_g}, v_g; \boldsymbol{\theta}_{g i_g})$ is the bivariate copula density of U_{i_g} and V_g for each group member i_g in group g with d_g members, for $g = 1, \dots, G$; and $c_{V_g, V_0}(v_g, v_0; \boldsymbol{\theta}_{0g})$ is the bivariate copula density of V_g and V_0 ; and $\boldsymbol{\theta} = [\boldsymbol{\theta}'_{01}, \dots, \boldsymbol{\theta}'_{0G}, \boldsymbol{\theta}'_{11}, \dots, \boldsymbol{\theta}'_{G d_G}]'$ is the set of copula parameters.

Similar to the two-factor copula, the conditional density function for the bi-factor copula model is given by,

$$\begin{aligned} p(u_1, \dots, u_d | v_0, \dots, v_G; \boldsymbol{\theta}) &= \prod_{g=1}^G \prod_{i_g=1}^{d_g} p(u_{i_g} | v_0, v_g; \boldsymbol{\theta}) \\ &= \prod_{g=1}^G \prod_{i_g=1}^{d_g} c_{U_{i_g}, V_g | V_0}(u_{i_g} | v_0, v_g; \boldsymbol{\theta}_{g i_g}) c_{U_{i_g}, V_0}(u_{i_g}, v_0; \boldsymbol{\theta}_{0 i_g}), \end{aligned} \quad (11)$$

where $u_{i_g} | v_0 = F_{U_{i_g} | V_0}(u_{i_g} | v_0)$ and $\boldsymbol{\theta} = [\boldsymbol{\theta}'_{01}, \dots, \boldsymbol{\theta}'_{0d}, \boldsymbol{\theta}'_{11}, \dots, \boldsymbol{\theta}'_{G d_G}]'$ is the vector of copula parameters.

A.4 Factor-vine copula models

We follow Kurowicka and Cooke (2006), Kurowicka and Cooke (2006), Brechmann et al. (2012), Krupskii and Joe (2013) to derive the joint density of the truncated factor vine copula model,

$$p(v_0, u_1, \dots, u_d; \boldsymbol{\theta}) = \prod_{i=0}^K \prod_{e(j,k) \in E_i} c_{j,k|D(e)}(u_{j|D(e)}, u_{k|D(e)}; \boldsymbol{\theta}) \quad (12)$$

where $u_{j|D(e)} = F_{U_j|U_{D(e)}}(u_j|u_{D(e)})$, $u_{k|D(e)} = F_{U_k|U_{D(e)}}(u_k|u_{D(e)})$ and $u_{D(e)}$ is a sub-vector of u with the conditioning set $D(e)$. The pseudo $u_{j|D(e)}$ can be computed recursively using the formula derived in Aas et al. (2009). In order to calculate the copula density of observable variables, one can take the integral of the conditional copula density over the latent space,

$$p(u_1, \dots, u_d; \boldsymbol{\theta}) = \int_0^1 \prod_{i=0}^K \prod_{e \in E_i} c_{j,k|D(e)}(u_{j|D(e)}, u_{k|D(e)}; \boldsymbol{\theta}) dv_0. \quad (13)$$

For example, the conditional density of a truncated factor-vine copula with two levels can be derived as

$$\begin{aligned} p(u_1, \dots, u_d | v_0; \boldsymbol{\theta}) &= \prod_{e\{j,k\} \in E_1} p(u_{j|v_0}, u_{k|v_0}; \boldsymbol{\theta}_{jk}) \prod_{i=1}^d p(u_i | v_0; \boldsymbol{\theta}_{0i}) \\ &= \prod_{e\{j,k\} \in E_1} c_{U_j, U_k | V_0}(u_{j|v_0, t}, u_{k|v_0, t}; \boldsymbol{\theta}_{jk}) \prod_{i=1}^d c_{U_i, V_0}(u_i, v_0; \boldsymbol{\theta}_{0i}), \end{aligned}$$

where E_1 is the set of edges $e(j, k)$ that connects variable j and variable k at level 2; and $\boldsymbol{\theta} = [\boldsymbol{\theta}'_{01}, \dots, \boldsymbol{\theta}'_{0d}, \boldsymbol{\theta}'_{jk}, \forall e(j, k) \in E_1]'$ is the vector of copula parameters.

B Variational Bayes: a schematic representation

Data: Copula data $U = \{u_i\}$ and structured copula model

Result: The bivariate copula types, and its free parameter λ

Initialization $\lambda^{(0)}$;

while *Any change in copula types do*

while *Change in ELBO is above some threshold do*

 Draw M samples from the variational distribution $q(\Theta; \lambda)$;

 Approximate the noisy gradient $\nabla_{\lambda} \text{ELBO}$;

 Update $\lambda^{(i+1)} \leftarrow \lambda^{(j)} + \rho^{(j)} \nabla_{\lambda} \text{ELBO}$ where $\rho^{(j)}$ is the adaptive step size;

 Incremental iteration (i) ;

end

Select best bivariate copula u_i and v based on BIC ;

 Reassign the copulas and estimate ;

end

Return copula structure and the parameters of proposal distribution;

Algorithm 1: Modified of ADVI in (Kucukelbir et al., 2017)

C Connectedness of European and U.S. banks

The co-movements of the CDS spreads can be written as a system of the reduced-form vector autoregression (VAR):

$$\begin{aligned} \mathbf{s}_t &= \Phi_1 \mathbf{s}_{t-1} + \dots + \Phi_4 \mathbf{s}_{t-4} + \mathbf{D}_t^{1/2} \boldsymbol{\eta}_t, \\ \mathbf{s}_t &= \Theta(L) \mathbf{a}_t, \text{ where } \mathbf{a}_t = \mathbf{D}_t^{1/2} \boldsymbol{\eta}_t \\ \Theta(L) &= (\mathbf{I} - \Phi_1 L - \dots - \Phi_4 L^4)^{-1} = \Theta_0 + \Theta_1 L + \Theta_2 L^2 + \dots, \end{aligned} \quad (14)$$

where $\mathbf{s}_t = (s_{1,t}, \dots, s_{d,t})'$ are the CDS log-differences, $\Phi_1, \dots, \Phi_4, \mathbf{D}_t^{1/2}$ are diagonal matrices that stack the parameters of the AR(4)-GJR-GARCH(1,1) marginals from Eq.3, and Θ_h is the coefficient matrix multiplying the h -lagged shock vector in the infinite MA representation. Diebold and Yilmaz (2014) consider the connectedness measures as the shares $\mathbf{H}^t = [\hat{h}_{ij}^t]$ of variance decomposition of K -step-ahead errors in forecasting \mathbf{s}_i due to the shocks to \mathbf{s}_j ,

$$\hat{h}_{ij}^t = \frac{h_{ij}^t}{\sum_{j=1}^d h_{ij}^t} \quad \text{with} \quad h_{ij}^t = \frac{\sigma_{jj,t}^{-1} \sum_{k=0}^{K-1} \left(\mathbf{e}_i' \Theta_k \Sigma_t \mathbf{e}_j \right)^2}{\sum_{k=0}^{K-1} \left(\mathbf{e}_i' \Theta_k \Sigma_t \Theta_k' \mathbf{e}_i \right)}, \quad (15)$$

where \mathbf{e}_j is a d -dimensional vector with the j th element a unity and zeros elsewhere, $\boldsymbol{\Sigma}_t = \text{cov}(\mathbf{a}_t)$ and σ_{jj} is the j -th diagonal element of $\boldsymbol{\Sigma}_t$.

Using the factor copula models, it is straightforward to simulate the forecast errors and estimate the connectedness at any given point in time. For illustration purposes, we estimate bank-wise connectedness at the point near the collapse of several U.S. banks in 2023. Tables 5 presents estimation results for the one-factor copula (systematic contagion channel only), and Table 6 presents estimation results for the factor-vine copula (systematic and idiosyncratic contagion channels). The 15 x 15 matrix in the upper left part of the both tables show the percentage of the changes in CDS spread of bank in row i due to the shocks in bank in column j . For example, according to the one-factor model (Table 5), 8% of the changes in the CDS spreads of Barclays bank (BCS) are due to the changes in the CDS spreads of Credit Suisse (CS) bank (first row, third column).

Table 5: Connectedness of European and U.S. banks using the one-factor copula model.

	BCS	BNP	CS	GLE	ISP	SAN	UBS	UCG	AXP	BAC	C	GS	JPM	MS	WFC	From
BCS	12.7	7.7	8.0	7.5	4.6	7.4	6.2	6.2	1.9	7.2	6.9	6.6	7.0	2.7	7.3	87.3
BNP	7.2	12.0	8.2	7.7	4.7	7.4	6.3	6.3	1.9	7.3	7.0	6.7	7.0	2.8	7.4	88.0
CS	7.0	7.7	14.2	7.6	4.6	7.2	6.2	6.1	1.9	7.1	6.9	6.6	6.8	2.8	7.3	85.8
GLE	7.2	7.8	8.2	11.6	4.7	7.4	6.4	6.3	1.9	7.3	7.0	6.8	7.0	2.8	7.5	88.4
ISP	7.3	7.9	8.3	7.8	7.6	7.5	6.4	6.4	2.0	7.4	7.1	6.7	7.1	2.8	7.5	92.4
SAN	7.2	7.8	8.1	7.7	4.7	12.2	6.3	6.3	1.9	7.2	7.0	6.6	7.0	2.8	7.4	87.8
UBS	7.0	7.6	8.0	7.6	4.6	7.2	12.6	6.2	1.9	7.2	6.8	6.5	6.8	2.8	7.3	87.4
UCG	7.3	7.9	8.2	7.8	4.8	7.5	6.4	9.6	1.9	7.4	7.1	6.8	7.1	2.8	7.5	90.4
AXP	7.1	7.7	8.0	7.5	4.6	7.3	6.3	6.2	7.8	7.2	6.8	6.5	6.9	2.8	7.3	92.2
BAC	7.2	7.8	8.0	7.7	4.7	7.3	6.3	6.3	1.9	12.0	7.0	6.7	7.0	2.8	7.4	88.0
C	7.2	7.8	8.2	7.7	4.7	7.4	6.3	6.3	1.9	7.3	10.9	6.8	7.1	2.8	7.5	89.1
GS	6.9	7.5	7.9	7.4	4.5	7.1	6.0	6.0	1.8	7.0	6.8	14.6	6.7	2.7	7.1	85.4
JPM	7.2	7.7	8.0	7.6	4.7	7.4	6.2	6.3	1.9	7.3	7.0	6.6	12.1	2.8	7.3	87.9
MS	7.2	7.8	8.2	7.8	4.7	7.4	6.4	6.3	1.9	7.4	7.0	6.8	7.1	6.6	7.4	93.4
WFC	7.2	7.7	8.1	7.7	4.7	7.4	6.3	6.2	1.9	7.2	7.0	6.7	6.9	2.8	12.2	87.8
To	100.3	108.3	113.3	107.1	65.3	102.9	88.0	87.5	26.7	101.5	97.5	93.3	97.5	38.9	103.1	88.8
Net	13.0	20.3	27.5	18.6	-27.0	15.1	0.6	-2.9	-65.5	13.5	8.4	8.0	9.6	-54.5	15.2	

Notes: Connectedness of European and U.S. banks using the one-factor copula model on March 16th, 2023. The 15 x 15 elements in the upper left show the shares $\mathbf{H}^f = [\hat{h}_{ij}^f]$ of variance decomposition of 5-step-ahead error in forecasting \mathbf{s}_i due to the shocks to \mathbf{s}_j . ‘From’ column reports total directional connectedness from all others to bank i and ‘To’ row reports total directional connectedness from bank j to all others. And ‘Net’ row reports the difference in total directional connectedness ‘To’-‘From’. The lower right number reports the average total connectedness.

Table 6: Connectedness of European and U.S. banks using the factor-vine copula model.

	BCS	BNP	CS	GLE	ISP	SAN	UBS	UCG	AXP	BAC	C	GS	JPM	MS	WFC	From
BCS	15.2	8.5	7.6	9.4	3.7	7.3	5.1	5.2	1.8	7.1	6.7	6.3	6.7	2.5	6.9	84.8
BNP	7.4	13.3	8.3	10.6	3.8	7.7	5.4	5.4	1.8	7.1	6.7	6.3	6.8	2.5	6.9	86.7
CS	6.4	8.0	15.8	8.9	3.6	7.1	8.8	5.1	1.7	6.8	6.4	6.0	6.4	2.4	6.6	84.2
GLE	8.1	10.6	9.1	12.3	4.0	8.6	5.7	5.8	1.7	6.7	6.3	5.9	6.4	2.4	6.5	87.7
ISP	5.8	6.9	6.8	7.3	9.1	9.1	4.8	9.9	1.9	7.5	7.0	6.7	7.2	2.7	7.3	90.9
SAN	6.6	8.1	7.7	9.0	5.2	13.4	5.1	7.8	1.7	6.9	6.5	6.1	6.6	2.5	6.8	86.6
UBS	5.8	7.0	11.7	7.4	3.4	6.3	14.8	4.8	1.8	7.2	6.8	6.3	6.9	2.6	7.0	85.2
UCG	6.0	7.2	7.0	7.7	7.3	9.9	4.9	11.1	1.8	7.3	6.8	6.5	7.0	2.6	7.1	88.9
AXP	5.7	6.5	6.4	6.2	3.8	6.1	5.1	5.1	7.9	9.2	8.6	8.3	8.8	3.3	9.0	92.1
BAC	5.8	6.7	6.6	6.3	3.9	6.2	5.2	5.2	2.4	12.0	9.2	8.6	9.2	3.4	9.4	88.0
C	5.8	6.7	6.6	6.3	3.9	6.2	5.3	5.1	2.4	9.8	11.0	8.6	9.3	3.5	9.5	89.0
GS	5.5	6.4	6.3	6.0	3.8	5.9	5.0	5.0	2.3	9.3	8.7	14.8	8.8	3.3	9.0	85.2
JPM	5.8	6.7	6.5	6.3	3.9	6.2	5.3	5.2	2.4	9.6	9.1	8.5	11.9	3.4	9.3	88.1
MS	5.7	6.6	6.5	6.2	3.9	6.2	5.2	5.2	2.4	9.4	9.1	8.4	9.0	6.9	9.3	93.1
WFC	5.8	6.6	6.6	6.3	3.9	6.2	5.2	5.2	2.4	9.6	9.1	8.5	9.1	3.4	12.3	87.7
To	86.3	102.4	103.7	103.9	58.1	98.9	76.3	79.8	28.4	113.7	106.8	101.0	108.1	40.3	110.6	87.9
Net	1.5	15.7	19.4	16.2	-32.7	12.4	-8.9	-9.1	-63.7	25.7	17.8	15.8	20.0	-52.8	22.8	

Notes: Connectedness of European and U.S. banks using the factor-vine copula model on March 16th, 2023. The 15 x 15 elements in the upper left show the shares $\mathbf{H}^i = [\hat{h}_{ij}^i]$ of variance decomposition of 5-step-ahead error in forecasting s_i due to the shocks to s_j . ‘From’ column reports total directional connectedness from all others to bank i and ‘To’ row reports total directional connectedness from bank j to all others. And ‘Net’ row reports the difference in total directional connectedness ‘To’-‘From’. The lower right number reports the average total connectedness.

References

- K. Aas, C. Czado, A. Frigessi, and H. Bakken. Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2):182–198, 2009.
- D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2 2015. ISSN 19447981. doi: 10.1257/aer.20130456.
- A. Agosto, D. F. Ahelegbey, and P. Giudici. Tree networks to assess financial contagion. *Economic Modelling*, 85:349–366, 2 2020. ISSN 02649993. doi: 10.1016/j.econmod.2019.11.005.
- M. Akhtaruzzaman, S. Boubaker, and J. W. Goodell. Did the collapse of Silicon Valley Bank catalyze financial contagion? *Finance Research Letters*, 56:104082, 9 2023. ISSN 15446123. doi: 10.1016/j.frl.2023.104082.
- R. Aloui, M. S. B. Aïssa, and D. K. Nguyen. Global financial crisis, extreme interdependencies, and contagion effects: The role of economic structure? *Journal of Banking and Finance*, 35(1):130–141, 1 2011. ISSN 03784266. doi: 10.1016/j.jbankfin.2010.07.021.
- A. Alter and A. Beyer. The dynamics of spillover effects during the European sovereign debt turmoil. *Journal of Banking and Finance*, 42(1):134–153, 5 2014. ISSN 03784266. doi: 10.1016/j.jbankfin.2014.01.030.

- A. Alter and Y. S. Schüler. Credit spread interdependencies of European states and banks during the financial crisis. *Journal of Banking and Finance*, 36(12):3444–3468, 12 2012. ISSN 03784266. doi: 10.1016/j.jbankfin.2012.08.002.
- V. Arakelian and P. Dellaportas. Contagion determination via copula and volatility threshold models. *Quantitative Finance*, 12(2):295–310, 2012. doi: 10.1080/14697680903410023.
- L. Ballester, B. Casu, and A. González-Urteaga. Bank fragility and contagion: Evidence from the bank CDS market. *Journal of Empirical Finance*, 38:394–416, 9 2016. ISSN 09275398. doi: 10.1016/j.jempfin.2016.01.011.
- T. Bedford and R. M. Cooke. Vines: A new graphical model for dependent random variables. *Annals of Statistics*, pages 1031–1068, 2002.
- J. Beirne and M. Fratzscher. The pricing of sovereign risk and contagion during the European sovereign debt crisis. *Journal of International Money and Finance*, 34:60–82, 4 2013. ISSN 02615606. doi: 10.1016/j.jimonfin.2012.11.004.
- M. B. Bjerregård, J. K. Møller, and H. Madsen. An introduction to multivariate probabilistic forecast evaluation. *Energy and AI*, 4:100058, 6 2021. ISSN 26665468. doi: 10.1016/j.egyai.2021.100058.
- D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- T. Bollerslev. Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *The Review of Economics and Statistics*, 72(3):498–505, 1990.
- E. C. Brechmann and H. Joe. Parsimonious parameterization of correlation matrices using truncated vines and factor analysis. *Computational Statistics & Data Analysis*, 77:233–251, 2014.
- E. C. Brechmann, C. Czado, and K. Aas. Truncated regular vines in high dimensions with application to financial data. *Canadian Journal of Statistics*, 40(1):68–85, 2012.
- E. C. Brechmann, K. Hendrich, and C. Czado. Conditional copula simulation for systemic risk stress testing. *Insurance: Mathematics and Economics*, 53(3):722–732, 2013.
- V. D. Bruyckere, M. Gerhardt, G. Schepens, and R. V. Vennet. Bank/sovereign risk spillovers in the european debt crisis. *Journal of Banking and Finance*, 37(12):4793–4809, 2013. ISSN 03784266. doi: 10.1016/j.jbankfin.2013.08.012.

- M. Cerrato, J. Crosby, M. Kim, and Y. Zhao. The joint credit risk of UK global-systemically important banks. *Journal of Futures Markets*, 37(10):964–988, 10 2017. ISSN 10969934. doi: 10.1002/fut.21855.
- W. Chen, K. C. Ho, and L. Yang. Network structures and idiosyncratic contagion in the European sovereign credit default swap market. *International Review of Financial Analysis*, 72, 11 2020. ISSN 10575219. doi: 10.1016/j.irfa.2020.101594.
- V. Constancio. Contagion and the European debt crisis. *Financial Stability Review*, 16: 109–121, 2012.
- J. Coval, J. Jurek, and E. Stafford. The economics of structured finance, 2009.
- D. D. Creal and R. S. Tsay. High dimensional dynamic stochastic copula models. *Journal of Econometrics*, 189(2):335–345, 2015.
- C. Czado and T. Nagler. Vine Copula Based Modeling. *Annual Review of Statistics and Its Application*, 9:453–477, 3 2022. ISSN 2326-8298. doi: 10.1146/annurev-statistics-040220-101153.
- F. X. Diebold and K. Yilmaz. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119–134, 2014.
- J. Dissmann, E. C. Brechmann, C. Czado, and D. Kurowicka. Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics & Data Analysis*, 59:52–69, 2013.
- R. Dornbusch, Y. C. Park, and S. Claessens. Contagion: understanding how it spreads. *The World Bank Research Observer*, 15(2):177–97, 2000.
- Y. Duan, S. E. Ghouli, O. Guedhami, H. Li, and X. Li. Bank systemic risk around covid-19: A cross-country analysis. *Journal of Banking and Finance*, 133:106299, 12 2021. ISSN 03784266. doi: 10.1016/j.jbankfin.2021.106299.
- M. Dungey and D. Gajurel. Contagion and banking crisis - International evidence for 2007-2009. *Journal of Banking and Finance*, 60:271–283, 2015. ISSN 03784266. doi: 10.1016/j.jbankfin.2015.08.007.
- X. Fan and H. Joe. High-dimensional factor copula models with estimation of latent variables. *Journal of Multivariate Analysis*, page 105263, 2023.

- C. Fernández and M. F. Steel. On bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association*, 93(441):359–371, 1998.
- Financial Stability Oversight Council. 2013 annual report. 2013.
- K. J. Forbes and R. Rigobon. No contagion, only interdependence: Measuring stock market comovements. *Journal of Finance*, 57(5):2223–2261, 2002. ISSN 00221082. doi: 10.1111/0022-1082.00494.
- P. Glasserman and H. P. Young. How likely is contagion in financial networks? *Journal of Banking & Finance*, 50:383–399, 2015.
- P. Hartmann, S. Straetmans, and C. de Vries. Banking system stability: A cross-atlantic perspective. *NBER Working Paper Series*, 11698:1–87, 2005.
- H. Joe. Asymptotic efficiency of the two-stage estimation method for copula-based models. *Journal of Multivariate Analysis*, 94(2):401–419, 2005.
- H. Joe. *Dependence Modeling with Copulas*. CRC Press, Boca Raton, FL, 2014.
- D. Kenourgios, A. Samitas, and N. Paltalidis. Financial crises and stock market contagion in a multivariate time-varying asymmetric framework. *Journal of International Financial Markets, Institutions and Money*, 21(1):92–106, 2 2011. ISSN 10424431. doi: 10.1016/j.intfin.2010.08.005.
- J. Kleinow and F. Moreira. Systemic risk among European banks: A copula approach. *Journal of International Financial Markets, Institutions and Money*, 42:27–42, 5 2016. ISSN 10424431. doi: 10.1016/j.intfin.2016.01.002.
- P. Krupskii and H. Joe. Factor copula models for multivariate data. *Journal of Multivariate Analysis*, 120:85 – 101, 2013.
- P. Krupskii and H. Joe. Structured factor copula models: Theory, inference and computation. *Journal of Multivariate Analysis*, 138:53–73, 2015.
- P. Krupskii and H. Joe. Flexible copula models with dynamic dependence and application to financial data. *Econometrics and Statistics*, 16:148–167, 10 2020. ISSN 24523062. doi: 10.1016/j.ecosta.2020.01.005.
- A. Kucukelbir, D. Tran, R. Ranganath, A. Gelman, and D. M. Blei. Automatic differentiation variational inference. *The Journal of Machine Learning Research*, 18(1):430–474, 2017.

- D. Kurowicka and R. M. Cooke. *Uncertainty analysis with high dimensional dependence modelling*. John Wiley & Sons, 2006.
- R.-J. Lange, A. Lucas, and A. H. Siegmann. Score-Driven Systemic Risk Signaling for European Sovereign Bond Yields and CDS Spreads. In M. Billio, L. Pelizzon, and R. Savona, editors, *Systematic Risk Tomography*, pages 129–150. Elsevier, 1 edition, 2016.
- A. Lucas, B. Schwaab, and X. Zhang. Conditional euro area sovereign default risk. *Journal of Business & Economic Statistics*, 32(2):271–284, 2014.
- S. Lyócsa, M. Halousková, and E. Haugom. The US banking crisis in 2023: Intraday attention and price variation of banks at risk. *Finance Research Letters*, 57, 11 2023. ISSN 15446123. doi: 10.1016/j.frl.2023.104209.
- A. M. Martins. Stock market effects of Silicon Valley bank and Credit Suisse failure: evidence for a sample of European listed banks. *Finance Research Letters*, 58(A):104296, 2023. ISSN 15446123. doi: 10.1016/j.frl.2023.104296.
- H. Nguyen. Bayesian inference for high dimensional factor copula models. *Doctoral Thesis, Universidad Carlos III de Madrid*, pages 1–138, 2019.
- H. Nguyen, M. C. Ausín, and P. Galeano. Variational inference for high dimensional structured factor copulas. *Computational Statistics and Data Analysis*, 151:107012, 11 2020. ISSN 01679473. doi: 10.1016/j.csda.2020.107012.
- D. H. Oh and A. J. Patton. Modelling dependence in high dimensions with factor copulas. *Journal of Business & Economic Statistics*, 35(1):139–154, 2017.
- D. H. Oh and A. J. Patton. Time-varying systemic risk: Evidence from a dynamic copula model of cds spreads. *Journal of Business and Economic Statistics*, 36(2):181–195, 2018. ISSN 15372707. doi: 10.1080/07350015.2016.1177535.
- D. H. Oh and A. J. Patton. Dynamic factor copula models with estimated cluster assignments. *Journal of Econometrics*, 237:105374, 2023.
- N. Paltalidis, D. Gounopoulos, R. Kizys, and Y. Koutelidakis. Transmission channels of systemic risk and contagion in the european financial network. *Journal of Banking & Finance*, 61:S36–S52, 2015.
- J. C. Rodriguez. Measuring financial contagion: a copula approach. *Journal of Empirical Finance*, 14(3):401–423, 2007. doi: 10.1016/j.jempfin.2006.07.002.

- M. A. Segoviano and C. Goodhart. Banking stability measures. *IMF Working Paper*, WP/09 (4):1–56, 2009.
- M. Sklar. *Fonctions de repartition a n dimensions et leurs marges*. Universite Paris 8, 1959.
- B. Stepankova and P. Teply. Consistency of banks’ internal probability of default estimates: Empirical evidence from the covid-19 crisis. *Journal of Banking and Finance*, 154:106969, 9 2023. ISSN 03784266. doi: 10.1016/j.jbankfin.2023.106969.
- G. Torri, R. Giacometti, and S. Paterlini. Robust and sparse banking network estimation. *European Journal of Operational Research*, 270(1):51–65, 2018.
- X. Ye, F. Yu, and R. Zhao. Credit derivatives and corporate default prediction. *Journal of Banking and Finance*, 138(106418):106418, 5 2022. ISSN 03784266. doi: 10.1016/j.jbankfin.2022.106418.
- D. M. Zimmer. The role of copulas in the housing crisis. *The Review of Economics and Statistics*, 94:607–620, 2012.