

Dynamic Factor Copula Models with Estimated Cluster Assignments*

Dong Hwan Oh Andrew J. Patton
Federal Reserve Board Duke University

First draft: 3 November 2020. This draft: 14 April 2022.

Abstract

This paper proposes a dynamic multi-factor copula for use in high dimensional time series applications. A novel feature of our model is that the assignment of individual variables to groups is estimated from the data, rather than being pre-assigned using SIC industry codes, market capitalization ranks, or other *ad hoc* methods. We adapt the *k*-means clustering algorithm for use in our application and show that it has excellent finite-sample properties. Applying the new model to returns on 110 US equities, we find around 20 clusters to be optimal. In out-of-sample forecasts, we find that a model with as few as five estimated clusters significantly outperforms an otherwise identical model with 21 clusters formed using two-digit SIC codes.

Keywords: high-dimensional models, risk management, multivariate density forecasting.

J.E.L. codes: C32, C38, C58.

*We thank participants at the 2020 EC² conference (Paris), Econometrics and Business Analytics conference (St. Petersburg), and the International Workshop on Financial Econometrics (Federal University of Rio Grande do Sul, Brazil). The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. We are grateful to Forrest Denson for his excellent computing support. Email: donghwan.oh@frb.gov, andrew.patton@duke.edu. A Matlab toolbox to implement the methods proposed in this paper will be available at <http://www.econ.duke.edu/~ap172>.

1 Introduction

Models for the dependence structure of a large collection of variables play an important role in risk management and regulation, yet there is a relative paucity of such models. A key impediment is that these models need to be parsimonious enough to deal with the inevitable curse of dimensionality that arises in high-dimensional applications, yet flexible enough to capture the time-varying and potentially asymmetric nature of the dependence between economic variables.

We propose a multi-factor, high-dimensional, copula model where the assignment of individual variables to groups or clusters is estimated from the data. Existing approaches for similar problems (see Creal and Tsay, 2015, Bester and Hansen, 2016, and Opschoor *et al.*, 2020, for example) use pre-specified cluster assignments, based on SIC industry codes, or market capitalization deciles, or similar. In the absence of a computationally feasible data-driven alternative such approaches are reasonable, however it is not obvious that such assignments are optimal empirically. We propose a method based on k -means clustering (see, e.g., Hastie *et al.*, 2009) to estimate the optimal assignments of variables to clusters, and we model dynamics in the conditional copula using a “generalized autoregressive score” (GAS) model (Creal *et al.*, 2013; Harvey 2013).

The estimation of the optimal cluster assignments for a high-dimensional dynamic copula model requires us to overcome two computational hurdles. Firstly, rather than the simulation-based factor copula model of Oh and Patton (2017), we adopt and extend the model of Opschoor *et al.* (2020), which has a closed-form likelihood and is thus much faster to estimate. Our extension enables us to capture asymmetric dependencies which can be important for equity returns, see Ang and Chen (2002), Hong *et al.* (2007) and Patton (2013) amongst many others. Secondly, we exploit the fact that the presence of clusters in the dynamic model implies the presence of clusters in the (misspecified) static version of the model. The static version of the model is naturally much faster to estimate than the dynamic version. These two techniques, combined with extensive use of parallel processing, make the estimation of optimal cluster assignments feasible.

We prove the consistency of the estimated cluster assignments under very mild conditions, and we find in realistically-designed simulations that our estimation method is remarkably accurate.

We apply the new model to daily returns on 110 U.S. equities over the period 2010-2019, and consider a range of choices for the number of clusters in the model. We find that the BIC optimal number of clusters is around 20, and moreover find that a model with just *five* estimated clusters outperforms an otherwise identical model based on 21 clusters formed using two-digit SIC groupings. In out-of-sample forecast comparisons, we find that the model with estimated cluster assignments significantly outperforms one with clusters formed using two-digit SIC codes.

This paper bridges two lines in the extant literature. Most directly, this paper is related to the literature on high-dimensional methods for financial risk measurement. Early work focused on improved methods for estimating large covariance matrices. For example, Fan *et al.* (2008, 2013) propose using a factor model where the number of factors grows with the number of variables, with the latter of these papers also accommodating approximate factor models. Tao *et al.* (2011) consider high-dimensional covariance matrix estimation based on a combination of high- and low-frequency data, also using a factor model. Hautsch *et al.* (2012) propose a method to estimate covariance matrices using high frequency data from assets with varying degrees of liquidity. More recent work in this area has included a focus on copula-based models, such as Creal and Tsay (2015) who proposed a high-dimensional stochastic copula with a factor structure, and Oh and Patton (2018) and Opschoor *et al.* (2020) who consider factor copulas with dynamics driven by a GAS specification. Christoffersen *et al.* (2018) propose a high-dimensional dynamic copula model with DCC (Engle, 2002) type dynamics. As far as we know, our paper is the first to consider a high-dimensional copula model with estimated group assignments.

This paper is also related to the fast-growing area of clustering and classification methods in economics and finance. Lin and Ng (2012) and Bonhomme and Manresa (2015) consider linear panel models with unknown group assignments which are estimated using k -means clustering. Su *et al.* (2016, 2019) consider panel models with group assignments estimated using a new type of LASSO estimator. The latter of these papers allows the parameters of the panel model to vary nonparametrically with time. Vogt and Linton (2020) also consider nonparametric regression for a panel of data with unknown group assignments. Francis *et al.* (2017) cluster countries by their business cycle patterns, and Patton and Weller (2019) consider clustering stocks by the risk premia

they generate. This research area is very active and this review is surely incomplete already.

The remainder of the paper is structured as follows. In Section 2 we present the dynamic copula models considered in this paper, and in Section 3 we discuss how we can optimally assign variables to clusters. Section 4 presents the results of a simulation study of the finite-sample performance of the proposed model and estimation method. Section 5 applies the new methods to a collection of 110 stock returns. Section 6 concludes, and the appendix contains proofs and technical details. A web appendix contains additional analyses and material.

2 A dynamic skewed t factor copula model

A copula is an N -dimensional distribution function with $Unif(0, 1)$ margins, and even when N is only moderately-sized the curse of dimensionality arises. A common approach to overcome this in other contexts is to impose some sort of factor structure, and recent work on high-dimensional copula models has moved in this direction, see Oh and Patton (2017, 2018), Creal and Tsay (2015) and Opschoor *et al.* (2020). An attractive feature of the latter two papers is that the copula likelihood is available in closed form. Motivated by previous work showing that equity returns exhibit asymmetric dependence (see, e.g., Ang and Chen, 2002, Hong *et al.*, 2007, and Patton, 2013), we consider an extension of the model proposed by Opschoor *et al.* (2020) to allow for asymmetric dependence, namely a skewed t factor copula:

$$u_{i,t} = T_{skew}(x_{i,t}; \nu, \zeta), \quad i = 1, \dots, N, \quad (1)$$

$$x_{i,t} = \sqrt{W_t} \left(\tilde{\boldsymbol{\lambda}}_{i,t}' \mathbf{z}_t + \sigma_{i,t} \epsilon_{i,t} \right) + \zeta W_t, \quad (2)$$

$$\text{where } \mathbf{z}_t \sim iid N(\mathbf{0}, \mathbf{I}_k), \quad \epsilon_{i,t} \sim iid \mathcal{N}(0, 1), \quad (3)$$

$$W_t \sim iid IG\left(\frac{\nu}{2}, \frac{\nu}{2}\right), \quad W_t \perp \mathbf{z}_t \perp \epsilon_{i,t} \quad (4)$$

where $T_{skew}(\cdot; \nu, \zeta)$ denotes the univariate skewed t CDF of $x_{i,t}$, with degrees of freedom parameter $\nu \in (2, \infty]$ and asymmetry parameter $\zeta \in [-1, 1]$.¹ $\tilde{\boldsymbol{\lambda}}_{i,t}$ is a vector of scaled factor loadings, \mathbf{z}_t is a

¹Creal and Tsay (2010) describe this copula but do not implement it or present results on its likelihood and scores. As that paper notes, when $\zeta \neq 0$ the function $T_{skew}(\cdot; \nu, \zeta)$ is not available in closed form, and Creal and Tsay (2010)

vector of common latent factors and $\epsilon_{i,t}$ is an idiosyncratic shock, both Normally distributed, and W_t is an inverse gamma variable. We define the vector $\tilde{\boldsymbol{\lambda}}_{i,t}$ and scalar $\sigma_{i,t}$ as

$$\tilde{\boldsymbol{\lambda}}_{i,t} = \frac{\boldsymbol{\lambda}_{i,t}}{\sqrt{1 + \boldsymbol{\lambda}'_{i,t} \boldsymbol{\lambda}_{i,t}}}, \quad \sigma_{i,t}^2 = \frac{1}{1 + \boldsymbol{\lambda}'_{i,t} \boldsymbol{\lambda}_{i,t}} \quad (5)$$

for a factor loading $\boldsymbol{\lambda}_{i,t}$ to maintain the unit variance of $\tilde{\boldsymbol{\lambda}}'_{i,t} \mathbf{z}_t + \sigma_{i,t} \epsilon_{i,t}$. The skewed t copula nests the Student's t copula when $\zeta = 0$, and the Gaussian copula when $\zeta = 0$ and $\nu \rightarrow \infty$. Given this structure, and since each element of the vector $\left[\tilde{\boldsymbol{\lambda}}'_{1,t} \mathbf{z}_t + \sigma_{1,t} \epsilon_{1,t}, \dots, \tilde{\boldsymbol{\lambda}}'_{N,t} \mathbf{z}_t + \sigma_{N,t} \epsilon_{N,t} \right]$ has unit variance, its covariance matrix is a correlation matrix, \mathbf{R}_t , with the form:

$$\mathbf{R}_t = \tilde{\mathbf{L}}'_t \tilde{\mathbf{L}}_t + \mathbf{D}_t \quad (6)$$

where $\tilde{\mathbf{L}}_t = \left[\tilde{\boldsymbol{\lambda}}_{1,t}, \dots, \tilde{\boldsymbol{\lambda}}_{N,t} \right]$ and $\mathbf{D}_t = \text{diag} \left(\sigma_{1,t}^2, \dots, \sigma_{N,t}^2 \right)$. The skewed t copula then contains time-varying factor loadings $\left[\boldsymbol{\lambda}'_{1,t}, \dots, \boldsymbol{\lambda}'_{N,t} \right]$ and static shape parameters (ν, ζ) . Creal and Tsay (2015) show that a factor copula structure of the sort in equation (2) facilitates the evaluation of the copula density even for high dimensions since the inverse and determinant of \mathbf{R}_t are available in closed form and require only lower-dimension inversions and determinant calculations:

$$\begin{aligned} \mathbf{R}_t^{-1} &= \mathbf{D}_t^{-1} - \mathbf{D}_t^{-1} \tilde{\mathbf{L}}'_t \left(\mathbf{I}_k + \tilde{\mathbf{L}}_t \mathbf{D}_t^{-1} \tilde{\mathbf{L}}'_t \right)^{-1} \tilde{\mathbf{L}}_t \mathbf{D}_t^{-1} \\ |\mathbf{R}_t| &= \left| \mathbf{I}_k + \tilde{\mathbf{L}}_t \mathbf{D}_t^{-1} \tilde{\mathbf{L}}'_t \right| \cdot |\mathbf{D}_t|. \end{aligned}$$

We consider a factor structure determined by a $(G + 1)$ vector \mathbf{z}_t of common latent factors and a loading matrix $\tilde{\mathbf{L}}_t$. Specifically, we allow for one common factor, shared by all variables, and G cluster-specific factors, shared only by members of that cluster.² For example, assuming there are G groups and each group has only two members, \mathbf{z}_t and $\tilde{\mathbf{L}}_t$ are determined by:

omit it from their analysis. The presence of this parameter raises no theoretical difficulties, only a computational one. In Appendix A.1 we describe a simple and computationally tractable method to overcome this difficulty, making the likelihood of this copula *quasi*-closed form (up to a simple one-dimensional numerical integral).

²This corresponds to the “multi-factor” (MF) model in Opschoor *et al.*, (2020), which is the second-most flexible factor structure considered in that paper. The more flexible “lower-triangular MF” model is less amenable to the estimation of group assignments, which is the key focus of this paper, and we do not consider that structure here.

$$\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{G+1})$$

$$\tilde{\mathbf{L}}'_t = \begin{pmatrix} \tilde{\lambda}_{1,t}^M & \tilde{\lambda}_{1,t}^C & 0 & 0 & \cdots & 0 \\ \tilde{\lambda}_{2,t}^M & 0 & \tilde{\lambda}_{2,t}^C & 0 & \cdots & 0 \\ \tilde{\lambda}_{3,t}^M & 0 & 0 & \tilde{\lambda}_{3,t}^C & 0 & 0 \\ \vdots & \vdots & \vdots & 0 & \ddots & 0 \\ \tilde{\lambda}_{G,t}^M & 0 & 0 & 0 & 0 & \tilde{\lambda}_{G,t}^C \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad (7)$$

where \otimes denotes the Kronecker product. Note that the loadings on the common factor and the cluster-specific factor can only be separately identified if each group has at least two members; we impose this condition when estimating the model.

Next, we formulate the dynamics of $2G$ distinct factor loadings based on the generalized autoregressive score model proposed by Creal *et al.* (2013) and Harvey (2013). Specifically, we model those dynamics by:

$$\lambda_{g,t+1}^M = \omega_g^M + \alpha^M \frac{\partial \log \mathbf{c}_{Skew,t,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \lambda_{g,t}^M} + \beta^M \lambda_{g,t}^M, \quad \text{for } g = 1, \dots, G \quad (8)$$

$$\lambda_{g,t+1}^C = \omega_g^C + \alpha^C \frac{\partial \log \mathbf{c}_{Skew,t,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \lambda_{g,t}^C} + \beta^C \lambda_{g,t}^C, \quad \text{for } g = 1, \dots, G$$

where $\mathbf{x}_t = T_{skew}^{-1}(\mathbf{u}_t; \nu, \zeta)$, $\mathbf{c}_{Skew,t,t}(\cdot; \mathbf{R}_t, \nu, \zeta)$ is the conditional skewed t copula density and $[\omega_1^M, \dots, \omega_G^M, \omega_1^C, \dots, \omega_G^C, \alpha^M, \beta^M, \alpha^C, \beta^C]'$ is the vector of parameters determining the dynamics of time varying factor loadings.³ Obviously the key component is the score of the conditional copula $\partial \log \mathbf{c}_{Skew,t,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta) / \partial \boldsymbol{\eta}_t$ where $\boldsymbol{\eta}_t$ is a $(2G \times 1)$ vector of all dynamic factor loadings:

$$\boldsymbol{\eta}_t = [\lambda_{1,t}^M, \dots, \lambda_{G,t}^M, \lambda_{1,t}^C, \dots, \lambda_{G,t}^C]'. \quad (9)$$

The skewed t copula density and the analytical derivation of its score are given in Appendix A.1. and A.2. respectively.

While our model has factor loadings that vary across time, we assume that the group assign-

³As in Opschoor *et al.* (2020) and Oh and Patton (2018), we use a unit scaling of the score in equation (8), rather than the inverse Hessian or its square-root, to reduce the computational burden of estimating the model.

ments are stable. Custodio João *et al.* (2022) and Lumsdaine *et al.* (2022) consider models with time-varying group assignments, and generalizing our framework to allow this is an interesting extension for future research.

3 Clustering and factor copulas

3.1 Clustering via a misspecified model

While the closed-form density and GAS equations presented in equation (8) greatly reduce the computational burden of estimating a dynamic high-dimensional copula model, this model is still too costly to use when combined with an EM algorithm to estimate group assignments from the data. In this section we show that the structure of our model is such that we can estimate group assignments based on a simpler, misspecified, model, overcoming this hurdle.

Firstly, consider a static skew t factor copula. The factor loading vectors ($\boldsymbol{\lambda}_i$) obey a cluster structure, in that all variables in the same cluster have the same loading vector. From equation (5) above, given the factor loadings we can obtain the normalized loadings and idiosyncratic variances, $\tilde{\boldsymbol{\lambda}}_{i,t}$ and $\sigma_{i,t}^2$, and from those we obtain the correlation matrix:

$$\mathbf{R} = \tilde{\mathbf{L}}'\tilde{\mathbf{L}} + \mathbf{D}$$

where $\tilde{\mathbf{L}} = [\tilde{\boldsymbol{\lambda}}_1, \dots, \tilde{\boldsymbol{\lambda}}_N]$ and $\mathbf{D} = \text{diag}(\sigma_1^2, \dots, \sigma_N^2)$. The cluster structure embedded in $\boldsymbol{\lambda}_i$ implies that \mathbf{R} exhibits a block structure, which, as discussed above, can be used to speed up matrix inverse and determinant calculations. Further, we note that the block structure in \mathbf{R} holds *regardless* of the shape parameters (ν, ζ) . Thus a Normal factor copula, where the shape parameters are incorrectly fixed at $(\nu, \zeta) = (\infty, 0)$ will exhibit the same cluster structure as the more complicated skew t factor copula. This means that the cluster assignments implied by the Normal factor copula are identical to the skew t factor copula, permitting us to use the simpler model to estimate cluster assignments, with the usual caveat that these estimates are likely less precise than those based on the true model.

Next consider a time-varying skew t factor copula. In this case the time-varying correlation matrix $\mathbf{R}_t = \widetilde{\mathbf{L}}_t' \widetilde{\mathbf{L}}_t + \mathbf{D}_t$ obeys a block structure, and while the values taken by the elements of \mathbf{R}_t vary over time, the block structure is constant due to the maintained assumption that group assignments are stable. The conditional marginal copula of any pair $(u_{i,t}, u_{j,t})$ is determined completely by $(R_{i,j,t}, \nu, \zeta)$, and any pair of variables (i, j) belonging to groups (g_1, g_2) will have the same distribution as any other pair (i', j') belonging to the same two groups. The unconditional marginal copula is just an integral of the conditional marginal copula, and so the unconditional rank correlation matrix, $\bar{\rho} \equiv \text{Corr}[\mathbf{u}_t]$, exhibits the same cluster structure as the conditional correlation matrix \mathbf{R}_t , opening up the possibility of using a constant Normal factor copula to estimate group assignments for a dynamic skew t factor copula.

One complication arises when using a static copula to determine group assignments for a dynamic DGP: since we are taking time series averages, it is possible that the unconditional rank correlation matrix $\bar{\rho}$ is more homogeneous than the conditional correlation matrix \mathbf{R}_t , making it harder to identify group assignments. That is, clusters may not be as well separated in the approximating model as in the true model. The concept of “well separatedness” is a finite-sample issue, and we examine it in detail in our simulation study. To preview our findings, our simulations indicate that this is not a significant concern here.

3.2 Estimation of cluster assignments and copula parameters

The main advantage of using a factor copula comes from the dimension reduction enabled by classifying variables into a relatively small number of clusters and assuming identical factor loadings within each cluster. In the existing literature, variables are clustered according to observable characteristics, such as SIC industry classifications. Given those cluster assignments, the factor copula can be estimated via maximum likelihood under standard conditions, however, the *ex ante* assignments of variables to clusters may not provide the best fit to the data.

We propose an iterative method which estimates cluster assignments, and copula parameters, directly from the data, exploiting an expectation-maximization (EM) algorithm. This algorithm cycles between (1) estimating copula parameters given cluster assignments and (2) estimating clus-

ter assignments given the estimated copula parameters. Let $\Gamma = [\gamma_1, \dots, \gamma_N]$ where $\gamma_i \in \{1, \dots, G\}$ for $i = 1, \dots, N$, denote the vector of cluster assignments, and let $\boldsymbol{\theta} = [\lambda_1^M, \dots, \lambda_G^M, \lambda_1^C, \dots, \lambda_G^C]$ be the vector of market and cluster-specific factor loadings used to obtain the correlation matrix parameter for the static Gaussian factor copula, with log-likelihood denoted $\log \mathbf{c}(\cdot)$. Given an estimate of the cluster assignment vector, $\hat{\Gamma}^{(s)}$ the log-likelihood of the copula model is maximized over the copula parameters $\boldsymbol{\theta}$ to yield:

$$\hat{\boldsymbol{\theta}}^{(s+1)} = \arg \max_{\boldsymbol{\theta}} \hat{Q}_T(\boldsymbol{\theta}, \hat{\Gamma}^{(s)}) \quad (10)$$

$$\text{where } \hat{Q}_T(\boldsymbol{\theta}, \Gamma) \equiv \sum_{t=1}^T \log \mathbf{c}(\mathbf{u}_t; \boldsymbol{\theta}, \Gamma) \quad (11)$$

Then, given copula parameter $\hat{\boldsymbol{\theta}}^{(s+1)}$, the log-likelihood is maximized over cluster assignments γ_i for $i = 1, \dots, N$:

$$\hat{\gamma}_i^{(s+1)} = \arg \max_{g \in \{1, \dots, G\}} \hat{Q}_T(\hat{\boldsymbol{\theta}}^{(s+1)}, \tilde{\Gamma}_{i,g}^{(s)}) \quad (12)$$

where $\tilde{\Gamma}_{i,g}^{(s)}$ is equal to $\hat{\Gamma}^{(s)}$ except that the i^{th} element is set equal to g .

The copula parameter in equation (10) is estimated through a typical gradient-based optimization. We update each variable's cluster assignment (equation 12) by re-optimizing the cluster assignments one variable at a time, motivated by the method underlying k -means clustering. This latter step requires only $G \times N$ likelihood evaluations, making cluster assignment estimation feasible and fast.⁴ The iteration between equation (10) and equation (12) continues until convergence. Convergence to a local optimum is guaranteed, and we use 10 randomly-chosen starting values to improve the accuracy of the estimator. Our simulation study below confirms this to be a sufficient number of starting values. Denote the resulting estimates as $(\hat{\boldsymbol{\theta}}_T, \hat{\Gamma}_T)$.

We next provide conditions under which the estimated cluster assignments, $\hat{\Gamma}_T$, are consistent for the true cluster assignments, Γ_0 . This is a non-standard estimation problem as the parameter Γ_0 is discrete: each of its N elements can take one of only G values. Let \mathcal{G} denote the parameter space

⁴Other estimation algorithms for k -means type problems have been proposed in the computer science/machine learning literature. Given the very good finite-sample performance we find for the algorithm described here, when a sufficient number of starting values is used, we did not consider any alternatives.

for Γ .⁵ Since the labels attached to clusters are arbitrary (i.e., the objective function is invariant to relabeling the clusters), there is a *set* of correct cluster labels, rather than just a singleton; let \mathcal{G}_0 denote this set. To state the assumptions we define the following:

$$\tilde{\boldsymbol{\theta}}^*(\Gamma) = \arg \min_{\boldsymbol{\theta} \in \Theta} \mathbb{E} [\log \mathbf{c}(\mathbf{u}_t; \boldsymbol{\theta}, \Gamma)] \quad (13)$$

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta} \in \Theta} \mathbb{E} [\log \mathbf{c}(\mathbf{u}_t; \boldsymbol{\theta}, \Gamma_0)] \quad (14)$$

Note that the parameter $\boldsymbol{\theta}^*$ is a pseudo-true parameter: it is the optimal parameter for the misspecified static Gaussian copula model. We obtain this parameter as a by-product of estimating the cluster assignments, but we have no subsequent use for it.

Assumption 1: $\{\mathbf{u}_t\}$ is a stationary ergodic sequence.

Assumption 2: For each $\Gamma \in \mathcal{G}$, (a) $|\log \mathbf{c}(\mathbf{u}_t; \boldsymbol{\theta}, \Gamma)|_1 < \infty \forall \boldsymbol{\theta} \in \Theta$, (b) $\left\| \nabla_{\boldsymbol{\theta}} \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma) \right\|_1 < \infty$, and (c) $\left\| \nabla_{\boldsymbol{\theta}\boldsymbol{\theta}} \log \mathbf{c}(\mathbf{u}_t; \boldsymbol{\theta}, \Gamma) \right\|_1 < \infty \forall \boldsymbol{\theta} \in \Theta$.

Assumption 3: (a) For each $\Gamma \in \mathcal{G}$, $\limsup_{T \rightarrow \infty} \left[\hat{Q}_T(\tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma) - \hat{Q}_T(\boldsymbol{\theta}, \Gamma) \right] > 0 \forall \boldsymbol{\theta} \in \Theta \setminus \eta_T(\varepsilon)$, where $\eta_T(\varepsilon)$ is an ε -neighborhood of $\tilde{\boldsymbol{\theta}}^*(\Gamma)$, and (b) $\limsup_{T \rightarrow \infty} \left[\hat{Q}_T(\boldsymbol{\theta}^*, \Gamma_0) - \hat{Q}_T(\boldsymbol{\theta}, \Gamma) \right] > 0 \forall (\boldsymbol{\theta}, \Gamma) \in \{\Theta \setminus \eta_T(\varepsilon)\} \times \{\mathcal{G} \setminus \mathcal{G}_0\}$.

Assumption 1 allows for general forms of serial dependence in the data (e.g., mixing). Importantly, given that we expect the static Gaussian copula model to be misspecified, it does *not* require correct specification of the conditional copula. Assumption 2, combined with Assumption 1, ensures that the log-likelihood and its first and second derivatives each obey a law of large numbers. Assumption 3 is a standard “identifiable uniqueness” assumption required for estimation, see Definition 3.3 of White (1994). In our application, it requires that the clusters are “well separated.” If the clusters are too close together, then identification of the clusters breaks down. A similar assumption is made in, e.g., Hahn and Moon (2010) and Bonhomme and Manresa (2015). The proof of the following theorem is in Appendix A.4.

⁵Recall that for identification of our model we require all clusters to have at least two members. We restrict \mathcal{G} to impose this condition.

Theorem 1 *Under Assumptions 1-3 we have $\Pr \left[\hat{\Gamma}_T \in \mathcal{G}_0 \right] \xrightarrow{p} 1$ as $T \rightarrow \infty$.*

Results from related contexts suggest that if the series $\{\mathbf{u}_t\}$ generated by equation (1) satisfies certain mixing properties, a large deviations principle may be applied (e.g., see Hahn and Moon, 2010, Choirat and Seri, 2012, and Bonhomme and Manresa, 2015). This enables obtaining a rate result, refining the consistency result in Theorem 1. Specifically, estimated cluster assignments have been shown in some applications to be superconsistent, with estimation errors taking the form:

$$\Pr \left[\hat{\Gamma}_T \notin \mathcal{G}_0 \right] \leq C_1 \exp \{-C_2 T^\kappa\} \quad (15)$$

for some constants $C_1, C_2, \kappa > 0$.⁶ The simulation results presented in the next section reveal that cluster assignments are indeed estimated extremely well, in line with a superconsistent rate of convergence, though, unfortunately, general results on the mixing properties of GAS processes are not yet available in the literature, and so we do not pursue a theoretical result of this nature here.⁷ A result of the form in equation (15) implies that estimation error in estimated cluster assignments vanishes faster than the usual \sqrt{T} rate, and standard errors on the remaining model parameters can be computed as though group assignments were known. If errors in estimated group assignments are of the same asymptotic order as those in the remaining model parameters, then standard errors on the remaining parameters need to be adjusted.⁸ The primary research questions of this paper do not require us to take a stand on the rate of convergence of the estimated group assignments.

With the estimated the cluster assignments $\hat{\Gamma}_T$ in hand, we estimate the parameters of the skewed t copula with GAS dynamics:

$$\hat{\boldsymbol{\psi}}_T = \arg \max_{\boldsymbol{\psi}} \sum_{t=1}^T \log \mathbf{c}_{Skewt,t} \left(\mathbf{u}_t; \boldsymbol{\psi} | \hat{\Gamma}_T \right)$$

⁶For example, Hahn and Moon (2010) provide conditions under which alpha mixing implies $\kappa = 1/2$, and phi mixing implies $\kappa = 1$. The constants C_1, C_2 vary with the specifics of the application.

⁷Related to the GAS context considered here, Carrasco and Chen (2002) and Hafner and Preminger (2009) show that univariate and multivariate GARCH processes, respectively, are beta mixing. Some results on the stationarity and ergodicity of univariate GAS processes are presented in Blasques et al. (2014).

⁸The supplemental appendix of Bonhomme and Manresa (2015) discusses a bootstrap approach for their application, however they note that it is borderline computationally prohibitive even in their linear, non-dynamic, model. To the best of our knowledge, the literature does not yet contain results allowing for the theoretical analysis of a bootstrap method for the dynamic time series applications considered here.

where $\boldsymbol{\psi} = [\omega_1^M, \dots, \omega_G^M, \omega_1^C, \dots, \omega_G^C, \alpha^M, \beta^M, \alpha^C, \beta^C, \nu, \zeta]'$. As the parameter $\boldsymbol{\psi}$ is large, we adopt a “variance targeting” approach to separately estimate the intercept parameters $[\omega_1^M, \dots, \omega_G^C]$, leaving us with only six parameters that require difficult numerical optimization. Details on this method are described in Appendix A.3. Once $\hat{\boldsymbol{\psi}}_T$ is obtained, the time series of factor loadings, $\boldsymbol{\lambda}_t$, can be computed using equation (8).

4 Simulation study

We investigate the finite-sample performance of the estimation method proposed above in a simulation study designed to match the key features of our empirical application below. We consider a sample size of $T = 1000$ and a collection of $N = 100$ variables, and three different factor copulas: a Gaussian factor copula, a t factor copula, and a skew t factor copula, corresponding to $[\nu, \zeta] = [\infty, 0], [5, 0], [5, -0.1]$ respectively. For illustration, a sample of bivariate data from these three copulas, as well as a skew Normal copula which we omit from the simulation study, is presented in Figure 1. In all cases the linear correlation is 0.5, and to aid the interpretation we transform draws from these copulas using the inverse Normal CDF, and so these four distributions all have $N(0, 1)$ marginal distributions. In the upper-left panel of Figure 1, we see the familiar bivariate Normal distribution, with low dependence in the tails and displaying radial symmetry. The upper-right panel displays the Student’s t copula, which is also radially symmetric but exhibits tail dependence, which manifests in this figure as realizations that lie close to the main diagonal in the upper and lower joint tails. The lower two panels present asymmetric copulas, with dependence being stronger in the lower tail than the upper tail, particularly for the skew t copula which exhibits non-zero tail dependence.

[INSERT FIGURE 1 ABOUT HERE]

We consider two cases for the dynamics of the copula: the benchmark static case, in which the conditional copula is constant, and the case of empirical interest, where the parameters of the copula evolve according to the GAS model introduced in Section 2. We set the number of clusters, G , to be 10 or 20, with an equal number of variables allocated to each cluster. In the static case, we

assume that the loadings across groups on the market factor range from 0.25 to 2.50 in increments of 0.25, while the loadings on the group specific factors range from 2.5 to 0.25 in increments of -0.25. This implies that some groups are more influenced by the common market factor than their group factor, while the reverse is true for other groups, roughly mimicking the differences between industries like manufacturing and mining/construction. Naturally, in this case the GAS dynamic parameters $(\alpha^M, \beta^M, \alpha^C, \beta^C)$ are all zero.

In the dynamic case, we set the intercept parameters (ω_g^M, ω_g^C) equal to 0.04 for all groups, which, combined with the common values for the GAS dynamic parameters $(\alpha^M, \beta^M, \alpha^C, \beta^C) = (0.02, 0.9, 0.02, 0.9)$, means that all groups have the same average loading on the market factor and on their group-specific factor. This homogeneity of loadings makes the estimation problem more difficult than if the loadings had different long-run averages, and is designed to further interrogate the ability of our clustering method to correctly assign variables to groups.

Table 1 presents the results for the static copula case with $G = 10$. In Panel A we see that the estimated parameters are centered on the true values, for all three copulas, and the standard errors on the factor loadings increase slightly (on average) as we move from Gaussian to t to skew t copulas.

Panel B of Table 1 reports the striking result that in 100% of the simulations there were *zero* variables assigned to an incorrect group. That is, in every simulation the clustering algorithm was able to correctly allocate variables to their groups.⁹ In the Gaussian case, the clustering step is done using the correct model (a static Gaussian copula) while in the other two cases the model used in the clustering step is misspecified. Panel B reveals that this misspecification leads to no errors in the classification of these variables.¹⁰ This is consistent with the exponential convergence rate (see equation 15) found in other contexts for cluster assignment estimators.

Panel C of Table 1 reports the average estimation time (using a machine with an Intel Xeon Gold 6132 processor, with ten cores and clock speed of 2.60GHz) and number of EM iterations

⁹Recall that groups are identified only up to a re-labelling; we account for this when computing the accuracy of the estimated group assignments.

¹⁰The clustering algorithm is not, of course, infallible: its accuracy depends on the structure of the DGP and the data available. In situations where the clusters are close together relative to sampling variation, estimated cluster assignments will inevitably contain errors. In our realistically-calibrated simulation design, the clusters appear to be sufficiently well separated that cluster assignments can be very accurately estimated.

required for convergence, and reveals no large differences in the difficulty of estimation across these models.

[INSERT TABLE 1 ABOUT HERE]

Table 2 presents the results for the dynamic copula case with $G = 10$. We again see that the estimated parameters are centered on the true values, and in Panel B we see the remarkable result that the clustering algorithm described above is able to correctly assign every variable to its group in 100% of simulations. Recall that the estimated cluster assignments are based on a static Gaussian copula model, which is misspecified in all three cases considered in Table 2. That model is shown in Table 2 to be rich enough to reveal the true clusters in the data even though it is misspecified, confirming the discussion in Section 3.1. Panel C of Table 2 shows that the clustering step for the dynamic model is almost as fast as for the static case (where one source of model misspecification is removed), while the copula parameter estimation step is naturally slower.

Table S1 in the Supplemental Appendix considers a design analogous to that for Table 2, except that we set the number of clusters to be 20 rather than 10. This is a more challenging estimation problem, and the time required for the cluster assignment estimation is greater (around 22 minutes compared with around 9 minutes for the $G = 10$ case), as is the time required for the copula estimation (around 60 minutes compared with around 40 minutes). In this design we also observe one, but only one, case where a variable is missclassified in the cluster assignment step, for the skew t copula DGP.

Overall, the results in Tables 1, 2 and S1 provide strong reassurance that the models and estimation methods proposed in Section 3 work well in finite samples, enabling us to take these to real data in the next section.

[INSERT TABLE 2 ABOUT HERE]

5 Empirical application

5.1 Data and summary statistics

We study daily equity returns over the period from January 4, 2010 to December 31, 2019, a total of $T = 2516$ trade days. Every stock that was ever a constituent of the S&P 100 index during this sample, and which traded for the full sample period, is included in the data set, yielding a total of $N = 110$ firms. A list of those firms, including their names, ticker symbols, and two-digit Standard Industrial Classification (SIC) codes, are provided in Table S2 in the supplemental appendix.

Table 3 presents summary statistics of the data and parameter estimates for the mean, variance and marginal distribution models. Panel A presents unconditional sample moments of the daily returns for each stock, and these moments are comparable to those observed in other studies. Given the skewness and kurtosis estimates reported in Panel A, our marginal distribution model combines an AR(1) for the conditional mean, GJR-GARCH(1,1) for the conditional variance, and a skewed t for the marginal distribution of the standardized residuals:

$$\begin{aligned} r_{i,t} &= \phi_{0i} + \phi_{1i}r_{i,t-1} + \epsilon_{i,t} \\ h_{i,t} &= \varpi_i + \beta_i h_{i,t-1} + \alpha_i \epsilon_{i,t-1}^2 + \kappa_i \epsilon_{i,t-1}^2 \mathbf{1}\{\epsilon_{i,t-1} \leq 0\} \\ \frac{\epsilon_{i,t}}{\sqrt{h_{i,t}}} &\sim iid \text{Skew } t(\xi_i, \psi_i) \end{aligned}$$

where $h_{i,t}$ is the conditional variance at time t for firm i and *Skew t* is the univariate skewed t distribution of Hansen (1994) with the tail parameter ξ_i and the asymmetry parameter ψ_i . Using quasi-maximum likelihood, we estimate the conditional mean and variance models, then given those estimated standardized residuals, we estimate the skewed t parameters. Panel B of Table 3 provides the estimation results of the marginal distribution model, and the values there are consistent with those reported in the empirical finance literature (see, e.g., Bollerslev, Engle, and Nelson 1994). The standardized residuals still indicate substantial skewness ($\hat{\psi} = -0.027$ on average) and kurtosis ($\hat{\xi} = 5.089$ on average). Given the marginal model parameters we obtain the probability integral transforms, u_{it} , used in the estimation of the copula.

Panel C of Table 3 presents Pearson’s linear correlations and Spearman’s rank correlations between those standardized residuals whose quantiles between 5% and 95% range from 0.17 to 0.49 and from 0.20 to 0.53, respectively, indicating heterogeneous pairwise dependence, and motivating our flexible factor copula specification presented in Section 2.

[INSERT TABLE 3 ABOUT HERE]

5.2 Estimated cluster assignments

We firstly use the method described in Sections 3 to estimate the group assignments for each variable. To determine the optimal number of groups, we use the BIC for the fitted static Gaussian copula model.¹¹ The value of the BIC for each choice of G is plotted in Figure 2, along with the values of the BIC obtained when using one-digit or two-digit SIC codes to determine group assignments. In our sample there are seven one-digit SIC groups and 21 two-digit SIC groups.¹² Figure 2 reveals that the BIC from a model using only four estimated group assignments dominates the seven one-digit SIC groups, and a model with just five estimated group assignments beats the 21-group model based on two-digit SIC codes. These rankings reveal the gains available from a data-driven assignment of stocks to groups, rather than assignments based on SIC codes.¹³

The optimal number of estimated groups, according to the BIC, is 21, which is coincidentally the same as the number of two-digit SIC groups.¹⁴ We note that the BIC curve is relatively flat near the optimum, indicating that choosing G between 20 and 25 leads to approximately the same fit; i.e., there is some robustness to the specific choice of G .

¹¹The BIC is computed as $BIC(G) = -2 \sum_{t=1}^T \log c(\mathbf{u}_t; \hat{\boldsymbol{\theta}}_T^{(G)}, \hat{\Gamma}_T^{(G)}) + 2G \log(T)$, where G denotes the number of clusters, leading to $2G$ parameters to be estimated. We use the notation $(\hat{\boldsymbol{\theta}}_T^{(G)}, \hat{\Gamma}_T^{(G)})$ to emphasize that the parameters of the copula vary with G .

¹²Our model cannot accommodate groups with only one member, and when estimating with SIC-based clusters we address this by moving stocks that are a singleton in their group to the SIC group with which they have the highest correlation. Specifically, in the one-digit clustering model, Weyerhaeuser (WY) is the only stock in the one-digit SIC group 0, and we move it to SIC group 3. In the two-digit clustering model, FCX (10), NKE (30), WY (08), FDX (45) and V(61) are all singletons, and those are moved into the two-digit SIC groups 13, 37, 37, 42, and 60, respectively.

¹³Opschoor *et al.* (2020) compare cluster assignments based on SIC codes with those based on some other common characteristics: market capitalization (size), the book-to-market ratio (value), and past returns (momentum). They find that SIC-based assignments easily dominate these alternatives.

¹⁴We used a set of 100 random starting values for Γ , the cluster assignment vector, in estimation, and did not use information from SIC codes at all in the EM-based model.

[INSERT FIGURE 2 ABOUT HERE]

Table 4 presents the estimated group assignments for the 110 stocks in our sample, along with each stock’s SIC code. Some of the estimated groups line up closely with a two-digit SIC group. For example, the largest group (Group 1) is comprised of 13 stocks, ten of which have SIC code 28 (“Chemical & Allied Products” manufacturing). The three other stocks (Baxter, Medtronic and United Health) have different SIC codes, but are clearly broadly in the same category as the rest of this group. Group 5, as another example, looks clearly like a “Tech” group, and all but two members have SIC code 73 (“Business Services”). The two listed with other codes are Apple (listed as 35, “Industrial Machinery & Equipment” manufacturing) and Netflix (listed as 78, “Motion Pictures”). Despite the different SIC codes, most investors would agree that Apple and Netflix fit neatly in a cluster containing Google, Amazon and Ebay. Among the smaller clusters, we see some obvious pairs of stocks grouped together: AT&T and Verizon; Lowes and Home Depot; Mastercard and Visa; McDonalds and Starbucks.

Overall, the group assignments in Table 4 look economically plausible, in addition to representing a much better statistical fit according to the BIC. In Section 5.4 we conduct formal out-of-sample forecast comparison tests to determine whether the improved in-sample fit leads to significantly better out-of-sample forecasts, and in Section 5.5 we study the economic environments in which this feature of the model is most helpful.

[INSERT TABLE 4 ABOUT HERE]

5.3 Estimated dependence time series

We now compare the fitted dependence time series from the two-digit SIC factor copula model and the factor copula model with estimated group assignments. We use rank correlations as a summary measure for the strength and direction of the dependence between assets implied by these models. With a fully-specified copula model such as the ones employed here, it is also possible to extract other dependence measures, such as tail dependence or probabilities of joint tail events, see e.g. the measures in Giesecke and Kim (2011) and Oh and Patton (2018).

The complete rank correlation matrix is 110×110 , and even just focusing on the blocks implied by the factor structure embedded in the model the matrix is 21×21 . As an initial summary measure, we firstly consider the conditional rank correlation for pairs in the same group. Figure 3 plots these for three groups, along with the two-digit SIC group that best matches the estimated group.^{15,16} The top panel compares estimated group 3 with SIC group 13. We observe that the two conditional rank correlation paths track each other quite closely, but the rank correlations based on estimated group assignments appear to adjust more quickly to news, and the SIC-based estimates look somewhat like a rolling average of the path from the model with estimated group assignments. A similar picture arises in the middle panel, comparing estimated group 7 with SIC group 36. It appears that by getting group assignments that better match the data, the model is more quickly able to react to information that suggests dependence has gone up or down.

The lower panel of Figure 3 compares estimated group 9 and SIC group 49, and represents a particularly interesting comparison. Group 9 contains six members, and all of them are from SIC group 49 (“Electric, Gas, & Sanitary Services,” in the “Transportation & Public Utilities” group). There is just one other SIC group 49 stock in our sample (Williams, ticker WMB), and this stock was estimated to belong to group 3, which is dominated by SIC group 13 members (SIC 13 is “Oil & Gas Extraction” in the “Mining” group). From the firm’s description on its website, it conducts a mix of activities captured by these SIC labels, and it turns out that our cluster assignment algorithm estimates it to be a better match with mining firms than with utilities firms. The lower panel of Figure 3 shows that by removing just this one stock the within-group rank correlation rises from around 0.55 to around 0.68. Moreover, we again see that the conditional rank correlations are more dynamic in the model with estimated group assignments.

[INSERT FIGURE 3 ABOUT HERE]

The plots of conditional rank correlations in Figure 3 allow us to see differences in pairwise

¹⁵Figures S1-S2 in the supplemental appendix present other comparisons of fitted rank correlations from the two models.

¹⁶For example, estimated Group 3 has eleven members, including all eight of the SIC group 13 stocks. Estimated Group 7 has seven members including all five members of SIC group 36. Estimated Group 9 has six members and all of them belong to SIC group 49; the single other SIC group 49 member was estimated to belong to Group 3.

dependence implied by the two models. For a more complete depiction of the differences implied by the model in the upper panel of Figure 4 we plot the QLIKE distance measure between the full 110×110 rank correlation matrices implied by the two models.¹⁷ When this measure is lower, the rank correlation matrices are more similar. We see that the difference is largest in mid 2011, and also large in late 2015, while it was relatively low in 2012. The middle panel of Figure 4 presents the normalized sum of the first 22 eigenvalues of the model-implied rank correlation matrices. Both of the models are based on a 22-factor model (one common factor and 21 group-specific factors), and the sum of the first 22 eigenvalues provides a summary for how informative the factors are.¹⁸ We see that the sum is uniformly greater for the model with estimated group assignments than for the model based on SIC group assignments. Note that the period when the two sums are furthest apart corresponds to the period when the QLIKE distance is also the greatest, indicating that this is one reason for the increased QLIKE distance. The lower panel of Figure 4 plots cross-sectional dispersion in pairwise rank correlations. We see that this dispersion has been broadly increasing over the sample period, and that periods when the two models differ most in the degree of dispersion also correspond to times when the QLIKE distance is larger.

[INSERT FIGURE 4 ABOUT HERE]

5.4 Out-of-sample forecast performance

We next compare the out-of-sample (OOS) forecasts of the factor copula models using SIC-based group assignments with those using estimated group assignments. To do so, we split our sample period in half, using data from 2010 to 2014 to estimate the models, and data from 2015 to 2019 to evaluate the models. Given the computational complexity of the models, we estimate the models only once, on the last day of the in-sample period, and retain those parameters throughout the OOS period.

¹⁷The QLIKE distance between two $(N \times N)$ matrices is $QLIKE(A, B) = \text{tr}(A^{-1}B) - \log|A^{-1}B| - N$.

¹⁸Figure S3 in the supplemental appendix presents corresponding results using just the largest eigenvalue, or the sum of the first three eigenvalues. The largest eigenvalues from each of the models are roughly equal, although similar to the pairwise rank correlation plots, the time series from the model with estimated group assignments appears more dynamic. The plot of the sum of the largest three eigenvalues reveals not only more dynamics, but a slight gap in the level, though it is not as large and not uniform as it is for the sum of the first 22 eigenvalues.

In Table 5 we use OOS forecast performance to determine the optimal shape of the copula (Gaussian, t , or skew t), as well as the optimal choice of dynamics (static vs. GAS).¹⁹ We do this for a range of choices for the number of groups, to determine the robustness of the conclusions, and also for the two SIC-based group assignments. In all cases we compare the models using their out-of-sample likelihoods, which is a consistent scoring rule for ranking density forecasts, see Gneiting and Raftery (2007). We test for the significance of the differences in OOS likelihoods using a Diebold and Mariano (1995) test with a Newey-West (1987) estimator of the standard error based on 10 lags.

The left panel of Table 5 clearly indicates that including GAS dynamics in the model improves the fit: in all cases the t -statistic is positive, and the smallest t -statistic across all configurations is 6.5, indicating strong evidence in favor of the GAS model over the static model. The right panel of Table 5 uses GAS dynamics in all cases, and we compare the choice of copula shape across various choices of the number of groups. We find in all cases that the t and skew t models outperform the Gaussian factor copula, with t -statistics all greater than 7.7. This is consistent with previous work in the literature (see, e.g., Patton, 2004, 2013, and Amengual and Sentana, 2020) that the Normal copula is not a good description of equity return dependence. In the last column of Table 5 we compare the t and skew t copulas, and we find that the t -statistics are all negative, and generally significant, indicating that the estimation of the additional skewness parameter in the skew t copula leads to worse OOS performance than the symmetric t factor copula. This is in contrast with the in-sample parameter estimates (presented in Tables S3 and S4 in the supplemental appendix) where the copula asymmetry parameter is significantly negative.²⁰ These conflicting results can be reconciled by the fact that OOS forecast comparisons tend to carry a strong implicit penalty for estimation error, and so unless the new parameter is far from zero and precisely estimated, better forecasts may be obtained by setting it to zero.

¹⁹In addition to being economically interesting in their own right, using OOS forecast performance to make these comparisons allows us to conduct formal statistical tests without having to make assumptions about the error rate in the estimated group assignments (see Section 3.2) that cannot be verified.

²⁰These two tables present parameter estimates and standard errors for models using one-digit SIC codes (seven clusters) or using estimated cluster assignments (using the BIC-optimal number of clusters, 21). In both cases we see that the copula asymmetry parameter is significantly negative. Standard errors in this table are computed assuming that estimation error from cluster assignments is negligible, as discussed in Section 3.2.

[INSERT TABLE 5 ABOUT HERE]

In Table 6 we compare the OOS performance of t factor copulas with GAS dynamics that use different numbers of groups. Consistent with the BIC rankings of models presented in Figure 2, the model based on one-digit SIC groupings is beaten by every other model except the estimated group assignment model with only 3 groups. Similarly, the model based on two-digit SIC groupings is beaten by every other model except the the one-digit SIC model and the estimated group assignment models with only 3 or 4 groups. Amongst the models with estimated group assignments, the model with $G = 20$ groups performs the best in terms of OOS likelihoods, significantly beating every other model, including the $G = 21$ model which was selected as being optimal over the full sample according to the BIC. That the optimal model for OOS forecasting is smaller than the optimal model for in-sample fitting is consistent with the abovementioned predilection of OOS forecasts for parsimonious models.

[INSERT TABLE 6 ABOUT HERE]

5.5 Economic determinants of forecast performance

We next seek to understand the economic environments in which specific features of the model lead to meaningful gains in forecast performance. We focus on three key model features: the model for *dynamics* in the conditional copula, the model for the *shape* of the copula, and the method for assigning stocks to *groups* in the factor structure. To summarize the economic environment we use three conditioning variables. Firstly, we use the Chicago Board of Exchange’s volatility index or “VIX” (see Whaley, 2009), which is a widely-used measure of market volatility. Next we use the cross-sectional standard deviation of stock returns on each day, which is a common measure of the degree of idiosyncratic risk (see Goyal and Santa-Clara, 2003, for example).²¹ Finally, we use the absolute value of the cross-sectional average of the difference between realized returns and the return predicted by the capital asset pricing model (CAPM), known as “alpha.”²² This measure

²¹We also considered the idiosyncratic risk relative to the CAPM for these stocks. That series had correlation of 0.99 with simple cross-sectional dispersion, and so we present results only using the latter.

²²We estimate the CAPM beta for each stock just once, using the full sample of data.

is interpretable as a measure of the degree of mispricing relative to the CAPM on a given day, or, more robustly, as a measure of how important non-market factors were for determining stock returns on a given day. These three measures each provide a different view of economic environment on a given day.

We use two complementary methods in this analysis. Firstly, we use the parametric approach of Giacomini and White (2005) (GW), where differences in out-of-sample forecast performance are regressed on one or more of the conditioning variables summarizing the economic environment. Secondly, we use the nonparametric “conditional superior predictive ability” (CSPA) test of Li *et al.* (2021). The GW test allows us to determine whether the conditioning variable(s) help explain the relative performance of the competing models, while the CSPA test determines whether one model has expected performance above/below the other model for *all* values of the conditioning variable. It is possible for neither, one, or both of these tests to reject their respective null hypotheses, and thus they provide complementary information about relative performance.

Table 7 presents the results for the three conditional model comparisons. In the GW tests, we consider the variables separately and jointly. We de-mean all regressors so that the intercept is interpretable as the expected difference in log-likelihoods when the regressors equal their average value. In all cases the intercepts are positive and strongly significant (consistent with Tables 5 and 6), indicating that in each comparison the more flexible model is preferred to the simpler model. Our interest in this analysis is primarily the slope coefficients, which affords us insight into whether there are specific economic environments where one model outperforms the other. The penultimate and antepenultimate rows show the p -values from joint tests of all coefficients in the regression, or only the slope coefficients.

In the left panel of Table 7 we compare a model with no dynamics in the conditional copula with a model that has GAS dynamics, and we see that the most useful variable is dispersion, which has a t -statistic of over three. The positive sign of this coefficient reveals that the gains from allowing for dynamics in the conditional copula are particularly great when stock returns exhibit higher idiosyncratic risk.

The middle panel of Table 7 presents results comparing a Gaussian copula with a Student’s t

copula, both with GAS dynamics. Here we see that VIX is the most useful explanatory variable, with a t -statistic of 2.8. The positive slope of this coefficient implies that when volatility is high the gains from allowing for joint fat tails, as provided by the Student’s t copula, are larger than in low-volatility times. We note also that dispersion and absolute alpha also have positive and significant coefficients in this panel, indicating that joint fat tails are also important when cross-sectional dispersion is high, and when non-market factors are particularly important for the cross-section of realized returns.²³

The right panel of Table 7 examines when the gains from using estimated group assignments rather than two-digit SIC groups are greatest. In this panel the most significant individual variable is absolute alpha, revealing that data-driven groupings are particularly valuable when non-market factors drive the cross-section of realized returns.

[INSERT TABLE 7 ABOUT HERE]

The bottom row of Table 7 presents the p -value from the nonparametric CSPA test for each of these comparisons.²⁴ We find p -values of less than 0.005 in all comparisons, consistent with (but generalizing) the strongly significant intercepts in the GW regressions, indicating that the more flexible model in each comparison dominates the more restrictive specification uniformly on the support of the conditioning variable.

To understand these nonparametric relationships better, Figure 5 presents simple nonparametric kernel-smoothed estimates of the differences in forecast performance as a function of each of the conditioning variables, for values between the 0.01 and 0.99 sample quantiles of each variable.²⁵ We see that the gains from allowing for dynamics in the conditional copula (top row) are roughly flat in VIX (consistent with the GW regression), while they are generally increasing in dispersion and absolute alpha. In particular, they increase almost linearly as dispersion increases from its first percentile (0.6) to about 1.8 (corresponding to its 88th percentile), and then flatten beyond that.

²³In the joint regression none of the three slope coefficients are individually significant, though they are strongly jointly significant (p -value less than 0.005), consistent with the presence of multicollinearity. The correlations between the regressors range from 0.21 to 0.50, indicating moderate multicollinearity.

²⁴Being a nonparametric test, the CSPA test suffers from the curse of dimensionality. For this reason, we implement it only for each variable separately, not for all variables jointly.

²⁵The estimate and confidence intervals are computed using Theorem 2.2 of Li and Racine (2007).

The gains from allowing for joint fat tails in the copula (second row of figures) are roughly flat when the three conditioning variables are in the lower half of their support, and then sharply increase beyond that, indicating that in times of high volatility, high idiosyncratic risk, or high importance of non-market factors, joint fat tails are particularly useful in the model. The gains from allowing for estimating group assignments (bottom row) are effectively unrelated to volatility, as measured by VIX, but they are strongly increasing in dispersion and absolute alpha. It is noteworthy that even when all three of these conditioning variables are at their first percentile, that is, when markets are calm, not disperse, and the market factor is dominant, the gains from allowing for dynamics, joint fat tails, and estimated group assignments are each significantly larger than zero. This is strong support for the importance of these features of the conditional copula.

[INSERT FIGURE 5 ABOUT HERE]

6 Conclusion

This paper proposes a new dynamic factor copula model for use in high dimensional time series applications. Our model does not require variables to be grouped according to *ex ante* information, like SIC industry codes or similar; instead we estimate the optimal assignment of variables to groups from the data using a *k*-means type approach. Our clustering method exploits the fact that clusters can be consistently estimated from a *static* version of the copula model, rather than the more computationally-challenging dynamic version that is of primary interest, making the clustering problem feasible. We show via an extensive simulation study that group assignments can be accurately estimated in finite samples. In an application to 110 U.S. equity returns over the period 2010-2019 we find evidence that a model with estimated group assignments significantly outperforms an otherwise identical model with group assignments determined using SIC codes. The improvement in fit appears to come from a better assignment of stocks that are labeled with one SIC code but comove more like stocks from a different SIC code, which allows the dynamic model to react to new information more quickly.

The methods in this paper suggest at least two directions for future research. First, one could

consider how to allow for time variation in estimated cluster assignments. For example, Lumsdaine *et al.* (2022) considers estimated group assignments that are subject to structural breaks, while Custodio João *et al.* (2022) model group assignments as evolving according to a hidden Markov model. Adapting these ideas to models of the form considered in this paper, which are already computationally intensive, is an interesting and challenging problem. A second direction is to consider extensions to “vast” dimensional data sets. Our analysis considers a cross-section of 110 stocks, which is large in absolute terms but small relative to our time series of around 2500 observations. Applications involving cross-sections comparable in size to the time series may require some new methods, e.g., adapting methods from random matrix theory (as in Fan *et al.*, 2013, for example) as well as alternative, faster, methods for estimating group assignments, e.g., hierarchical clustering (see, e.g., Hastie, *et al.*, 2009). We leave these interesting extensions for future work.

Appendix

A.1 The probability density of the skewed t copula

We adopt the skewed t copula discussed in Demarta and McNeil (2005) and Christoffersen *et al.* (2012). Specifically, it is the copula embedded in the multivariate skewed t distribution of \mathbf{X}_t , where:

$$\mathbf{X}_t = \sqrt{W_t} \mathbf{Z}_t + \zeta W_t \quad (16)$$

where ζ is a $N \times 1$ asymmetry parameter vector filled with an identical scalar ζ , W_t is an inverse gamma variable $W_t \sim IG\left(\frac{\nu}{2}, \frac{\nu}{2}\right)$, \mathbf{Z}_t is a $N \times 1$ normal variable $\mathbf{Z}_t \sim \mathcal{N}(0, \mathbf{R}_t)$, and W_t and \mathbf{Z}_t are independent. The probability density function of the skewed t copula is given by

$$c(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta) = \frac{2^{\frac{(\nu-2)(N-1)}{2}} K_{\frac{\nu+N}{2}} \left(\sqrt{(\nu + \mathbf{x}'_t \mathbf{R}_t^{-1} \mathbf{x}_t) \zeta' \mathbf{R}_t^{-1} \zeta} \right) \exp(\mathbf{x}'_t \mathbf{R}_t^{-1} \zeta)}{\ddot{\Gamma}\left(\frac{\nu}{2}\right)^{1-N} |\mathbf{R}_t|^{\frac{1}{2}} \left(\sqrt{(\nu + \mathbf{x}'_t \mathbf{R}_t^{-1} \mathbf{x}_t) \zeta' \mathbf{R}_t^{-1} \zeta} \right)^{-\frac{\nu+N}{2}} \left(1 + \frac{1}{\nu} \mathbf{x}'_t \mathbf{R}_t^{-1} \mathbf{x}_t\right)^{\frac{\nu+N}{2}}} \times \prod_{i=1}^N \frac{\left(\sqrt{(\nu + x_{it}^2) \zeta^2} \right)^{-\frac{\nu+1}{2}} \left(1 + \frac{x_{it}^2}{\nu}\right)^{\frac{\nu+1}{2}}}{K_{\frac{\nu+1}{2}} \left(\sqrt{(\nu + x_{it}^2) \zeta^2} \right) \exp(x_{it} \zeta)} \quad (17)$$

where $\ddot{\Gamma}$ is the Gamma function, $K(\cdot)$ is the modified Bessel function of the third kind (also called the modified Bessel function of the second kind, or the modified Hankel function), and $\mathbf{x}_t = [x_{1,t}, \dots, x_{N,t}]' = [T_{skew}^{-1}(u_{1,t}; \nu, \zeta), \dots, T_{skew}^{-1}(u_{N,t}; \nu, \zeta)]'$ are obtained by applying the inverse of the univariate skewed t distribution from equation (16) defined by

$$T_{skew}(y; \nu, \zeta) = \int_{-\infty}^y \frac{2^{1-0.5(\nu+1)} K_{\frac{\nu+1}{2}} \left(\sqrt{(\nu+x^2)\zeta^2} \right) \exp(x\zeta)}{\ddot{\Gamma}\left(\frac{\nu}{2}\right) \sqrt{\pi\nu} \left(\sqrt{(\nu+x^2)\zeta^2} \right)^{-\frac{\nu+1}{2}} \left(1 + \frac{x^2}{\nu}\right)^{\frac{\nu+1}{2}}} dx.$$

Since $T_{skew}^{-1}(\cdot; \nu, \zeta)$ is not available in closed form, we generate 1,000,000 random draws from equation (16) and use linear interpolation to approximate $T_{skew}^{-1}(\cdot; \nu, \zeta)$ on $(0, 1)$. Note that $T_{skew}(\cdot; \nu, \zeta)$ is identical across the cross sectional dimension and also over time because the shape parameters of this distribution are assumed constant, this means that we can approximate T_{skew}^{-1} just once per parameter set $[\nu, \zeta]$ and apply it to all copula inputs $\{\mathbf{u}_t\}_{t=1}^T$, making the likelihood evaluation of the skewed t copula very fast. An alternative to this simulation-based approach for finding an inverse CDF is to use quadrature-based methods, though in our initial analyses of such methods we found them too slow for use in the computationally-intensive estimation problem considered here.

A.2 Derivation of the score

From equation (17), the log-likelihood of the skewed t copula is obtained by

$$\begin{aligned} \log \mathbf{c}_{Skewt,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta) &= -\frac{1}{2} \log |\mathbf{R}_t| - \frac{\nu + N}{2} \log \left(1 + \frac{1}{\nu} \mathbf{x}_t' \mathbf{R}_t^{-1} \mathbf{x}_t \right) \\ &\quad + \log \left(K_{\frac{\nu+N}{2}} \left(\sqrt{(\nu + \mathbf{x}_t' \mathbf{R}_t^{-1} \mathbf{x}_t) \zeta' \mathbf{R}_t^{-1} \zeta} \right) \right) \\ &\quad + \mathbf{x}_t' \mathbf{R}_t^{-1} \zeta + \frac{\nu + N}{2} \log \left(\sqrt{(\nu + \mathbf{x}_t' \mathbf{R}_t^{-1} \mathbf{x}_t) \zeta' \mathbf{R}_t^{-1} \zeta} \right) + const(\nu, \zeta) \end{aligned} \quad (18)$$

where $const(\nu, \zeta)$ contains any components that do not depend on \mathbf{R}_t , and recall that

$$\mathbf{R}_t = \tilde{\mathbf{L}}_t' \tilde{\mathbf{L}}_t + \mathbf{D}_t, \quad \tilde{\mathbf{L}}_t = [\tilde{\lambda}_{1,t}, \dots, \tilde{\lambda}_{N,t}], \quad \mathbf{D}_t = diag(\sigma_{1,t}^2, \dots, \sigma_{N,t}^2)$$

where

$$\tilde{\boldsymbol{\lambda}}_{i,t} = \frac{\boldsymbol{\lambda}_{i,t}}{\sqrt{1 + \boldsymbol{\lambda}'_{i,t} \boldsymbol{\lambda}_{i,t}}}, \quad \sigma_{i,t}^2 = \frac{1}{1 + \boldsymbol{\lambda}'_{i,t} \boldsymbol{\lambda}_{i,t}}$$

To derive the score, we first define $\mathbf{L}_t = (\boldsymbol{\lambda}_{1,t}, \dots, \boldsymbol{\lambda}_{N,t}) \in \mathbb{R}^{k \times N}$ where $\boldsymbol{\lambda}_{i,t}$ is a $k \times 1$ vector of factor loadings for variable i . In the example of equation (7), the number of factors denoted by k is $G + 1$ and if the variable i belongs to Group 3, then $\boldsymbol{\lambda}_{i,t} = [\lambda_{M,3,t}, 0, 0, \lambda_{3,3,t}, 0, \dots, 0]'$. By the chain rule, the derivative of equation (18) with respect to $\boldsymbol{\eta}_t$ (given in equation 9) can be written as product of three factors

$$\underbrace{\frac{\partial \log \mathbf{c}_{Skewt,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \boldsymbol{\eta}'_t}}_{(1 \times 2G)} = \underbrace{\frac{\partial \log \mathbf{c}_{Skewt,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \text{vec}(\mathbf{R}_t)'}}_{(1 \times N^2)} \cdot \underbrace{\frac{\partial \text{vec}(\mathbf{R}_t)}{\partial \text{vec}(\mathbf{L}_t)'}}_{(N^2 \times (G+1)N)} \cdot \underbrace{\frac{\partial \text{vec}(\mathbf{L}_t)}{\partial \boldsymbol{\eta}'_t}}_{((G+1)N \times 2G)} \quad (19)$$

where $\text{vec}(\cdot)$ stacks the columns of the matrix on top of one another to form a vector. The first factor of equation (19) can be written as

$$\frac{\partial \log \mathbf{c}_{Skewt,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \text{vec}(\mathbf{R}_t)'} = \left(\text{vec} \left(\frac{\partial \log \mathbf{c}_{Skewt,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \mathbf{R}_t} \right) \right)'$$

so we focus on

$$\begin{aligned} \frac{\partial \log \mathbf{c}_{Skewt,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \mathbf{R}_t} &= -\frac{1}{2} \cdot \frac{\partial \log |\mathbf{R}_t|}{\partial \mathbf{R}_t} - \frac{\nu + N}{2} \frac{\partial \log \left(1 + \frac{1}{\nu} \mathbf{x}'_t \mathbf{R}_t^{-1} \mathbf{x}_t \right)}{\partial \mathbf{R}_t} \\ &\quad + \frac{\partial \log \left(K_{\frac{\nu+N}{2}} \left(\sqrt{(\nu + \mathbf{x}'_t \mathbf{R}_t^{-1} \mathbf{x}_t) \zeta' \mathbf{R}_t^{-1} \zeta} \right) \right)}{\partial \mathbf{R}_t} \\ &\quad + \frac{\partial \mathbf{x}'_t \mathbf{R}_t^{-1} \zeta}{\partial \mathbf{R}_t} + \frac{\nu + N}{2} \cdot \frac{\partial \log \left(\sqrt{(\nu + \mathbf{x}'_t \mathbf{R}_t^{-1} \mathbf{x}_t) \zeta' \mathbf{R}_t^{-1} \zeta} \right)}{\partial \mathbf{R}_t} \end{aligned}$$

Two useful formulas from the matrix differentials are

$$\frac{d(\mathbf{v}' \mathbf{M}^{-1} \mathbf{w})}{d\mathbf{M}} = -(\mathbf{M}^{-1})' \mathbf{v} \mathbf{w}' (\mathbf{M}^{-1})' \quad \text{and} \quad \frac{d \log |\mathbf{M}|}{d\mathbf{M}} = \mathbf{M}^{-1}$$

where \mathbf{M} is a symmetric non-singular matrix, and \mathbf{v} and \mathbf{w} are vectors conformable with \mathbf{M} . With

those formulas we calculate each component separately to obtain

$$\begin{aligned} \frac{\partial \log \mathbf{c}_{Skewt,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \mathbf{R}_t} &= -\frac{1}{2} \mathbf{R}_t^{-1} + \frac{\nu + N}{2} \frac{\mathbf{R}_t^{-1} \mathbf{x}_t \mathbf{x}_t' \mathbf{R}_t^{-1}}{(\nu + \mathbf{x}_t' \mathbf{R}_t^{-1} \mathbf{x}_t)} + \frac{K'_{\frac{\nu+N}{2}}(A_t)}{K_{\frac{\nu+N}{2}}(A_t)} \cdot \frac{B_t}{2A_t} \\ &\quad - \mathbf{R}_t^{-1} \mathbf{x}_t \zeta' \mathbf{R}_t^{-1} + \left(\frac{\nu + N}{4} \right) \frac{B_t}{A_t^2} \end{aligned} \quad (20)$$

where

$$\begin{aligned} A_t &= \sqrt{(\nu + \mathbf{x}_t' \mathbf{R}_t^{-1} \mathbf{x}_t) \zeta' \mathbf{R}_t^{-1} \zeta} \\ B_t &= -(\nu + \mathbf{x}_t' \mathbf{R}_t^{-1} \mathbf{x}_t) (\mathbf{R}_t^{-1} \zeta \zeta' \mathbf{R}_t^{-1}) - (\mathbf{R}_t^{-1} \mathbf{x}_t \mathbf{x}_t' \mathbf{R}_t^{-1}) (\zeta' \mathbf{R}_t^{-1} \zeta) \\ K'_{\frac{\nu+N}{2}}(\cdot) &= -\frac{1}{2} \left[K_{\frac{\nu+N}{2}-1}(\cdot) + K_{\frac{\nu+N}{2}+1}(\cdot) \right]. \end{aligned}$$

As $\zeta \rightarrow 0$, equation (20) boils down to the derivative of the log density of the Student t copula,

$$\frac{\partial \log \mathbf{c}_{Student-t,t}(\mathbf{x}_t; \mathbf{R}_t, \nu)}{\partial \mathbf{R}_t} = -\frac{1}{2} \mathbf{R}_t^{-1} + \frac{\nu + N}{2} \frac{\mathbf{R}_t^{-1} \mathbf{x}_t \mathbf{x}_t' \mathbf{R}_t^{-1}}{(\nu + \mathbf{x}_t' \mathbf{R}_t^{-1} \mathbf{x}_t)}$$

and in addition, as $\nu \rightarrow \infty$, it becomes that of the Gaussian copula,

$$\frac{\partial \log \mathbf{c}_{Gaussian,t}(\mathbf{x}_t; \mathbf{R}_t)}{\partial \mathbf{R}_t} = -\frac{1}{2} \mathbf{R}_t^{-1} + \frac{1}{2} \mathbf{R}_t^{-1} \mathbf{x}_t \mathbf{x}_t' \mathbf{R}_t^{-1}.$$

To derive the second and third factors in equation (19) we closely follow the logic and notations of Opschoor *et al.* (2020). The second factor is re-written as

$$\begin{aligned} \frac{\partial \text{vec}(\mathbf{R}_t)}{\partial \text{vec}(\mathbf{L}_t)'} &= \frac{\partial \text{vec}(\tilde{\mathbf{L}}_t' \tilde{\mathbf{L}}_t)}{\partial \text{vec}(\tilde{\mathbf{L}}_t)'} \cdot \frac{\partial \text{vec}(\tilde{\mathbf{L}}_t)}{\partial \text{vec}(\mathbf{L}_t)'} + \frac{\partial \text{vec}(\mathbf{D}_t)}{\partial \text{vec}(\mathbf{L}_t)'} \\ &= (\mathbf{I}_{N^2} + J_{N,N}) (\mathbf{I}_N \otimes \tilde{\mathbf{L}}_t') \cdot \frac{\partial \text{vec}(\tilde{\mathbf{L}}_t)}{\partial \text{vec}(\mathbf{L}_t)'} + \frac{\partial \text{vec}(\mathbf{D}_t)}{\partial \text{vec}(\mathbf{L}_t)'} \end{aligned} \quad (21)$$

with another useful formula

$$\frac{\partial \text{vec}(\mathbf{M}'\mathbf{M})}{\partial \text{vec}(\mathbf{M})'} = (\mathbf{I}_{n^2} + J_{n,n}) (\mathbf{I}_n \otimes \mathbf{M}')$$

where $\mathbf{M} \in \mathbb{R}^{m \times n}$, and $J_{m,n} \in \mathbb{R}^{mn \times mn}$ is the vectorized transpose matrix, i.e. $vec(\mathbf{M}') = J_{m,n}vec(\mathbf{M})$. Furthermore, we have

$$\frac{\partial vec(\tilde{\mathbf{L}}_t)}{\partial vec(\mathbf{L}_t)'} = \begin{pmatrix} \mathbf{Q}_{1,t} & 0 & \cdots & 0 \\ 0 & \mathbf{Q}_{2,t} & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{Q}_{N,t} \end{pmatrix}, \quad \mathbf{Q}_{i,t} = \frac{\mathbf{I}_k}{\sqrt{1 + \lambda'_{i,t}\lambda_{i,t}}} - \frac{\lambda_{i,t}\lambda'_{i,t}}{(1 + \lambda'_{i,t}\lambda_{i,t})^{3/2}}$$

for $i = 1, \dots, N$. (The denominator on the final term in the above equation corrects a typo in Opschoor et al., 2020). The sparsity of $(\mathbf{I}_N \otimes \tilde{\mathbf{L}}_t')$ and $\partial vec(\tilde{\mathbf{L}}_t) / \partial vec(\mathbf{L}_t)'$ simplifies product of those two factors to

$$(\mathbf{I}_N \otimes \tilde{\mathbf{L}}_t') \cdot \frac{\partial vec(\tilde{\mathbf{L}}_t)}{\partial vec(\mathbf{L}_t)'} = \begin{pmatrix} \tilde{\mathbf{L}}_t' \mathbf{Q}_{1,t} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \tilde{\mathbf{L}}_t' \mathbf{Q}_{N,t} \end{pmatrix}. \quad (22)$$

We define T_{diag} as a $N^2 \times N^2$ transformation matrix such that $vec(A) = T_{diag} \cdot a$ where A is a $N \times N$ diagonal matrix with a $N \times 1$ vector a on the diagonal. Then,

$$\frac{\partial vec(\mathbf{D}_t)}{\partial vec(\mathbf{L}_t)'} = T_{diag} \cdot \frac{\partial [\sigma_{1,t}^2, \dots, \sigma_{N,t}^2]'}{\partial vec(\mathbf{L}_t)'} = T_{diag} \cdot \begin{pmatrix} \frac{-2\lambda'_{1,t}}{(1 + \lambda'_{1,t}\lambda_{1,t})^2} & \cdots & 0 \\ 0 & \ddots & 0 \\ 0 & \cdots & \frac{-2\lambda'_{N,t}}{(1 + \lambda'_{N,t}\lambda_{N,t})^2} \end{pmatrix}. \quad (23)$$

For the last factor in equation (19), recall that $\boldsymbol{\eta}_t$ in equation (9) is a vector of distinct factor loadings and $\eta_{t,i}$ denotes i -th element of $\boldsymbol{\eta}_t$. \mathbf{L}_t is written as

$$\mathbf{L}_t = \sum_{i=1}^{2G} \eta_{t,i} \cdot S'_i$$

$$S_i = \begin{pmatrix} \delta_{1,i\iota_{N_1}} & \delta_{G+1,i\iota_{N_1}} & \mathbf{0} & \cdots & \mathbf{0} \\ \delta_{2,i\iota_{N_2}} & \mathbf{0} & \delta_{G+2,i\iota_{N_2}} & & \vdots \\ \vdots & \vdots & & \ddots & \mathbf{0} \\ \delta_{G,i\iota_{N_G}} & \mathbf{0} & \cdots & \mathbf{0} & \delta_{2G,i\iota_{N_G}} \end{pmatrix} \in \mathbb{R}^{N \times (G+1)}$$

where ι_p is a $p \times 1$ vector filled with ones, $\delta_{i,j} = 1$ if $i = j$ and zero otherwise, N_g for $g = 1, \dots, G$ is the number of members in group g such that $N = \sum_{g=1}^G N_g$. Then

$$\frac{\partial \text{vec}(\mathbf{L}_t)}{\partial \boldsymbol{\eta}'_t} = \left(\text{vec}(S'_1), \dots, \text{vec}(S'_G), \text{vec}(S'_{G+1}), \dots, \text{vec}(S'_{2G}) \right) \in \mathbb{R}^{(G+1)N \times 2G}. \quad (24)$$

Thus, the score expressed in equation (19) is obtained by combining equation (20), (21), (22), (23), and (24).

A.3 Variance targeting

The number of parameters to estimate in the proposed skewed t copula with GAS dynamics is $2G + 6$, and when G is large estimating all the parameters at once is not feasible. We adopt a two-step approach, so-called “variance targeting,” to eliminate the need to numerically optimize over the intercept parameters $[\omega_1^M, \dots, \omega_G^M, \omega_1^C, \dots, \omega_G^C]$. Specifically, under the stationarity assumption, the unconditional expectation of all distinct factor loadings in equation (8), $\boldsymbol{\eta}_t = [\lambda_{M,1,t}, \dots, \lambda_{M,G,t}, \lambda_{1,1,t}, \dots, \lambda_{G,G,t}]'$ is

$$\bar{\boldsymbol{\eta}} = \boldsymbol{\omega} + \mathbf{B} \cdot \bar{\boldsymbol{\eta}}$$

where $\bar{\boldsymbol{\eta}} \equiv \mathbb{E}[\boldsymbol{\eta}_t]$, $\boldsymbol{\omega} \equiv [\omega_1^M, \dots, \omega_G^M, \omega_1^C, \dots, \omega_G^C]'$ and $\mathbf{B} = \text{diag}(\beta^M, \dots, \beta^M, \beta^C, \dots, \beta^C) \in \mathbb{R}^{2G \times 2G}$, so if $\bar{\boldsymbol{\eta}}$ is estimated from data in a first step, $\boldsymbol{\omega}$ can be replaced with $(I_{2G} - \mathbf{B}) \hat{\bar{\boldsymbol{\eta}}}$ and only $[\alpha^M, \beta^M, \alpha^C, \beta^C, \nu, \zeta]$ are left to be estimated numerically in a second step.

Define a $(G + 1) \times G$ matrix $\bar{\mathbf{L}}$ filled with elements of $\bar{\boldsymbol{\eta}}$ as

$$\bar{\mathbf{L}}' = \begin{pmatrix} \frac{\mathbb{E}[\lambda_{M,1,t}]}{v_1} & \frac{\mathbb{E}[\lambda_{1,1,t}]}{v_1} & 0 & \cdots & 0 \\ \frac{\mathbb{E}[\lambda_{M,2,t}]}{v_2} & 0 & \frac{\mathbb{E}[\lambda_{2,2,t}]}{v_2} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\mathbb{E}[\lambda_{M,G,t}]}{v_G} & 0 & 0 & \cdots & \frac{\mathbb{E}[\lambda_{G,G,t}]}{v_G} \end{pmatrix}$$

where $v_i \equiv \sqrt{1 + \mathbb{E}[\lambda_{M,i,t}]^2 + \mathbb{E}[\lambda_{i,i,t}]^2}$, then the model implied correlation matrix of within- and across-group is obtained by $\bar{\mathbf{L}}'\bar{\mathbf{L}}$. The corresponding unconditional correlation matrix based on samples $\mathbf{x}_{it} = \Phi^{-1}(u_{it})$ is denoted by $\hat{\Omega} \in \mathbb{R}^{G \times G}$ where the g -th diagonal element is the average correlation of any pair of variables belonging to group g and a (i, j) element is the average correlation of any pair of variables belonging to group i and group j . Then $\bar{\boldsymbol{\eta}}$ is estimated by minimizing the difference between the sample ($\hat{\Omega}$) and model-implied correlation matrices:

$$\hat{\boldsymbol{\eta}} = \arg \min_{\bar{\boldsymbol{\eta}}} \left[\text{vech} \left(\hat{\Omega} - \bar{\mathbf{L}}'\bar{\mathbf{L}} \right) \right]' \text{vech} \left(\hat{\Omega} - \bar{\mathbf{L}}'\bar{\mathbf{L}} \right)$$

In most variance targeting applications, the estimation of the intercept can be done analytically. Here it requires numerical optimization, however it is extremely fast.

A.4 Proof of Theorem 1

Firstly, define the profile estimator

$$\tilde{\boldsymbol{\theta}}_T(\Gamma) = \arg \max_{\boldsymbol{\theta}} \frac{1}{T} \sum_{t=1}^T \log c(\mathbf{u}_t; \boldsymbol{\theta}, \Gamma) \quad (25)$$

Assumptions 1, 2(a) and 3(a) are sufficient for $\tilde{\boldsymbol{\theta}}_T(\Gamma) \xrightarrow{p} \tilde{\boldsymbol{\theta}}^*(\Gamma)$ for each $\Gamma \in \mathcal{G}$, see White (1994, Theorem 3.5) for example. Next define the sample and population profile likelihoods as:

$$\begin{aligned} \bar{Q}_T(\Gamma) &= \frac{1}{T} \sum_{t=1}^T \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}_T(\Gamma), \Gamma) \\ Q^*(\Gamma) &\equiv \mathbb{E} \left[\log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma) \right] \end{aligned}$$

Define the infeasible version of the sample likelihood using the population copula parameter as

$$\dot{Q}_T(\Gamma) \equiv \frac{1}{T} \sum_{t=1}^T \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma)$$

Consider a mean-value expansion of the sample objective function:

$$\begin{aligned} \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}_T(\Gamma), \Gamma) &= \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma) + \nabla_{\boldsymbol{\theta}} \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma)' (\tilde{\boldsymbol{\theta}}_T(\Gamma) - \tilde{\boldsymbol{\theta}}^*(\Gamma)) \\ &\quad + \frac{1}{2} (\tilde{\boldsymbol{\theta}}_T(\Gamma) - \tilde{\boldsymbol{\theta}}^*(\Gamma))' \nabla_{\boldsymbol{\theta}\boldsymbol{\theta}} \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma) (\tilde{\boldsymbol{\theta}}_T(\Gamma) - \tilde{\boldsymbol{\theta}}^*(\Gamma)) \\ \text{where } \tilde{\boldsymbol{\theta}}^* &= \lambda \tilde{\boldsymbol{\theta}}_T(\Gamma) + (1 - \lambda) \tilde{\boldsymbol{\theta}}^*(\Gamma) \text{ for some } \lambda \in [0, 1] \end{aligned} \quad (26)$$

Then summing the equation above over $t = 1, \dots, T$ we have

$$\begin{aligned} \bar{Q}_T(\Gamma) - \dot{Q}_T(\Gamma) &= \left(\frac{1}{T} \sum_{t=1}^T \nabla_{\boldsymbol{\theta}} \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma) \right)' (\tilde{\boldsymbol{\theta}}_T(\Gamma) - \tilde{\boldsymbol{\theta}}^*(\Gamma)) + \\ &\quad + \frac{1}{2} (\tilde{\boldsymbol{\theta}}_T(\Gamma) - \tilde{\boldsymbol{\theta}}^*(\Gamma))' \frac{1}{T} \sum_{t=1}^T \nabla_{\boldsymbol{\theta}\boldsymbol{\theta}} \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma) (\tilde{\boldsymbol{\theta}}_T(\Gamma) - \tilde{\boldsymbol{\theta}}^*(\Gamma)) \end{aligned} \quad (27)$$

Assumption 2(b) and 3(a) imply that $\frac{1}{T} \sum_{t=1}^T \nabla_{\boldsymbol{\theta}} \log \mathbf{c}(\mathbf{u}_t; \tilde{\boldsymbol{\theta}}^*(\Gamma), \Gamma) \xrightarrow{p} 0$, as usual for M -estimation. Assumption 3(c) ensures the Hessian term has a finite limit. Thus we have $\bar{Q}_T(\Gamma) - \dot{Q}_T(\Gamma) = o_p(1)$. Further, by Assumptions 1 and 2(a) we also have $\dot{Q}_T(\Gamma) - Q^*(\Gamma) = o_p(1)$, and so we have $\bar{Q}_T(\Gamma) - Q^*(\Gamma) = o_p(1)$. Thus the sample objective function is pointwise (in Γ) consistent for the population objective function. Since the parameter space is discrete, uniform convergence simplifies to pointwise convergence (see, e.g., Choirat and Seri, 2012). This implies the estimator obtained by maximizing the sample objective function is consistent for the parameter that maximizes the population objective function. Since the population objective function is unaffected by re-labeling of clusters, any element of \mathcal{G}_0 maximizes the population objective function. We thus have $\Pr \left[\hat{\Gamma}_T \in \mathcal{G}_0 \right] \xrightarrow{p} 1$ as $T \rightarrow \infty$, completing the proof.

References

- [1] Amengual D. and E. Sentana, 2020, Is a Normal Copula the Right Copula?, *Journal of Business & Economic Statistics*, 38(2), 350-366.
- [2] Ang, A. and J. Chen, 2002, Asymmetric Correlations of Equity Portfolios, *Journal of Financial Economics*, 63, 443-494
- [3] Bester, A. and C. Hansen, 2016, Grouped Effects Estimators in Fixed Effects Models, *Journal of Econometrics*, 190(1), 197-208.
- [4] Blasques, F., A. Lucas and S. J. Koopman, 2014, Stationarity and Ergodicity Conditions for Generalized Autoregressive Score Processes, *Electronic Journal of Statistics*, 8, 1088-1112.
- [5] Bollerslev, T., R. F. Engle, and D. B. Nelson, 1994, Chapter 49 ARCH Models, *Handbook of Econometrics*, Vol 4, Elsevier.
- [6] Bonhomme, S., and Manresa, E., 2015, Grouped patterns of heterogeneity in panel data, *Econometrica*, 83(3), 1147-1184.
- [7] Carassco, M. and X. Chen, 2002, Mixing and Moment Properties of Various GARCH and Stochastic Volatility Models, *Econometric Theory*, 18(1), 17-39.
- [8] Choirat, C. and R. Seri, 2012, Estimation in Discrete Parameter Models, *Statistical Science*, 27(2), 278-293.
- [9] Christoffersen, P., K. Jacobs, X. Jin and H. Langlois, 2018, Dynamic Dependence and Diversification in Corporate Credit, *Review of Finance*, 22(2), 521-560.
- [10] Creal, D.D. and R. Tsay, 2015, High-dimensional Dynamic Stochastic Copula Models, *Journal of Econometrics*, 189(2), 335-345.
- [11] Creal, D.D., S.J. Koopman, and A. Lucas, 2013, Generalized Autoregressive Score Models with Applications, *Journal of Applied Econometrics*, 28(5), 777-795.
- [12] Custodio João, I., A. Lucas, J. Schaumburg, and B. Schwaab, 2022, Dynamic Clustering of Multivariate Panel Data, *Journal of Econometrics*, forthcoming.
- [13] Demarta, S. and A. J. McNeil, 2005, The t Copula and Related Copulas. *International Statistical Review*, 73(1), 111-129.
- [14] Diebold, F.X. and R.S. Mariano, 1995, Comparing Predictive Accuracy, *Journal of Business & Economic Statistics*, 13(3), 253-263.
- [15] Engle R.F., 2002, Dynamic conditional correlation: A simple class of multivariate GARCH models, *Journal of Business & Economic Statistics*, 20, 339-350.
- [16] Fan, J., Y. Fan and J. Lv, 2008, High dimensional covariance matrix estimation using a factor model, *Journal of Econometrics*, 147, 186-197.
- [17] Fan, J., Y. Liao and M. Micheva, 2013, Large Covariance Estimation by Thresholding Principal Orthogonal Complements, *Journal of Royal Statistical Society Series B*, 75, 603-680.

- [18] Francis, N., M.T. Owyang, and Ö. Savascin, 2017, An endogenously clustered factor approach to international business cycles, *Journal of Applied Econometrics*, 32, 1261-1276.
- [19] Giacomini, R. and H. White, 2006, Tests of conditional predictive ability, *Econometrica*, 74, 1545-1578.
- [20] Giesecke, K. B. and Kim, 2011, Systemic Risk: What Defaults are Telling Us, *Management Science*, 57, 1387-1405.
- [21] Goyal, A. and P. Santa-Clara, 2003. Idiosyncratic risk matters! *Journal of Finance*, 58, 975-1008.
- [22] Gneiting, T. and A.E. Raftery, 2007, Strictly Proper Scoring Rules, Prediction and Estimation, *Journal of the American Statistical Association*, 102, 358-378.
- [23] Hafner, C.M. and A. Preminger, 2009, On Asymptotic Theory for Multivariate GARCH Models, *Journal of Multivariate Analysis*, 100, 2044-2054.
- [24] Hahn, J. and R. Moon, 2010, Panel Data Models with Finite Number of Equilibria, *Econometric Theory*, 26(3), 863-881.
- [25] Hansen, B. E., 1994, Autoregressive Conditional Density Estimation, *International Economic Review*, 35, 705-730.
- [26] Harvey, A.C., 2013, *Dynamic Models for Volatility and Heavy Tails*, Econometric Society Monograph 52, Cambridge University Press, Cambridge.
- [27] Hastie, T., R. Tibshirani and J. Friedman, 2009, *The Elements of Statistical Learning: Data Mining, Inference and Prediction, Second Edition*, Springer, New York.
- [28] Hautsch, N., L.M. Kyj and R.C.A. Oomen, 2012, A blocking and regularization approach to high dimensional realized covariance estimation, *Journal of Applied Econometrics*, 27, 625-645, 2012.
- [29] Hong, Y., J. Tu and G. Zhou, 2007, Asymmetries in Stock Returns: Statistical Tests and Economic Evaluation, *Review of Financial Studies*, 20, 1547-1581.
- [30] Li, J., Z. Liao and R. Quaedvlieg, 2021, Conditional superior predictive ability, *Review of Economic Studies*, forthcoming.
- [31] Li, Q. and J.S. Racine, 2007, *Nonparametric Econometrics*, Princeton University Press, Princeton.
- [32] Lin, C.-C. and S. Ng, 2012, Estimation of Panel Data Models with Parameter Heterogeneity when Group Membership is Unknown, *Journal of Econometric Methods*, 1(1), 42-55.
- [33] Lumsdaine, R.L. R. Okui and W. Wang, 2022, Estimation of panel group structure models with structural breaks in group memberships and coefficients, *Journal of Econometrics*, forthcoming.
- [34] Newey, W.K. and D. McFadden, 1993, Large Sample Estimation and Hypothesis Testing, in R.F. Engle and D. McFadden (eds.), *Handbook of Econometrics, Vol. IV*, Elsevier.

- [35] Newey, W.K. and K.D. West, 1987, A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55(3), 703-708.
- [36] Oh, D.H. and A.J. Patton, 2017, Modelling Dependence in High Dimensions with Factor Copulas, *Journal of Business & Economic Statistics*, 35(1), 139-154.
- [37] Oh, D.H. and A.J. Patton, 2018, Time-Varying Systemic Risk: Evidence from a Dynamic Copula Model of CDS Spreads, *Journal of Business & Economic Statistics*, 36(2), 181-195.
- [38] Opschoor, A., A. Lucas, I. Barra & D. van Dijk, 2020, Closed-form multi-factor copula models with observation-driven dynamic factor loadings, *Journal of Business & Economic Statistics*, forthcoming.
- [39] Patton, A.J., 2004, On the Out-of-Sample Importance of Skewness and Asymmetric Dependence for Asset Allocation, 2004, *Journal of Financial Econometrics*, 2(1), 130-168.
- [40] Patton, A.J., 2013, Copula Methods for Forecasting Multivariate Time Series, in G. Elliott and A. Timmermann (eds.) *Handbook of Economic Forecasting, Vol 2*, Springer Verlag.
- [41] Patton, A.J. and B.M. Weller, 2019, Risk price variation: The missing half of empirical asset pricing, working paper, Duke University.
- [42] Su, L., Z. Shi, and P.C.B. Phillips, 2016, Identifying Latent Structures in Panel Data, *Econometrica*, 6(84), 2215-2264.
- [43] Su, L. X. Wang and S. Jin, 2019, Sieve Estimation of Time-Varying Panel Data Models With Latent Structures, *Journal of Business & Economic Statistics*, 37(2), 334-349.
- [44] Tao, M., Y. Wang, Q. Yao and J. Zou, 2011, Large Volatility Matrix Inference via Combining Low-Frequency and High-Frequency Approaches, *Journal of the American Statistical Association*, 106(495), 1025-1040.
- [45] Vogt, M. and O. Linton, 2020, Multiscale Clustering of Nonparametric Regression Curves, *Journal of Econometrics*, 216(1), 305-325.
- [46] Whaley, R.E., 2009, Understanding the VIX, *Journal of Portfolio Management*, 35(3), 98-105.
- [47] White, H., 1994, *Estimation, Inference and Specification Analysis*, Econometric Society Monographs No. 22, Cambridge University Press, USA.

Table 1: Simulation results for static copulas

Panel A: Parameter estimation accuracy							
		Gaussian		t		skew t	
	True	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
β_1^M	0.25	0.245	0.076	0.253	0.089	0.276	0.062
β_2^M	0.50	0.508	0.071	0.509	0.075	0.473	0.068
β_3^M	0.75	0.755	0.053	0.744	0.056	0.740	0.066
β_4^M	1.00	1.000	0.052	1.002	0.048	0.954	0.189
β_5^M	1.25	1.246	0.033	1.243	0.035	1.246	0.040
β_6^M	1.50	1.500	0.021	1.498	0.034	1.501	0.035
β_7^M	1.75	1.747	0.020	1.752	0.027	1.755	0.036
β_8^M	2.00	2.001	0.021	1.998	0.027	1.996	0.030
β_9^M	2.25	2.253	0.019	2.252	0.028	2.255	0.031
β_{10}^M	2.50	2.496	0.020	2.502	0.029	2.502	0.030
β_1^C	2.50	2.499	0.026	2.495	0.030	2.503	0.031
β_2^C	2.25	2.248	0.023	2.246	0.031	2.251	0.036
β_3^C	2.00	2.001	0.031	2.003	0.030	1.998	0.038
β_4^C	1.75	1.749	0.030	1.747	0.034	1.771	0.065
β_5^C	1.50	1.497	0.030	1.501	0.029	1.497	0.031
β_6^C	1.25	1.246	0.028	1.252	0.026	1.252	0.030
β_7^C	1.00	1.005	0.023	0.999	0.023	1.003	0.019
β_8^C	0.75	0.756	0.022	0.749	0.023	0.743	0.018
β_9^C	0.50	0.504	0.019	0.499	0.023	0.499	0.026
β_{10}^C	0.25	0.241	0.039	0.234	0.048	0.230	0.050
ν	5.00			5.007	0.069	4.858	0.369
ζ	-0.10					-0.095	0.015

Panel B: Group assignment estimation accuracy			
Number incorrect			
0	100	100	100
≥ 1	0	0	0

Panel C: Estimation details						
	Clustering	Copula	Clustering	Copula	Clustering	Copula
Time (min)	8.7	0.47	8.9	6.37	8.8	7.13
EM (iter)	79.54	–	80.12	–	80.56	–

Notes: This table presents results from 100 simulations from static Gaussian, t , and skew t factor copulas with 10 groups. Panel A presents results on estimation accuracy of the copula parameters, Panel B presents results on estimation accuracy of the group assignments, and Panel C presents average estimation time (for the two stages of estimation) and EM iterations using a machine with an Intel Xeon processor, with ten cores and clock speed of 2.60GHz.

Table 2: Simulation results for time-varying copulas

Panel A: Parameter estimation accuracy							
		Gaussian		t		skew t	
	True	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
ω_1^M	0.04	0.042	0.007	0.042	0.007	0.044	0.008
ω_2^M	0.04	0.042	0.007	0.042	0.007	0.044	0.008
ω_3^M	0.04	0.042	0.007	0.042	0.007	0.043	0.007
ω_4^M	0.04	0.042	0.007	0.042	0.007	0.044	0.008
ω_5^M	0.04	0.042	0.007	0.042	0.007	0.043	0.007
ω_6^M	0.04	0.042	0.007	0.042	0.007	0.044	0.008
ω_7^M	0.04	0.042	0.006	0.042	0.007	0.043	0.008
ω_8^M	0.04	0.041	0.007	0.041	0.007	0.044	0.008
ω_9^M	0.04	0.042	0.007	0.042	0.007	0.044	0.008
ω_{10}^M	0.04	0.042	0.007	0.042	0.007	0.044	0.008
ω_1^C	0.04	0.043	0.007	0.043	0.007	0.042	0.007
ω_2^C	0.04	0.043	0.007	0.043	0.008	0.042	0.007
ω_3^C	0.04	0.042	0.007	0.043	0.007	0.042	0.007
ω_4^C	0.04	0.043	0.007	0.044	0.008	0.042	0.007
ω_5^C	0.04	0.043	0.007	0.043	0.008	0.041	0.007
ω_6^C	0.04	0.043	0.007	0.043	0.007	0.042	0.007
ω_7^C	0.04	0.043	0.008	0.043	0.007	0.041	0.006
ω_8^C	0.04	0.043	0.007	0.043	0.008	0.042	0.006
ω_9^C	0.04	0.043	0.008	0.043	0.007	0.042	0.007
ω_{10}^C	0.04	0.043	0.007	0.043	0.006	0.041	0.007
α^M	0.02	0.020	0.002	0.020	0.002	0.020	0.002
β^M	0.90	0.894	0.015	0.894	0.014	0.893	0.017
α^C	0.02	0.020	0.002	0.020	0.002	0.020	0.002
β^C	0.90	0.896	0.016	0.894	0.016	0.898	0.014
ν	5.00			5.014	0.071	5.016	0.108
ζ	-0.10					-0.100	0.007

Panel B: Group assignment estimation accuracy

Number incorrect			
0	100	100	100
≥ 1	0	0	0

Panel C: Estimation details

	Clustering	Copula	Clustering	Copula	Clustering	Copula
Time (min)	8.8	21.6	9.1	37.4	9.1	41.7
EM (iter)	91.32	-	91.78	-	90.83	-

Notes: This table presents results from 100 simulations from Gaussian, t , and skew t factor copulas with 10 groups and GAS dynamics. Panel A presents results on estimation accuracy of the copula parameters, Panel B presents results on estimation accuracy of the group assignments, and Panel C presents average estimation time (for the two stages of estimation) and EM iterations using a machine with an Intel Xeon

processor, with ten cores and clock speed of 2.60GHz.

Table 3: Summary statistics

	Cross-sectional distribution					
	Mean	5%	25%	Median	75%	95%
Panel A: Marginal moments						
Mean	0.001	0.000	0.001	0.001	0.001	0.001
Std	0.016	0.010	0.012	0.015	0.018	0.023
Skewness	-0.081	-0.748	-0.310	-0.091	0.092	0.648
Kurtosis	9.939	5.154	6.411	8.087	10.923	22.803
Panel B: Marginal model parameters						
Constant	0.001	0.000	0.001	0.001	0.001	0.001
AR(1)	-0.019	-0.068	-0.041	-0.017	0.000	0.031
$\varpi \times 10^4$	0.009	0.002	0.003	0.006	0.011	0.025
α	0.025	0.000	0.009	0.019	0.033	0.077
κ	0.099	0.029	0.064	0.095	0.131	0.179
β	0.885	0.756	0.864	0.904	0.932	0.958
ξ	5.089	3.401	4.234	4.846	5.798	7.256
ψ	-0.027	-0.087	-0.051	-0.025	-0.004	0.020
Panel C: Correlations of standardized residuals						
Pearson	0.322	0.170	0.256	0.314	0.378	0.492
Spearman	0.360	0.197	0.295	0.356	0.418	0.531

Notes: This table presents summary statistics on the 110 daily equity return series used in this paper. The sample period is January 2010 to December 2019. Panel A presents a summary of the cross-sectional distribution of the first four moments of these returns, Panel B presents a summary of the estimated AR(1)-GJR GARCH(1,1)-skew t model used for the marginal distributions, and Panel C presents a summary of the 5,995 pairwise correlations of the standardized residuals.

Table 4: Estimated group assignments

Group	Ticker	Name	SIC	Group	Ticker	Name	SIC	
1	ABT	Abbott Lab.	28	7	CSCO	Cisco Sys	36	
	AGN	Actavis	28		HPQ	Hewlett Pac	35	
	AMGN	Amgen	28		INTC	Intel	36	
	BAX	Baxter	38		MSFT	Microsoft	73	
	BIIB	Biogen	28		NVDA	Nvidia	36	
	BMY	Bristol-Myers	28		QCOM	Qualcomm	36	
	GILD	Gilead	28		TXN	Texas Instru	36	
	JNJ	Johnson & J	28	8	AIG	Ame Inter Group	63	
	LLY	Lilly Eli	28		ALL	Allstate	63	
	MDT	Medtronic	38		CMCSA	Comcast	48	
	MRK	Merck	28		DIS	Disney Walt	48	
	PFE	Pfizer	28		F	Ford	37	
	UNH	Unitedhealth	63		GE	Gen Electric	35	
					XRX	Xerox	35	
2	BAC	Bank Of Am	60	9	AEP	Ame Elec Pow	49	
	BK	Bank Of NY	60		DUK	Duke Energy	49	
	C	Citigroup Inc	60		ETR	Entergy Corp	49	
	COF	Capital One	60		EXC	Exelon	49	
	GS	Goldman Sachs	62		NEE	Nextera Energy	49	
	JPM	Jpmorgan	60		SO	Southern Co	49	
	MET	Metlife	63		10	COST	Costco	53
	MS	Morgan Stanley	60	CVS		C V S Health	59	
	RF	Regions Fin	60	TGT		Target	53	
	USB	U S Bancorp	60	WBA	Walgreens	59		
WFC	Wells Fargo	60	WMT	Walmart	53			
3	APA	Apache	13	11	GD	Gen Dynamics	37	
	BHI	Baker Hughes	35		LMT	Lockheed Martin	37	
	COP	Conocophillips	13		RTN	Raytheon	38	
	CVX	Chevron	13	12	AMT	American Tower	48	
	DVN	Devon	13		SPG	Simon Property	67	
	HAL	Halliburton	13		WY	Weyerhaeuser	8	
	NOV	Nat. Oilwell	35		13	BA	Boeing	37
	OXY	Occidental	13			FCX	Freeport Mcmo	10
	SLB	Schlumberger	13	NKE		Nike	30	
	WMB	Williams Co	49	14	ACN	Accenture	67	
	XOM	Exxon Mobil	13		IBM	IBM	35	
			ORCL		Oracle	73		
4	CAT	Caterpillar	35	15	AXP	Amex	60	
	EMR	Emerson Ele	35		BLK	Blackrock	62	
	FDX	Fedex	45	16	DHR	Danaher	38	
	HON	Honeywell Int	37		TMO	Thermo Fisher	38	
	MMM	3M	38		17	T	A T & T	48
	NSC	Norfolk South	40			VZ	Verizon	48
	UNP	Union Pacific	40			18	AVP	Avon Products
UPS	United Parcel	42	SNS	Steak N Shake	58			
5	AAPL	Apple	35	19	MA	Mastercard	73	
	ADBE	Adobe	73		V	Visa	61	
	AMZN	Amazon	73	20	MCD	Mcdonalds	58	
	CRM	Salesforce	73		SBUX	Starbucks	58	
	EBAY	Ebay	73		21	HD	Home Depot	52
	GOOGL	Google	73	LOW		Lowes	52	
	NFLX	Netflix	78					
	PCLN	Priceline	73					
	6	CL	Colgate Palmo	28				
CPB		Campbell Soup	20					
KO		Coca Cola	20					
MDLZ		Mondelez	20					
MO		Altria	21					
PEP		Pepsico	20					
PG		Procter Gamble	28					
PM		Philip Morris	21					

Notes: This table presets the estimated group assignments based on the BIC-optimal number of groups, $G=21$. The groups are ordered by the number of members.

Table 5: Comparing different copula specifications

	Static vs. GAS			Copula shape		
	Gaussian	t	skew t	G vs. t	G vs. skew t	t vs. skew t
SIC 1 digit	7.861	12.067	11.552	9.288	8.596	-2.975
SIC 2 digit	9.887	15.730	16.501	8.938	8.465	-2.849
3 groups	6.528	6.636	6.761	8.339	7.711	-2.257
4 groups	7.681	10.059	9.412	9.121	8.381	-2.351
5 groups	7.548	10.804	10.913	9.236	9.088	-1.717
18 groups	10.571	16.052	15.711	9.295	7.945	-3.550
19 groups	9.806	15.457	15.259	9.553	8.847	-2.065
20 groups	10.908	16.193	14.741	9.426	7.995	-3.797
21 groups	10.916	16.626	17.321	9.474	8.670	-2.366
22 groups	10.732	16.891	16.802	9.688	8.962	-2.714
25 groups	11.817	19.001	19.140	9.475	8.569	-2.355
27 groups	10.725	15.836	15.794	9.591	8.951	-1.676
30 groups	10.917	17.169	15.917	9.697	8.717	-2.962

Notes: This table presents Diebold-Mariano t -statistics on pairwise comparisons of models using their out-of-sample log-likelihood. The left panel compares models assuming no dynamics with those using GAS dynamics, for three different copula shapes (Gaussian, t , and skew t) and for a variety of choices for the number of groups. The right panel compares the different copula shapes, using GAS dynamics in all cases, across a variety of choices for the number of groups. In a comparison labeled “A vs B,” a positive t -statistic indicates that B is preferred; a negative t -statistic indicates that A is preferred. Note that there are 7 groups of firms using the 1-digit SIC, and 21 groups using the 2-digit SIC.

Table 6: Comparing different numbers of clusters

	SIC-1	SIC-2	3 groups	4 groups	5 groups	18 groups	19 groups	20 groups	21 groups	22 groups	30 groups
Panel A: t-statistics from pair-wise comparisons											
SIC-1		26.174	-4.783	14.180	21.898	33.495	32.074	33.127	33.292	32.973	29.975
SIC-2	-26.174		-21.579	-8.472	2.530	23.186	21.202	23.288	23.359	21.255	17.075
3 groups	4.783	21.579		18.347	24.175	31.398	30.848	30.926	31.247	31.529	28.722
4 groups	-14.180	8.472	-18.347		12.922	27.064	25.292	26.950	26.957	26.511	21.562
5 groups	-21.898	-2.530	-24.175	-12.922		23.175	20.562	22.535	22.689	20.069	15.187
18 groups	-33.495	-23.186	-31.398	-27.064	-23.175		-6.779	2.377	0.184	-6.754	-12.345
19 groups	-32.074	-21.202	-30.848	-25.292	-20.562	6.779		7.596	6.545	-3.270	-9.828
20 groups	-33.127	-23.288	-30.926	-26.950	-22.535	-2.377	-7.596		-2.915	-7.892	-13.062
21 groups	-33.292	-23.359	-31.247	-26.957	-22.689	-0.184	-6.545	2.915		-7.315	-12.798
22 groups	-32.973	-21.255	-31.529	-26.511	-20.069	6.754	3.270	7.892	7.315		-7.736
30 groups	-29.975	-17.075	-28.722	-21.562	-15.187	12.345	9.828	13.062	12.798	7.736	

Panel B: Out-of-sample log-likelihood values

$\log \mathcal{L}$	34074.5	37887.6	33175.8	36353.3	38303.6	42041.2	41564.0	42145.5	42050.1	41276.9	40559.2
--------------------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------

Notes: This table presents Diebold-Mariano t -statistics on pairwise comparisons of models using their out-of-sample log-likelihood. In all cases we use a t copula with GAS dynamics. A positive t -statistic indicates that the column model is preferred to the row model; a negative t -statistic indicates the opposite. Note that there are 7 groups of firms using the 1-digit SIC, and 21 groups using the 2-digit SIC.

Table 7: Economic determinants of forecast performance

	Static vs. GAS dynamics		Gaussian vs. Student's t shape		SIC vs. Estimated clusters	
<i>Intercept</i>	0.511	0.511	1.493	1.493	4.114	4.114
(s.e.)	(0.066)	(0.066)	(0.145)	(0.147)	(0.171)	(0.17)
[t-stat]	[7.741]	[7.791]	[10.316]	[10.142]	[24.133]	[24.218]
<i>VIX</i>	0.023	0.007	0.106	0.063	0.024	-0.061
(s.e.)	(0.018)	(0.020)	(0.038)	(0.020)	(0.039)	(0.020)
[t-stat]	[1.265]	[0.357]	[2.820]	[1.217]	[0.609]	[-1.300]
<i>Dispersion</i>						
(s.e.)	0.474	0.393	1.354	0.760	2.243	2.146
[t-stat]	(0.157)	(0.186)	(0.578)	(0.020)	(0.565)	(0.020)
	[3.018]	[2.108]	[2.342]	[1.073]	[3.969]	[3.232]
<i>Abs. alpha</i>						
(s.e.)		1.312	0.479	5.256	6.098	2.333
[t-stat]		(0.498)	(0.605)	(2.144)	(1.395)	(0.020)
		[2.635]	[0.793]	[2.451]	[4.371]	[1.462]
R^2 (%)	0.300	1.400	1.095	1.823	0.038	3.733
GW p -value _{ALL}	0.000	0.000	0.000	0.000	0.000	0.000
GW p -value _{SLOPES}	0.206	0.003	0.005	0.014	0.543	0.000
CSPA p -value	0.000	0.000	0.000	0.000	0.000	0.000

Notes: This table presents the results of Giacomini-White (2005) tests and Li et al. (2021) conditional superior predictive ability (CSPA) tests, comparing three pairs of models. The baseline model uses a Student's t copula with GAS dynamics and estimated assignments into 20 clusters, and is compared to a model that: imposes a static conditional copula (left panel); uses a Gaussian copula (middle panel); uses two-digit SIC codes to form clusters (right panel). The bottom row presents the CSPA p -value for the null hypothesis that the simpler model in each comparison is uniformly weakly better against the alternative that the more flexible model is uniformly strictly better. The rows labeled "GW p -value" test that all parameters in the regression, or only the slope coefficients, are equal to zero. The conditioning variables are the VIX index, the cross-sectional standard deviation of returns ("dispersion"), and the absolute value of the cross-sectional average CAPM alpha.

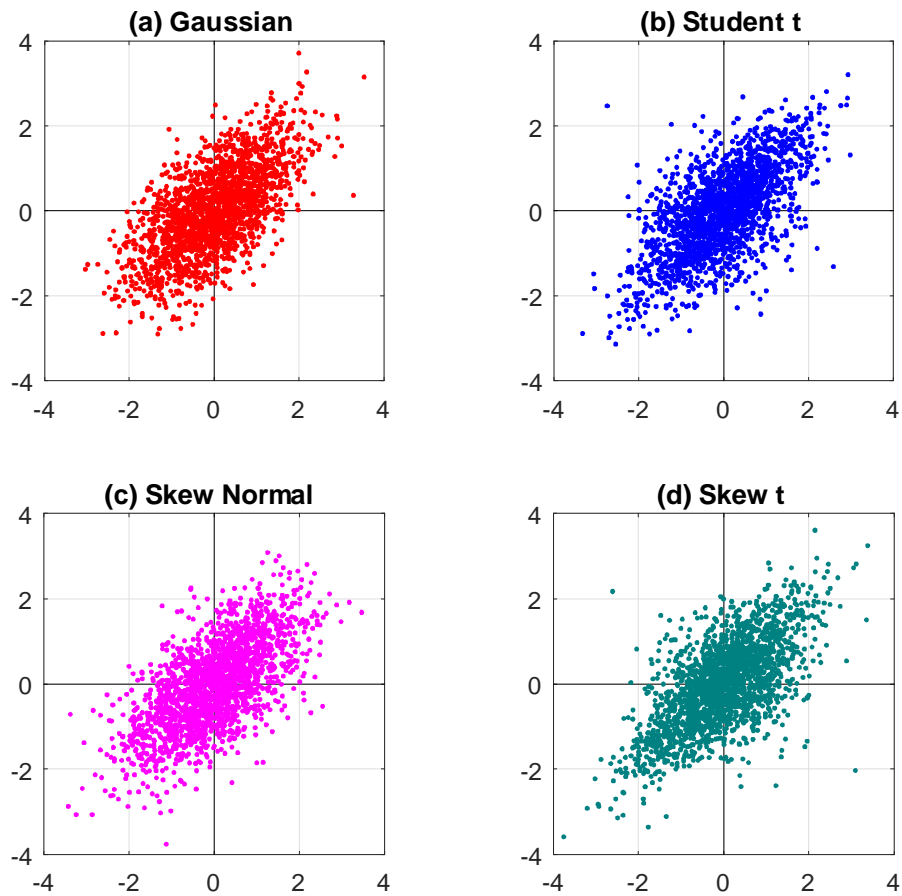


Figure 1: *This figure presents random draws from four joint distributions, all with standard Normal margins. Panel (a) uses a Gaussian copula, Panel (b) uses a Student's t copula, Panel (c) uses a skew Normal copula, Panel (d) uses a skew t copula. For all four copulas the correlation parameter is set to 0.5. For both t copulas the degrees of freedom parameter is set to 5. For both skewed copulas the skewness parameter is set to -0.1.*




Figure 2: Plot of BIC value as a function of the number of groups (G) for the EM-estimated model. The BIC values for the 1-digit and 2-digit SIC-based groups are also reported for comparison; these models have 7 and 21 groups respectively. As usual, lower BIC values are preferred. (Note the y -axis has been scaled by 10^{-4} for ease of presentation.)




Figure 3: *Time series plots of model-implied within-group rank correlations. The upper panel presents estimated group 3 and SIC group 13; the middle panel presents estimated group 7 and SIC group 36; the lower panel presents estimated group 9 and SIC group 49.*

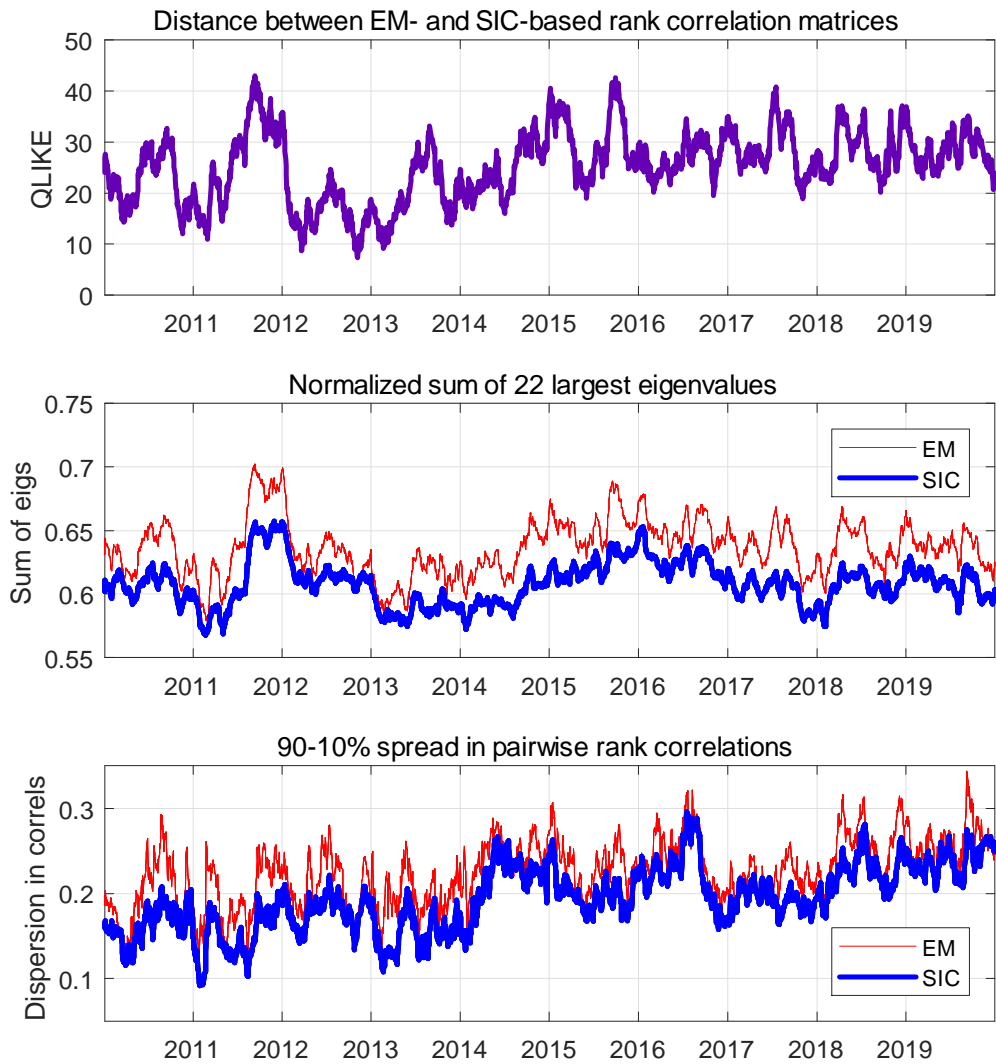


Figure 4: *The upper panel presents the QLIKE distance between the conditional rank correlation matrices implied by the 2-digit SIC-based model and the optimal EM-based factor copula model, both of which have a total of 22 factors. The middle panel presents the sum of the 22 largest eigenvalues of the conditional rank correlation matrices, divided by 110, the number of assets. The lower panel presents the difference between the 90% and 10% cross-sectional quantiles of all 5,995 pairwise rank correlations.*

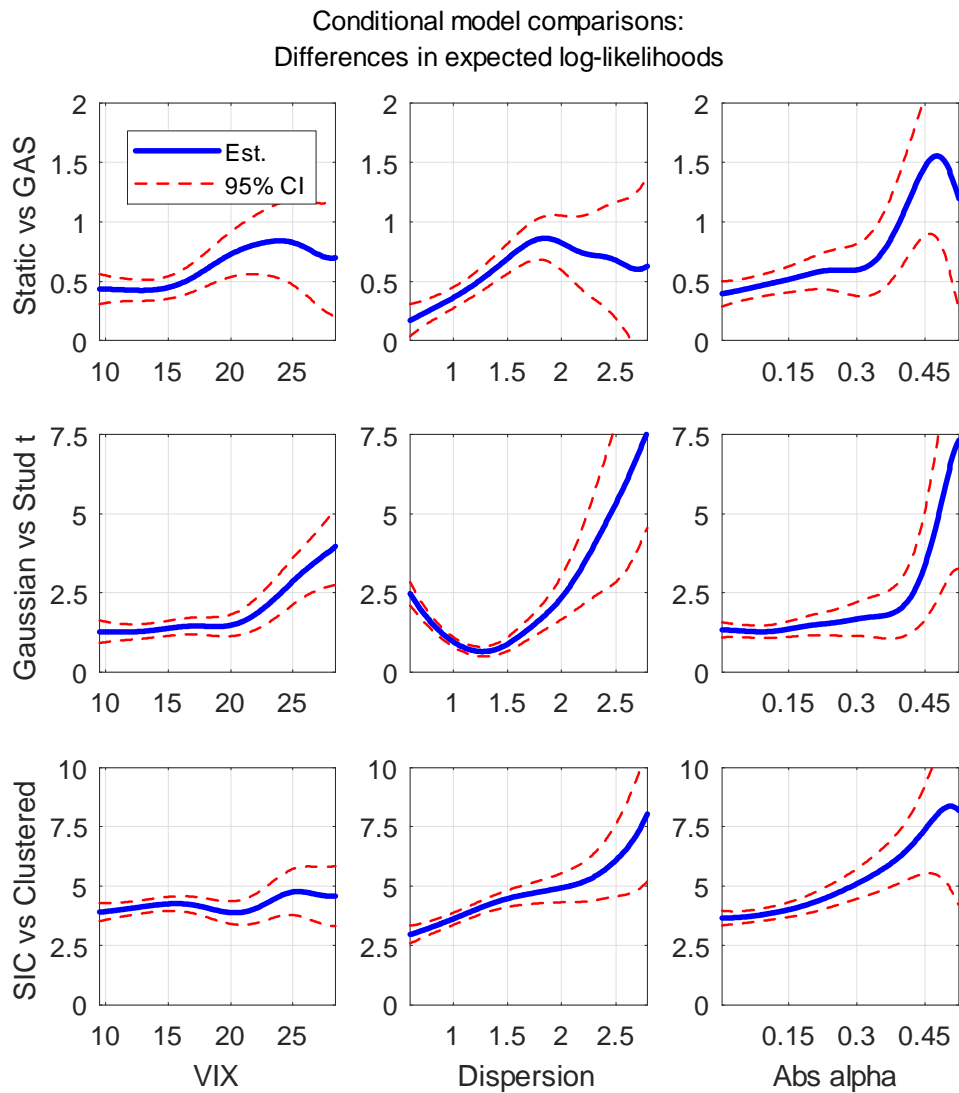


Figure 5: *This figure presents estimates of the expected difference in out-of-sample log likelihoods between the models listed in the y-axis label, conditioning on the variable given in the x-axis label. Positive values indicate that the second model in the comparison is preferred. Also presented are pointwise 95% confidence intervals for the estimated difference.*