

Closed-Form Multi-Factor Copula Models With Observation-Driven Dynamic Factor Loadings

Anne Opschoor, André Lucas, István Barra & Dick van Dijk

To cite this article: Anne Opschoor, André Lucas, István Barra & Dick van Dijk (2020): Closed-Form Multi-Factor Copula Models With Observation-Driven Dynamic Factor Loadings, Journal of Business & Economic Statistics, DOI: [10.1080/07350015.2020.1763806](https://doi.org/10.1080/07350015.2020.1763806)

To link to this article: <https://doi.org/10.1080/07350015.2020.1763806>



© 2020 The Author(s). Published with license by Taylor & Francis Group, LLC



[View supplementary material](#)



Published online: 01 Jun 2020.



[Submit your article to this journal](#)



Article views: 758



[View related articles](#)



[View Crossmark data](#)

Closed-Form Multi-Factor Copula Models With Observation-Driven Dynamic Factor Loadings

Anne Opschoor^{a,b}, André Lucas^{b,c}, István Barra^{a,b}, and Dick van Dijk^{b,d}

^aDepartment of Finance, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands; ^bTinbergen Institute, Amsterdam, The Netherlands; ^cDepartment of Econometrics, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands; ^dDepartment of Econometrics, Erasmus University Rotterdam, Rotterdam, The Netherlands

ABSTRACT

We develop new multi-factor dynamic copula models with time-varying factor loadings and observation-driven dynamics. The new models are highly flexible, scalable to high dimensions, and ensure positivity of covariance and correlation matrices. A closed-form likelihood expression allows for straightforward parameter estimation and likelihood inference. We apply the new model to a large panel of 100 U.S. stocks over the period 2001–2014. The proposed multi-factor structure is much better than existing (single-factor) models at describing stock return dependence dynamics in high-dimensions. The new factor models also improve one-step-ahead copula density forecasts and global minimum variance portfolio performance. Finally, we investigate different mechanisms to allocate firms into groups and find that a simple industry classification outperforms alternatives based on observable risk factors, such as size, value, or momentum.

ARTICLE HISTORY

Received February 2019
Accepted April 2020

KEYWORDS

Factor copulas; Factor structure; Multivariate density forecast; Score-driven dynamics

1. Introduction

Copulas are a key ingredient in many current applications in economics and finance (see, e.g., Patton 2009; Cherubini et al. 2011; Fan and Patton 2014; McNeil, Frey, and Embrechts 2015). In particular, time-varying copulas have turned out to be an important and flexible tool to describe dependence dynamics in an unstable environment (see Patton 2006; Manner and Reznikova 2012; Lucas, Schwaab, and Zhang 2014). Most copula applications deal with a cross-sectional dimension that is small to moderate (for an overview, see Patton 2013). Applications to high-dimensional datasets are scarce, mainly due to the “curse of dimensionality”: the number of parameters grows rapidly when the dimension increases.



Recently, Creal and Tsay (2015), Oh and Patton (2017, 2018), and Lucas, Schwaab, and Zhang (2017) put forward a general approach to modeling time-varying dependence in high cross-sectional dimensions using a factor copula structure. The factor copula structure describes the dependence between a large number of observed variables by a smaller set of latent variables (or factors) with time-varying loadings. This allows one to considerably limit the number of parameters required to flexibly describe the dynamics of high-dimensional dependence structures.

Dynamic factor copulas have mainly been implemented for the single-factor case; see the references above. This is predominantly driven by computational reasons. Though adding more factors with dynamic loadings is possible in principle, it would increase the computational burden substantially. In the approach of Oh and Patton (2018) this results from the fact that the densities of the common latent factors and of the idiosyncratic factors do not convolute easily. The copula density

is then not available in closed form and additional numerical methods are required for estimation. This requires considerable computational effort, particularly if multiple factors are used. Creal and Tsay (2015) faced a different challenge as they used a standard parameter driven recurrence equation for the factor loadings dynamics. This introduces additional stochastic components into the model that need to be integrated out (see also Hafner and Manner 2012). Bayesian simulation techniques are used for this integration step, which again becomes computationally expensive as the number of factors with dynamic loadings grows.

Though it is understandable from a computational point of view to restrict oneself to a single factor, it seems too restrictive for most empirical applications. For instance, for panels of equity returns a minimum of three to five factors seems to be the standard (see Fama and French 1993, 2016). A computationally simple yet flexible approach that can easily deal with both the multi-factor setting and dynamic loadings thus seems to be called for.

In this article, we develop exactly such a multi-factor copula model with dynamic loadings. For this purpose, we assume that the cross-sectional units can be grouped using observable characteristics, such as the industry of the firm, its headquarters location, or risk characteristics such as firm size, its book-to-market value, etc. Each of these groups is possibly subject to one or more common factors as well as to group-specific factors. We limit the number of parameters in the model by assuming that all units in a particular group have identical factor loadings. We allow the loadings for each of the factors to vary over time using score-driven dynamics as introduced by Creal, Koopman, and Lucas (2013) and Harvey (2013). Using appropriate

CONTACT Anne Opschoor  a.opschoor@vu.nl 

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/UBES.

© 2020 The Authors. Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

distributional assumptions for the latent common and group-specific factors as well as for the idiosyncratic components, we obtain a model with a tractable, closed-form likelihood expression. Hence, parameter estimation and inference are straightforward using maximum likelihood (ML) and the computational burden is kept to a minimum. In particular, a two-step targeting approach that combines a moment-based estimator and the ML approach leads to fast estimation of the parameters in our most flexible multi-factor copula model. The new multi-factor model can be implemented without any difficulty for high dimensions. In addition, the model easily allows for the inclusion of exogenous variables that help to describe the dynamics of the factor loadings.

As a typical high-dimensional financial dataset, we consider a panel of 100 U.S. daily equity returns across 10 different industries over the period 2001–2014. We group the stocks according to industry, and consider various single- and multi-factor specifications, with Gaussian and Student's t copulas. We compare the factor copula models with three popular multivariate GARCH (MGARCH) models: the cDCC model of Engle (2002) and Aielli (2013), and the DECO and block DECO models of Engle and Kelly (2012). Our comparison is based on in-sample and out-of-sample (density) forecasts. For the latter forecasts, we use the model confidence set (MCS) approach of Hansen, Lunde, and Nason (2011). In addition, we consider the economic performance of the models when used for constructing global minimum variance portfolios (GMVPs).

We find that for our panel of equity returns, both within-industry and between-industry dependence dynamics are key data features that need to be accommodated. Single-factor models and the standard DECO model have difficulty matching these two types of dependence dynamics simultaneously. Our multi-factor specification with Student's t copula, by contrast, outperforms all benchmarks considered in terms of density forecasts, both in-sample and one-step-ahead out-of-sample. When considering the joint lower tail of the multivariate distribution, we again find that the multi-factor Student's t copula model always belongs to the MCS.

For economic criteria, simpler models prevail, though the multi-factor model still belongs to the MCS. Meanwhile, our one-factor specification with heterogeneous dynamic loadings has the best ex-post GMVP performance. We attribute this difference to the character of the global *minimum* variance criterion: differences in *minimum* variance are harder to obtain and typically smaller, such that the increased flexibility of more complex models does not offset the associated estimation risk of the additional parameters used. This contrasts with the criterion based on the full density forecasts, where all dynamics play a more dominant role and the multi-factor specifications work best in-sample and out-of-sample.

As a final novelty in this article, we investigate whether industry classification provides the best grouping structure. We consider alternative classifications based on observable risk characteristics such as firm size, value, or momentum. This provides a further modeling challenge, as the group structure is allowed to vary over time, with corresponding changes in the factor loadings matrix. We find that group classifications based on observable risk characteristics do not outperform the simpler, static classification based on industry.

This article relates to various strands of the literature. First, there is an extensive literature on factor models and the estimation of large covariance matrices (see, e.g., Fan, Fan, and Lv 2008; Fan, Liao, and Mincheva 2011; Fan, Liao, and Liu 2016). Engle, Ng, and Rothschild (1990) developed factor ARCH models with an application to asset pricing with many assets. Factor copulas vis-à-vis factor ARCH models, however, offer more flexibility in choosing the factor structure and distributional assumptions with respect to both the marginals and the dependence structure. Second, factor copulas have recently been introduced by Krupskii and Joe (2013), Creal and Tsay (2015), Oh and Patton (2017), among others. Oh and Patton (2018) and Lucas, Schwaab, and Zhang (2017) are the first to introduce the score-driven framework of Creal, Koopman, and Lucas (2013) within factor copulas. Compared to their work, we consider specifications that yield closed-form densities and use a parameterization that is easily scalable to many factors and high cross-sectional dimensions. Third, we relate to a strand of literature on copula-MGARCH models, such as Christoffersen et al. (2012, 2014), who combine a skewed Student's t copula with a DCC model to study diversification benefits in a panel of more than 200 asset returns. These models suffer in general from the curse of dimensionality mentioned earlier and also require the repeated inversion of (large) covariance or correlation matrices during parameter estimation, which becomes computationally cumbersome and numerically problematic in high dimensions.

The rest of this article is organized as follows. Section 2 presents the multi-factor copula model with dynamic loadings. We carefully lay out the different aspects of our modeling approach, including various possible common factor specifications and the loadings dynamics. We also discuss important details concerning parameter estimation, using either full likelihood estimation, a two-step targeting approach, or composite likelihood (CL) methods. Section 3 studies the performance of the multi-factor copula models in a controlled environment. Section 4 provides the results for the empirical application. Section 5 concludes. An online appendix to this article provides more details on the derivations, as well as more empirical and simulation results for the new models.

2. The Modeling Framework

In this section, we develop the class of closed-form dynamic multi-factor copulas with score-driven loadings. The approach allows for time-varying dependence that remains tractable yet versatile in high-dimensional settings. Our aim is to characterize the conditional joint distribution $F_t(\mathbf{y}_t)$ of the vector $\mathbf{y}_t = (y_{1,t}, \dots, y_{N,t})^\top \in \mathbb{R}^N$ of asset returns in period t , $t = 1, \dots, T$, where the cross-sectional dimension N is possibly large. We decompose $F_t(\mathbf{y}_t)$ into N marginals and a conditional copula as in Patton (2006),

$$\begin{aligned} \mathbf{y}_t | \mathcal{F}_{t-1} &\sim F_t(\mathbf{y}_t) \\ &= C_t \left(F_{1,t}(y_{1,t}; \boldsymbol{\theta}_{M,1,t}), \dots, F_{N,t}(y_{N,t}; \boldsymbol{\theta}_{M,N,t}); \boldsymbol{\theta}_{C,t} \right), \end{aligned} \quad (1)$$

where $C_t(\cdot; \boldsymbol{\theta}_{C,t})$ is the conditional copula given the information set $\mathcal{F}_{t-1} = \sigma(\mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots)$ and the time-varying copula parameter vector $\boldsymbol{\theta}_{C,t}$, and $F_{i,t}(y_{i,t}; \boldsymbol{\theta}_{M,i,t})$, $i = 1, \dots, N$, denotes the conditional marginal distribution of asset i given \mathcal{F}_{t-1} and

the time-varying marginal distribution parameter vector $\theta_{M,i,t}$. We return to the choice of the marginals later. Note that the conditional copula $C_t(\cdot; \theta_{C,t})$ can also be interpreted as the conditional distribution $C_t(\mathbf{u}_t; \theta_{C,t})$ of the probability integral transforms (PITs) $\mathbf{u}_t = (u_{1,t}, \dots, u_{N,t})^\top$ of \mathbf{y}_t , where $u_{i,t} \equiv F_{i,t}(y_{i,t}; \theta_{M,i,t})$ for $i = 1, \dots, N$.

As is well known, decomposing the multivariate (conditional) distribution $F_t(\mathbf{y}_t)$ into its marginals and copula has several advantages. Particularly when the cross-sectional dimension N is large, splitting the modeling task into specifying the marginals and the copula may substantially reduce the computational burden as parameters can be estimated using a two-step approach. As modeling the univariate marginal distributions is relatively simple and fast even for large N , the main remaining challenge is to parsimoniously specify the conditional copula $C_t(\cdot; \theta_{C,t})$. This can be done using factor copulas or multivariate GARCH models like the DCC or DECO models.

2.1. Observation-Driven Dynamic Factor Copulas

The general literature on copula modeling is extensive (see, e.g., Patton 2009, 2013; Fan and Patton 2014 for partial overviews). However, the literature on how to deal with copulas in large cross-sectional dimensions is rather scarce. The main challenge in high dimensions is to keep the parameter space manageable, but at the same time to allow for sufficient flexibility in the dependence structure. To strike this balance, we use a multi-factor copula structure that we endow with score-driven parameter dynamics. Furthermore, we assume that the N asset returns can be clustered into G groups, with assets in the same group having identical factor loadings.

We start from the factor copula structure

$$\begin{aligned} x_{i,t} &= D_{x,i}(x_{i,t}; \tilde{\lambda}_{i,t}, \sigma_{i,t}, \psi_C), \quad i = 1, \dots, N, \quad (2) \\ x_{i,t} &= \tilde{\lambda}_{i,t}^\top \mathbf{z}_t + \sigma_{i,t} \epsilon_{i,t}, \\ \mathbf{z}_t &\stackrel{\text{iid}}{\sim} D_z(\mathbf{z}_t | \psi_C), \quad \epsilon_{i,t} \stackrel{\text{iid}}{\sim} D_\epsilon(\epsilon_{i,t} | \psi_C), \end{aligned}$$

where $\tilde{\lambda}_{i,t}$ is a $k \times 1$ vector of scaled factor loadings, \mathbf{z}_t is a $k \times 1$ vector of common latent factors and $\epsilon_{i,t}$ is an idiosyncratic shock. In addition, \mathbf{z}_t and $\epsilon_{i,t}$ are cross-sectionally and serially independent with distributions D_z and D_ϵ , respectively, depending on a static shape parameter vector ψ_C . Furthermore, D_z and D_ϵ have zero mean and unit (co)variance (matrix). Finally, $D_{x,i}(\cdot)$ denotes the implied marginal distribution of $x_{i,t}$ (see Creal and Tsay 2015). We define the vector $\tilde{\lambda}_{i,t}$ and scalar $\sigma_{i,t}$ as

$$\tilde{\lambda}_{i,t} = \frac{\lambda_{i,t}}{\sqrt{1 + \lambda_{i,t}^\top \lambda_{i,t}}}, \quad \sigma_{i,t}^2 = \frac{1}{1 + \lambda_{i,t}^\top \lambda_{i,t}} \quad (3)$$

for an unrestricted $k \times 1$ vector $\lambda_{i,t}$, such that $x_{i,t}$ has zero mean and unit variance by design. Further parameterization details can be added to ensure for instance that some elements of $\lambda_{i,t}$ are positive by design, for example, by taking exponential transformations or a multinomial logit parameterization. The copula parameter vector gathers all free parameters in $\theta_{C,t} = (\lambda_{1,t}^\top, \dots, \lambda_{N,t}^\top, \psi_C^\top)$.

The correlation matrix of $\mathbf{x}_t = (x_{1,t}, \dots, x_{N,t})^\top$ equals

$$\begin{aligned} \mathbf{R}_t &= \tilde{\mathbf{L}}_t^\top \tilde{\mathbf{L}}_t + \mathbf{D}_t, \\ \tilde{\mathbf{L}}_t &= (\tilde{\lambda}_{1,t}, \dots, \tilde{\lambda}_{N,t}), \\ \mathbf{D}_t &= \text{diag}(\sigma_{1,t}^2, \dots, \sigma_{N,t}^2), \end{aligned} \quad (4)$$

which satisfies all requirements of a correlation matrix, namely positive semidefiniteness and ones on the diagonal.

The factor copula structure in (2) comes with an important computational advantage (see also Creal and Tsay 2015), namely that the inverse and determinant of \mathbf{R}_t are available in closed form as

$$\begin{aligned} \mathbf{R}_t^{-1} &= \mathbf{D}_t^{-1} - \mathbf{D}_t^{-1} \tilde{\mathbf{L}}_t^\top (\mathbf{I}_k + \tilde{\mathbf{L}}_t \mathbf{D}_t^{-1} \tilde{\mathbf{L}}_t^\top)^{-1} \tilde{\mathbf{L}}_t \mathbf{D}_t^{-1}, \\ |\mathbf{R}_t| &= |\mathbf{I}_k + \tilde{\mathbf{L}}_t \mathbf{D}_t^{-1} \tilde{\mathbf{L}}_t^\top| \cdot |\mathbf{D}_t|, \end{aligned} \quad (5)$$

where \mathbf{I}_k denotes the k -dimensional identity matrix and \mathbf{D}_t a diagonal matrix. Computing the inverse of $\mathbf{I}_k + \tilde{\mathbf{L}}_t \mathbf{D}_t^{-1} \tilde{\mathbf{L}}_t^\top$ is relatively easy for two reasons. First, computing the inverse of a diagonal matrix \mathbf{D}_t is straightforward, and the subsequent matrix multiplications are sparse. In particular, $\tilde{\mathbf{L}}_t \mathbf{D}_t^{-1}$ can be computed directly by dividing each column of $\tilde{\mathbf{L}}_t$ by the corresponding diagonal element of \mathbf{D}_t . Second, as the number of common latent factors k is typically much smaller than the number of observed assets N , computing the inverse of the $k \times k$ matrix $\mathbf{I}_k + \tilde{\mathbf{L}}_t \mathbf{D}_t^{-1} \tilde{\mathbf{L}}_t^\top$ is much faster than computing the inverse of the $N \times N$ matrix \mathbf{R}_t .

The class of factor copulas is very flexible. We can vary the number and types of factors, the distributional assumptions of the common factors \mathbf{z}_t and idiosyncratic shocks $\epsilon_{i,t}$, and the dynamics of the factor loadings $\lambda_{i,t}$. The following subsections discuss each of these choices in more detail.

2.1.1. The Factor Structure

Our main goal in this article is to develop feasible dependence structures that allow for multiple factors in a flexible, dynamic way while still giving rise to a closed-form likelihood expression. With our focus on multiple factors, we extend earlier articles that emphasize single-factor implementations, such as Oh and Patton (2018) and Creal and Tsay (2015).

A key aspect of our approach is the assumption that we can split the N assets into G groups according to an observed characteristic such as industry, region, or riskiness, etc. Each group may be subject to several factors, where all assets within a specific group are assumed to have identical factor loadings. As in the block DECO model of Engle and Kelly (2012), this implies that (i) all assets within a group share the same dependence structure, and (ii) the dependence between any pair of assets in two specific, different groups is also the same (but varying across group combinations). This yields a flexible, yet highly parsimonious set-up.

For the sake of exposition, we take the example of $G = 4$ groups with 2 firms in each group throughout this subsection. In reality, of course, the number of groups and the number of firms per group is typically much larger. For instance, in our application in Section 4 we have $G = 10$ groups with up to 19 firms per group.

In our most general specification the loadings matrix is obtained from a lower-triangular matrix with columns containing group-specific loadings. The loadings matrix then takes the form

$$\tilde{\mathbf{L}}_t^\top = \begin{pmatrix} \tilde{\lambda}_{1,1,t} & 0 & 0 & 0 \\ \tilde{\lambda}_{1,2,t} & \tilde{\lambda}_{2,2,t} & 0 & 0 \\ \tilde{\lambda}_{1,3,t} & \tilde{\lambda}_{2,3,t} & \tilde{\lambda}_{3,3,t} & 0 \\ \tilde{\lambda}_{1,4,t} & \tilde{\lambda}_{2,4,t} & \tilde{\lambda}_{3,4,t} & \tilde{\lambda}_{4,4,t} \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad (6)$$

where \otimes denotes the Kronecker product. The first column vector can be interpreted as a common-factor with group-specific loadings, like different market betas. Overall, the loadings matrix could also be seen as a Cholesky decomposition of a 4×4 “quasi correlation” matrix containing the within and between group correlations. The number of unique factor loadings equals $G(G + 1)/2$. Note that the Cholesky decomposition could be sensitive to the different ordering of the groups. We show in Section 4 that in our empirical application this effect is small: estimated dependence measures hardly change when we reorder the variables. We label the model with the factor structure in (6) as the multi-factor lower-triangular (MF-LT) copula model.

A second, much more restricted version of our general specification combines a single common factor with (common) equi-loadings, and a set of G group-specific factors with corresponding group-specific loadings. This results in a loadings matrix with $G + 1$ unique factor loadings

$$\tilde{\mathbf{L}}_t^\top = \begin{pmatrix} \tilde{\lambda}_{1,t} & \tilde{\lambda}_{2,1,t} & 0 & 0 & 0 \\ \tilde{\lambda}_{1,t} & 0 & \tilde{\lambda}_{2,2,t} & 0 & 0 \\ \tilde{\lambda}_{1,t} & 0 & 0 & \tilde{\lambda}_{2,3,t} & 0 \\ \tilde{\lambda}_{1,t} & 0 & 0 & 0 & \tilde{\lambda}_{2,4,t} \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}. \quad (7)$$

For $G \geq 3$ and at least 2 firms in each group, this model meets the necessary requirement for identification. To see this, note that the correlation matrix \mathbf{R}_t for $G = 3$ has 3 within-group correlations and 3 between-group correlations, hence 6 free positions for the 4 different parameters in $\tilde{\mathbf{L}}_t$. For more groups and firms, the number of positions in \mathbf{R}_t increases quadratically, whereas the number of parameters in $\tilde{\mathbf{L}}_t$ increases linearly, thus allowing for overidentification. The first (equi)factor with common loadings $\tilde{\lambda}_{1,t}$ affects both the within-group and the between-group correlations. The group-specific factors with their group-specific loadings, on the other hand, only affect the within-group correlations and not the between-group correlations. We label this model as the multi-factor (MF) copula model.

A third specification is obtained by replacing the group-specific factors in (7) with a common factor with group-specific loadings. The loadings matrix $\tilde{\mathbf{L}}_t^\top$ is then given by

$$\tilde{\mathbf{L}}_t^\top = \begin{pmatrix} \tilde{\lambda}_{1,t} & \tilde{\lambda}_{2,1,t} \\ \tilde{\lambda}_{1,t} & \tilde{\lambda}_{2,2,t} \\ \tilde{\lambda}_{1,t} & \tilde{\lambda}_{2,3,t} \\ \tilde{\lambda}_{1,t} & \tilde{\lambda}_{2,4,t} \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}. \quad (8)$$

From an asset pricing point of view, this second common factor has different betas for each group. Although there are again $G + 1$ unique factor loadings, there is now less freedom to capture the

differences between within-group and between-group effects. Note that $\tilde{\lambda}_{2,1,t}$ cannot be rotated to zero without destroying the equi-loading structure of the first column of $\tilde{\mathbf{L}}_t^\top$, illustrating that the model is identified. We label the model in Equation (8) the 2-Factor (2F) copula model. Omitting the factor corresponding to $\tilde{\lambda}_{1,t}$ in (8) leads to the 1-Factor-Group (1F-Gr) model, which consists of a single factor but with G different group loadings. The 1F-Gr model has also been used in Lucas, Schwaab, and Zhang (2017) and Oh and Patton (2018). Similarly, if instead we drop the factor corresponding to $\tilde{\lambda}_{2,g,t}$ in (8), we obtain a single-factor model with common loadings. We label this special case the 1F-Equi copula model; see also the single-factor copula structures of Oh and Patton (2018) and Creal and Tsay (2015). It corresponds to a DECO correlation structure as in Engle and Kelly (2012), where each pairwise asset correlation is assumed to be the same. From an asset pricing perspective, the single factor can be seen as the market factor with identical betas for all assets.

Table 1 lists all the factor structures considered in this article with their corresponding properties and dimensions.

2.1.2. Distributional Assumptions

Given the various factor structures proposed in Section 2.1.1, the next step is to specify a distribution for the common, group-specific, and idiosyncratic factors in (2). Oh and Patton (2018) assumed a skewed and symmetric Student's t density for the common factor z_t and the idiosyncratic shock $\epsilon_{i,t}$, respectively. As a result, their copula density for $x_{i,t}$ is not available in closed-form. Hence, likelihood evaluation and parameter estimation become computationally involved. Also Creal and Tsay (2015) did not have a likelihood in closed form due to their choice of a new stochastic component in the transition equation for the factor loading $\lambda_{i,t}$. They solve the issue by employing Bayesian (numerical) techniques to estimate the parameters. Again, this is computationally costly for increasing dimensions, particularly in multi-factor settings.

In contrast to the above approaches, we retain tractability of the model and a closed form of the likelihood by two particular choices. First, we make convenient distributional assumptions for the factors z_t and $\epsilon_{i,t}$. Second, we consider a score-driven transition equation for the factor loadings $\lambda_{i,t}$. We discuss the latter in the next subsection.

To model z_t and $\epsilon_{i,t}$, we use the Student's t copula,

$$u_{i,t} = T(x_{i,t}; \nu_C), \quad i = 1, \dots, N, \quad (9a)$$

$$x_{i,t} = \sqrt{\zeta_t} \left(\tilde{\lambda}_{i,t}^\top z_t + \sigma_{i,t} \epsilon_{i,t} \right), \quad z_t \sim N(\mathbf{0}, \mathbf{I}_k),$$

$$\epsilon_{i,t} \sim N(0, 1), \quad (9b)$$

$$\zeta_t \sim \text{Inv-Gamma}(\nu_C/2, \nu_C/2). \quad (9c)$$

where $T(\cdot; \nu_C)$ denotes the cdf of the univariate Student's t distribution with ν_C degrees of freedom, location zero, and unit scale, and ζ_t denotes an independent inverse-gamma distributed random variable. Note that—in contrast to Creal and Tsay (2015) and Oh and Patton (2018)—our proposed factor structures of the previous subsection easily fit into the distributional framework above, while the copula density (and thus the likelihood) retains its analytical closed-form expression. For the special case $\nu_C \rightarrow \infty$, we obtain $\zeta_t \equiv 1$ and a Gaussian copula setting. The Gaussian copula, however, has no tail dependence

Table 1. Various factor structures and their properties.

Name	# factors	# unique factor loadings	Common factor with (equi) common loading	Common factor with group loadings	Group factors with group loadings	dim \tilde{L}_t^\top
1F-Equi	1	1	Yes	No	No	$N \times 1$
1F-Group	1	G	No	Yes	No	$N \times 1$
2F	2	$G + 1$	Yes	Yes	No	$N \times 2$
MF	$G + 1$	$G + 1$	Yes	No	Yes	$N \times (G + 1)$
MF-LT	G	$G(G + 1)/2$	No	Yes	Yes	$N \times G$

NOTE: This table summarizes the various factor structures that are proposed given that there are N assets allocated to G different groups. We show the number of factors, the number of unique factor loadings in \tilde{L}_t^\top , the dimension of the scaled factor loadings matrix and the existence of an equi-factor, group-specific factors and/or group-specific loadings.

(see McNeil, Frey, and Embrechts 2015) and may therefore be less suitable to describe the dependence structure in empirical applications involving financial data.

The copula in (9a)–(9c) is symmetric. Oh and Patton (2017, 2018), by contrast, developed a 1-Factor asymmetric copula model that allows for skewness. By adding an additional term $\boldsymbol{\gamma}\zeta_t$ for $\boldsymbol{\gamma} \in \mathbb{R}^N$ to the right-hand side of (9b) and letting ζ_t be generalized inverse Gaussian, we obtain the generalized hyperbolic (GH) copula class with skewness parameter $\boldsymbol{\gamma}$. A special case is the GH skewed Student’s t copula as used in, for instance, Lucas, Schwaab, and Zhang (2014, 2017). Such a generalization would come at a substantial increase in computational burden as the copula requires the numerical inversion of the marginal cdfs at each point in time for all coordinates. Preliminary experiments for the simplest model structures and a skewed t copula did not result in major in-sample likelihood increases or in substantial changes in the paths of the fitted dynamic dependence parameters. Therefore, we leave such further generalizations for future research and concentrate in this article on the value-added of the multi-factor structures.

2.1.3. Score-Driven Factor Loading Dynamics

To complete our dynamic factor copula specification, we formulate the dynamics of the unique factor loadings within the matrix \tilde{L}_t^\top . We gather these unique time-varying parameters in the vector \boldsymbol{f}_t , whose dimension and content varies across different factor model structures, see Table 1.

In general, there are two approaches to model time-variation in \boldsymbol{f}_t . The first approach is parameter-driven and assumes \boldsymbol{f}_t evolves as a stochastic process driven by its own innovation. This leads to so-called stochastic copula models as in Hafner and Manner (2012) and Creal and Tsay (2015). Estimating such models is typically computationally involved and requires integrating out the random innovations of the time-varying parameters in a numerically efficient way. The second approach is observation-driven and assumes the factor loadings depend on functions of past observables. Our proposal falls into this latter category and uses score-driven dynamics as introduced by Creal, Koopman, and Lucas (2013); see also Harvey (2013) and Oh and Patton (2018). As mentioned before, an advantage of the observation-driven approach is that the likelihood is available in closed-form via a standard prediction error decomposition. This substantially reduces the computational burden compared to a parameter-driven approach.

Score-driven dynamics use the score of the conditional copula density to drive \boldsymbol{f}_t . Intuitively, this adjust the loadings in

a steepest ascent direction using the local log-likelihood fit at time t as the criterion function. The approach has information theoretic optimality properties as argued in Blasques, Koopman, and Lucas (2015) and its generalizations in Creal et al. (2018). As an example in our context, consider a 1-Factor equicorrelation copula, such that $\tilde{L}_t^\top = \tilde{\lambda}_t \boldsymbol{1}_N$ for a scalar parameter $\tilde{\lambda}_t = \lambda_t / \sqrt{1 + \lambda_t^2}$ and $\boldsymbol{f}_t = \lambda_t$, such that $\tilde{\lambda}_t \in [-1, 1]$ by design, where $\boldsymbol{1}_N$ denotes an $N \times 1$ vector filled with ones. Then the score-driven dynamics for \boldsymbol{f}_t are given by $\boldsymbol{f}_{t+1} = \boldsymbol{\omega} + A s_t + B \boldsymbol{f}_t$, with $s_t = \partial \log c(\boldsymbol{x}_t; \tilde{\lambda}_t, \nu_C) / \partial \boldsymbol{f}_t$, and $c(\cdot; \tilde{\lambda}_t, \nu_C)$ the Student’s t copula density. We assume the same type of factor loading dynamics for vector-valued \boldsymbol{f}_t . In that case we allow the intercept vector $\boldsymbol{\omega}$ to have unit or group-specific elements, while we continue to treat A and B as scalars. Extensions to non-scalar A or B are straightforward, and some of these are investigated in the empirical application later on. Following Oh and Patton (2018), we use unit scaling for the score s_t in the sense of Creal, Koopman, and Lucas (2013) to reduce the computational burden of estimating a separate scaling function. As an alternative, the score could be scaled with the Information matrix. Explicit expressions for the score and Information matrix for all factor copula specifications used in our article are provided in online Appendix A.7.

2.2. Benchmarks, Marginals, and Parameter Estimation

We compare the dynamic factor copula models developed above against MGARCH alternatives, in particular the cDCC model (Engle 2002; with the correction of Aielli 2013) and the (block) DECO model of Engle and Kelly (2012); see online Appendix C for the implementation details of these models in our setting. To maintain a fair comparison in high dimensions, we also consider the MGARCH models in a copula framework and use the same marginal models for the MGARCH and multi-factor score-driven copulas.

For the marginal distributions, we use the univariate t -GAS volatility model of Creal, Koopman, and Lucas (2011, 2013). We also perform a robustness check with other marginals, such as univariate GARCH models with skewed t innovations. The results are qualitatively similar. For more details on the estimation results for the marginal models or for the copula results based on PITs from skewed marginal distributions, see online Appendices D and E, respectively.

Parameter estimation requires some further details, both for the factor copula and the MGARCH copula models. To

estimate the model parameters, we use a two-step likelihood based approach. First, we estimate the parameters of each of the marginals (separately). Second, we estimate the copula parameters conditional on the marginal parameter estimates. This approach follows directly from decomposing the joint likelihood as

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}) \equiv & \sum_{t=1}^T \log f_t(\mathbf{y}_t; \boldsymbol{\theta}_t) = \sum_{i=1}^N \sum_{t=1}^T \log f_{i,t}(y_{i,t}; \boldsymbol{\theta}_{M,i,t}) \\ & + \sum_{t=1}^T \log c_t \left(F_{1,t}(y_{1,t}; \boldsymbol{\theta}_{M,1,t}), \dots, \right. \\ & \left. F_{N,t}(y_{N,t}; \boldsymbol{\theta}_{M,N,t}); \boldsymbol{\theta}_{C,t} \right) \end{aligned} \quad (10)$$

with $f_{i,t}(\cdot; \boldsymbol{\theta}_{M,i,t})$ denoting the conditional marginal density corresponding to $F_{1,t}(\cdot; \boldsymbol{\theta}_{M,i,t})$, and $\boldsymbol{\theta}_t = \{\boldsymbol{\theta}_{M,t}, \boldsymbol{\theta}_{C,t}\}$. According to Patton (2013), the implied efficiency loss of the two-step approach compared to the one-step approach is small.

We assume a Student's t and Gaussian copula to model the dependence, as discussed before. For the factor copula specifications, inverses and determinants of \mathbf{R}_t are given in closed form by (5), which substantially reduces the computational burden in high dimensions. This enables us to estimate these models by maximum likelihood.

In case of our most general multi-factor copula model (the MF-LT), we potentially have $G(G+1)/2$ different elements in the vector of intercepts $\boldsymbol{\omega}$ in the score-driven dynamics of the factor loadings. A computational challenge may then arise if G becomes large. In that case, we suggest to estimate the copula parameters using the following two-step procedure. Assuming that the loading process is covariance stationary, and defining the unconditional mean of \mathbf{f}_t as $\bar{\mathbf{f}}$, we have

$$\bar{\mathbf{f}} = \mathbb{E}[\mathbf{f}_{t+1}] = \boldsymbol{\omega} + B \mathbb{E}[\mathbf{f}_t] \Leftrightarrow \bar{\mathbf{f}} = (1 - B)^{-1} \boldsymbol{\omega}. \quad (11)$$

In the first step, we match $\bar{\mathbf{f}}$ to the empirical within-group and between-group correlations using a moment estimator. We do so as follows. Let \mathbf{R}_M denote the $G \times G$ ‘‘quasi unconditional correlation matrix’’ based on $x_{it} = \Phi^{-1}(u_{it})$. The off-diagonal elements $\mathbf{R}_{M,g,h}$ equal the average correlation between all asset pairs from group g and h , $g, h = 1, \dots, G$, $g \neq h$, respectively. The diagonal element $\mathbf{R}_{M,g,g}$ holds the average of pairwise correlations within group g . The moment estimator is then obtained by minimizing

$$L_M = \text{vech}(\mathbf{R}_M - \bar{\bar{\mathbf{L}}}\bar{\bar{\mathbf{L}}}^\top)^\top \text{vech}(\mathbf{R}_M - \bar{\bar{\mathbf{L}}}\bar{\bar{\mathbf{L}}}^\top), \quad (12)$$

with $\bar{\bar{\mathbf{L}}}^\top$ a $G \times G$ lower triangular matrix as in (6) depending on $\bar{\mathbf{f}}$ via the same nonlinear transformation that maps \mathbf{f}_t into $\boldsymbol{\lambda}_{i,t}$ and subsequently into $\tilde{\boldsymbol{\lambda}}_{i,t}$ as in (3). In a second step, we estimate the remaining parameters A and B keeping $\bar{\mathbf{f}}$ fixed and setting $\boldsymbol{\omega} = (1 - B)\bar{\mathbf{f}}$. This two-step targeting procedure substantially decreases the computational burden. The moment estimator in the first step is computed quickly, while in the second step we only need to estimate the two remaining parameters A and B . Note that in the first step \mathbf{R}_M is based on the inverse normal cdf in case of the Gaussian copula. For the Student's t copula, we could use the inverse Student's t cdf, but we show in the

next section that the normal inverse cdf also works well for the Student's t copula case in the moments estimator.

In contrast to the multi-factor models, inverses and determinants of \mathbf{R}_t are not available in closed form for the block DECO and cDCC specifications. We therefore estimate the cDCC model by means of the CL method of Pakel et al. (2020). This technique is based on maximizing the sum of bivariate (copula) log-likelihood values to estimate A and B (and ν_C). In a second step the matrix $\boldsymbol{\Omega}$ is estimated by its sample analogue.

Finally, we also use a CL approach for the block DECO model of Engle and Kelly (2012) by extending their proposal from the Gaussian to the Student's t case. They consider the joint log-likelihood of all the firms in two separate groups $g \neq h$, with $g, h \in \{1, \dots, G\}$, that is,

$$\mathcal{L}_{g,h}^{\text{Stud}} = \sum_{t=1}^T \left[-\frac{1}{2} \log |\mathbf{R}_t| - \frac{\nu + n_g + n_h}{2} \log \left(1 + \frac{\mathbf{x}_t^\top \mathbf{R}_t^{-1} \mathbf{x}_t}{\nu - 2} \right) \right], \quad (13)$$

where $|\mathbf{R}_t|$ and \mathbf{R}_t^{-1} are given analytically for the 2-block case by Lemma 3.1 in Engle and Kelly (2012). The CL method now maximizes the sum of all log-likelihoods of each pair of blocks $g > h$,

$$\max \mathcal{L}_{CL} = \max \sum_{g>h} \mathcal{L}_{g,h}^{\text{Stud}}, \quad (14)$$

where the intercept $\boldsymbol{\Omega}$ is estimated by the unconditional correlation matrix of \mathbf{x}_t . Note that for $\nu \rightarrow \infty$, we recover the Gaussian block DECO model, which is the specification used in most of the literature. As argued before, however, the Gaussian copula lacks tail dependence and may therefore be less suitable for financial data.

3. Simulation Experiment

We briefly report the results of three Monte Carlo experiments, conducted to study the properties of the new method. Full details can be found in the online Appendix B.2.

In the first experiment, we investigate the accuracy of estimation and inference in the new model. Panel A of Table 2 presents the outcomes for a set-up with an $N = 100$ dimensional time series of length $T = 1000$ with $G = 10$ equally sized groups holding $N/G = 10$ individual cross-sectional units each. These settings roughly correspond to the data dimensions in our empirical application. The data-generating process (DGP) is the MF copula model from Equation (7). We only report results for A , B , and ν_C . Results for $\boldsymbol{\omega}$ and for smaller sample sizes can be found in the online appendix and are qualitatively similar. We find that all parameters are estimated near their true values. Comparing results over sample sizes (in the online appendix), we see that the standard deviation decreases approximately with a factor $\sqrt{2}$. By comparing the Monte Carlo standard error of the estimates (std column in Table 2) with the mean of the estimated standard error over all replications (mean(SE) column), we find that our computed standard errors fairly reflect estimation uncertainty. Overall, Panel A shows that the parameters and standard errors of the Gaussian and Student's t factor copulas with score-driven dynamic factor loadings can be accurately estimated if the model is correctly specified.

Table 2. Monte Carlo results of parameter estimates of the multi-factor copula.

Coef.	True	Normal			Student's <i>t</i>		
		Mean	Std	Mean(SE)	Mean	Std	Mean(SE)
Panel A: MF, $T = 1000$							
$A_{eq}(N)$	0.0085	0.0085	0.0009	0.0008			
$A_{gr,f}(N)$	0.0095	0.0093	0.0018	0.0018			
$A_{eq}(t)$	0.0150				0.0149	0.0020	0.0019
$A_{gr,f}(t)$	0.0100				0.0096	0.0016	0.0016
$B(N)$	0.8700	0.8626	0.0221	0.0248			
$B(t)$	0.9200				0.9149	0.0129	0.0126
ν_C	35.00				35.1760	1.8629	1.8821
Panel B: MF-LT, $T = 1000$							
A	0.015	0.0161	0.0006	0.0006	0.0161	0.0007	0.0006
B	0.970	0.9700	0.0025	0.0023	0.9697	0.0024	0.0024
ν_C	35.00				35.06	1.862	1.846

NOTE: This table provides Monte Carlo results of parameter estimates using the multi-factor (MF) Gaussian and *t*-copula model as given in (7), and the MF-LT model based on (6). For full details, see the online Appendix B.2. $B(N)$ and $B(t)$ denote the value of B in case of the Gaussian (N) and Student's *t* (*t*) factor copula model, respectively. The table reports the mean and standard deviation of the estimated coefficients, as well as the mean of the computed standard error. Results are based on 1000 Monte Carlo replications.

Panel B of Table 2 shows results for the MF-LT model from Equation (6). The results for A , B , and ν_C are similar to those of the MF model. Our two-step targeting approach for ω thus appears to work well both for estimation and inference. The estimates of ω as shown in the online appendix reveal that the standard deviations of moment-based estimators for ω_i are higher than the standard errors of the ML estimators for A , B , and ν_C . The two-step estimator thus implies a large computational gain at the expense of some efficiency loss in the estimation of ω . The assumed distribution does not appear to have a major impact on the performance.

Finally, we investigate the impact of misspecification of the factor structure on the estimated dependence structure. In this third experiment, we consider a DGP with $N = 25$, $T = 1000$, and the MF factor structure, using Student's *t*(35) distributed errors (with $\nu_C = 35$ based on the empirical application) for $G = 5$ different groups, each containing $N/G = 5$ units. Using different (possibly mis-specified) factor copula models, we compute the time average of the squared Frobenius norm of $\widehat{R}_t - R_t$, which is a consistent loss function according to Laurent, Rombouts, and Violante (2013). The results in Table 3 clearly indicate that underestimating the number of factors causes substantial discrepancies between the true and the fitted dependence dynamics, particularly for one-factor models with an equi-loading structure, or for the multi-factor models that ignore the different between-group dependencies. This holds irrespective of the distribution used.

4. Empirical Application

4.1. Data

In our high-dimensional empirical application, we investigate the daily open-to-close returns of 100 randomly chosen constituents of the S&P 500 index during the period January 2, 2001 until December 31, 2014 ($T = 3521$ days). Table B.1 in the online appendix provides an overview of the ticker symbols of

Table 3. Performance of misspecified factor copulas.

	MF-LT	MF	2F	1F-Group	1F-Equi
Student's <i>t</i>	0.058 (0.023)	1.692 (0.266)	1.600 (0.222)	0.891 (0.107)	2.541 (0.308)
Gaussian	0.080 (0.023)	1.699 (0.263)	1.567 (0.205)	0.914 (0.109)	2.559 (0.312)

NOTE: This table summarizes the mean and standard deviation of the average Euclidian distance between a simulated correlation matrix R_t ($t = 1, \dots, 1000$) from the MF-LT model with *t*(35)-distributed errors and the estimated \widehat{R}_t based on one, two, or multi-factor copula models with either Gaussian or a Student's *t* distributions. All results are based on 1000 replications.

all stocks. The same table shows the classification of the stocks into 10 groups based on the industry of the firm. In our sample, financials form the largest group with 19 firms, followed by consumer services and energy, respectively. Each industry group includes at least four firms.

To model the marginal characteristics of daily stock returns, we estimate univariate *t*-GAS volatility models as given in Equations (D.1) and (D.2) in the online Appendix D. For the conditional mean, we find at most two significant autoregressive (AR) lags. We therefore use an AR(2) conditional mean specification in the marginal models for all 100 stocks. Estimation results are summarized in Table D.1 in online Appendix D. We find that the mean of the estimated degrees of freedom parameter ν of the Student's *t* distribution equals 8.22, underlining the fat-tailed nature of daily stock returns even after filtering for time-varying volatility. The mean estimate of β (0.991) reflects the usual strong persistence in volatility.

We follow Creal and Tsay (2015) and evaluate the fit of the marginal distributions by transforming the PITs $\widehat{u}_{i,t}$ into Gaussian variables $\widehat{x}_{i,t} = \Phi^{-1}(\widehat{u}_{i,t})$, $t = 1, \dots, T$. We subsequently test each series $\widehat{x}_{i,t}$ for normality using the Kolmogorov-Smirnov test. Across the 100 firms, we only reject the null hypothesis of normality for 5 series at the 5% significance level. We conclude that the marginal models are adequate for our subsequent analysis.

As a robustness check, we also estimated univariate GARCH models with the skewed Student's t distribution of Hansen (1994) and compared this to GARCH- t models for all assets. The results for the skewed Student's t GARCH models are reported in online Appendix D. The comparison indicates that the average increase in the maximized log-likelihood relative to GARCH t models is a modest 1.3 points. Given this weak evidence for the presence of skewness, we therefore stick to the standard Student's t distribution for our main analysis.

4.2. Full-Sample Copula Comparisons

After estimating the parameters of the marginal distributions, we proceed to estimate the parameters of the score-driven factor copula models and the benchmark MGARCH copula models using the full sample of 3521 observations. The factor copulas are based on grouping firms into industries as laid out in Table B.1 of the online appendix.

Table 4 shows the parameter estimates and maximized log-likelihood values for all models. Panels A.1 and A.2 contain results for Gaussian and t factor copula specifications, respectively: a one-factor copula with homogeneous (1F-Equi) or with industry-specific (1F-Group) loadings, a two-factor copula (2F) with one factor with homogeneous loadings and one factor with industry-specific loadings, a multi-factor copula (MF) with 10 industry factors, and a multifactor model (MF-LT) with a triangular loadings matrix. Panels B.1 and B.2 contain results for Gaussian and t benchmark copulas from the MGARCH class: the cDCC, DECO, and block DECO models. In both multi-factor copula models, we assume that the B parameter in the GAS transition equation is the same for all factor loadings. For the MF-LT model we assume a common scalar A , while A is allowed to differ between the common factor and the industry-specific factors in the MF model. In the 2F model, we also allow for different A values for the two common factors, while assuming the same B value. Finally, for the 1F-Gr model we

Table 4. Parameter estimates of the full sample.

Model	ω^{eq}	A^{eq}	A^{gr}	B	ν_C	LogL	AIC	# para
Panel A.1: Gaussian factor copulas								
1F-Equi	0.017 (0.002)	0.005 (0.000)		0.975 (0.003)		65,934	-131,862	3
1F-Group			0.007 (0.001)	0.970 (0.006)		68,086	-136,148	12
2F	0.047 (0.006)	0.012 (0.000)	0.013 (0.001)	0.941 (0.009)		71,667	-143,306	14
MF	0.042 (0.005)	0.012 (0.000)	0.014 (0.001)	0.930 (0.009)		81,827	-163,626	14
MF-LT		0.009 (0.001)		0.964 (0.005)		83,226	-166,339	57
Panel A.2: t -factor copulas								
1F-Equi	0.062 (0.013)	0.012 (0.001)		0.918 (0.016)	36.52 (1.52)	69,679	-139,350	4
1F-Group			0.005 (0.000)	0.986 (0.001)	31.87 (1.11)	72,293	-144,560	13
2F	0.004 (0.002)	0.009 (0.001)	0.006 (0.001)	0.993 (0.002)	38.57 (1.64)	77,828	-155,627	15
MF	0.033 (0.002)	0.012 (0.001)	0.012 (0.001)	0.957 (0.002)	44.98 (1.78)	84,858	-169,687	15
MF-LT		0.004 (0.000)		0.990 (0.002)	36.22 (1.38)	86,433	-172,749	58
Panel B.1: Gaussian copula-MGARCH models								
cDCC (CL)		0.017 (0.001)		0.968 (0.003)		74,263	-138,623	4952
DECO		0.071 (0.001)		0.929 (0.001)		64,474	-119,044	4952
Block DECO		0.030 (0.002)		0.957 (0.003)		83,087	-156,270	4952
Panel B.2: t copula-MGARCH models								
cDCC (CL)		0.018 (0.001)		0.968 (0.002)	14.17 (0.58)	82,688	-155,470	4953
DECO		0.106 (0.000)		0.894 (0.000)	34.43 (0.80)	69,314	-128,721	4953
Block DECO		0.032 (0.002)		0.955 (0.003)	22.51 (0.60)	86,222	-162,537	4953

NOTE: This table reports maximum likelihood parameter estimates of various factor copula models, the (block) DECO model of Engle and Kelly (2012) and the cDCC model of Engle (2002), applied to daily returns of 100 stocks included in the S&P 500 index. We consider five different factor copula models, see Table 1 for the definition of their abbreviations. Panel A.1 presents the factor models with a Gaussian copula density, Panel A.2 presents the parameter estimates corresponding with the Student's t copula. Panel B.1 and B.2 present the estimates of the MGARCH class of models. In case of the cDCC and block DECO models, the table shows parameters estimates obtained by the composite likelihood (CL) method. Standard errors are provided in parenthesis and based on the (sandwich) robust covariance matrix estimator. We report the copula log-likelihood, the Akaike information criteria (AIC) as well as the number of estimated parameters for all models. The sample comprises daily returns from January 2, 2001 until December 31, 2014 (3521 observations).

assume a common A and B parameter for all different groups. To save space, we do not report all the different intercepts ω_g for all groups for the factor copulas with group-specific loadings. These detailed results are provided in the online Appendix E. Standard errors are based on the sandwich (robust covariance matrix) estimator $\hat{H}_0^{-1} \hat{G}_0 \hat{H}_0^{-1}$ with \hat{H}_0 the inverse Hessian of the likelihood, and \hat{G}_0 the outer product of gradients.

Five interesting results emerge from Table 4. First, in terms of the statistical fit, the MF-LT t model outperforms the other factor-copula models, as well as the MGARCH-copula models (cDCC, DECO and block DECO). The MF-LT model not only achieves the highest total log-likelihood value, but also performs best in terms of AIC, which takes into account the number of estimated parameters.

Second, multi-factor models provide a much better fit than one-factor copula models. For example, the log-likelihood difference between the MF-LT t copula and the 1F-Equi t copula is more than 15,000 points. The largest gain with respect to the factor structure is obtained by including industry factors, that is, extending the 1F-Equi model to the MF specification. This increases the log-likelihood by 15,000 points in both the Gaussian and Student's t case. Note that allowing for industry specific loadings in the single-factor model leads to a much more modest improvement in the log-likelihood of 2500 points. Extending the single-factor model with a second factor with industry-specific loadings performs better, but the increase in the log-likelihood is still only half of the improvement achieved by the MF specification.

Third, the Student's t factor copulas fit considerably better than their Gaussian counterparts. Log-likelihood differences range between 3000 and 6000 points, depending on the specification. Differences for the multi-factor specifications are

typically at the lower end of this range. This underlines that allowing for more than one factor also takes care of part of the tail clustering.

Fourth, we find strong persistence in the time-varying factor loadings with a value of $B \approx 0.97$ for most of the estimated (t -)factor copula models. This finding, as well as the previous one, confirms the empirical results of Oh and Patton (2018) using an entirely different dataset of log-differences of U.S. CDS spreads.

Finally, we note that the estimated degrees of freedom parameter ν_C is (much) lower for the block DECO t and cDCC t specifications than for the MF-LT t model or the DECO model. It seems that there is empirically some bias effect due to the use of the composite versus the ML approach to parameter estimation.

Our main results are robust against two variations in the estimation set-up. First, we re-estimate all models based on PITs obtained from estimating a skewed Student's t GARCH model for the marginals. Second, we investigate the sensitivity of the MF-LT t model with respect to the ordering of the industries by re-estimating the MF-LT t model for 50 different random industry orderings. Online Appendix E shows the results for both robustness checks and confirms that our conclusions continue to hold.

Figure 1 shows an example of within and between industry correlation differences for two industries. For clarity, each panel compares the MF-LT t model to one of its competitors. The upper-left and lower-right panels show that the 1-Factor specifications under-estimate within correlation levels, as they have to compromise within and between correlations using the same dynamic loadings. The upper panels also show that the MF-LT model results in much less noisy correlation estimates,

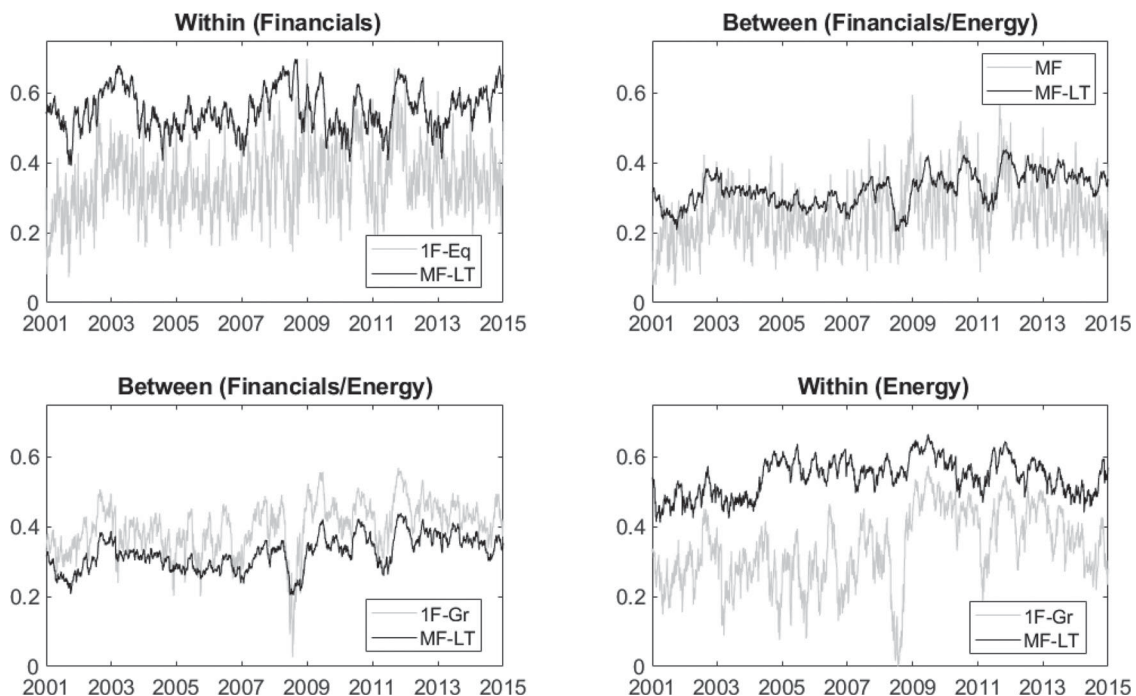


Figure 1. In-sample within and between industry correlations of the MF-LT t copula. This figure shows the fitted within industry correlations of financials and energy, as well as their between industry correlations. The panels compare the MF-LT t model output to that of the 1F-Equi t , 1F-Group t , and MF t model, respectively. The sample spans the period from January 2, 2001 until December 31, 2014 ($T = 3521$ days).

both with respect to the MF and the 1F-Eq model. The main take-away is that our multi-factor models pick up the within industry specific correlations that cannot be captured by a single factor model. This ability of the MF and MF-LT models explains their substantial increases in statistical fit as shown before.

4.3. Alternative Groups Based on Dynamic Risk Factors

So far, we have allocated firms into groups using their industry classification. Alternative group allocations are of course possible. Here we investigate an obvious alternative by forming groups based on key asset pricing risk factors, including firm size (market capitalization), value (book-to-market), and momentum (see Fama and French 1993; Carhart 1997). Due to data availability, this reduces the sample from 100 to 90 assets. For the size and value factors, we form 10 new groups in July each year based on deciles of sorted market capitalization and book-to-market values of the previous fiscal year. Similarly, for momentum we sort stocks into deciles in January each year based on their sorted past 12-month returns. This mimics the way these factors are constructed in typical asset pricing studies. Using risk factors of this type to form groups comes with an additional challenge, namely that group composition can now change from one period to the next. This can easily be accommodated in the factor copula approach introduced in this article by a straightforward but tedious bookkeeping exercise to account for possible switches in factor loadings of firms depending on their group allocation at time t . The use of such grouping criteria in a dynamic factor copula framework is new and can be seen as a separate contribution of this article.

Table 5 shows the estimation results for our preferred in-sample model, the MF-LT copula, using different grouping criteria. Given the time-varying group composition, we alter our targeting approach for ω by each year using the moment estimator (12) based on the unconditional correlation matrix R_M of $\hat{x}_{it} = \Phi^{-1}(\hat{u}_{it})$ for the 250 daily observations of the upcoming year. This smaller targeting sample may of course influence the accuracy of the estimates of ω . In a second step, we estimate the parameters A and B . The results show that the models with dynamic groups based on risk factors achieve a considerably worse statistical fit than the model with static

industry groups. The minimum loss in log-likelihood exceeds 14,000 points. Among the three risk factors, momentum seems to perform best, but differences with size and value are small. The conclusions on the preferred grouping structure do not depend on the distributional assumption, and are similar for the Gaussian and the Student's t case.

4.4. Multivariate Density Forecasts

As we have closed-form copula density expressions, a natural way to compare the out-of-sample (OOS) forecasting performance of factor copula models and copula MGARCH models is to consider multivariate density forecasts as in Salvatierra and Patton (2015). Because we use the same marginal distributions in all models, the density forecast comparison actually boils down to an evaluation and comparison of the OOS copula density forecasts.

We use a moving estimation window of 1000 observations (or roughly four calendar years), which leaves $P = 2521$ observations for the out-of-sample period, starting December 28, 2004. Hence, the OOS period includes the Great Financial Crisis. We re-estimate the parameters in all models after each 50 observations (or roughly 10 calendar weeks) and construct a one-step ahead copula density forecast each day.

We evaluate the copula density forecasts using two scoring rules. First, we consider accuracy using the densities' full support by means of the log scoring rule (see Mitchell and Hall 2005; Amisano and Giacomini 2007)

$$S_{l,t}(\hat{\mathbf{u}}_t, M_j) = \log c_t(\hat{\mathbf{u}}_t | \hat{\boldsymbol{\theta}}_{C,t}, M_j), \quad (15)$$

where $c_t(\cdot | \hat{\boldsymbol{\theta}}_{C,t}, M_j)$ is the Gaussian or Student's t conditional copula density obtained from model M_j and $\hat{\mathbf{u}}_t$ denotes the vector of corresponding PITs. Note that the PITs in \mathbf{u}_t are based on the same marginal distributions for both model specifications in any log score comparison, and that the marginal densities therefore drop out from a difference in log scores between two models. We therefore omit the marginals from the log score expression in (15). This underlines that we are really comparing the forecasting quality of the copula part.

Second, we focus on the joint lower region of the copula support by using the conditional likelihood (cl) scoring rule

Table 5. Forming groups in the MF-LT model.

	Gaussian copula			t copula			
	A	B	LogL	A	B	ν_C	LogL
Value	0.009 (0.001)	0.965 (0.005)	61,025	0.004 (0.000)	0.995 (0.001)	29.16 (0.98)	64,950
Size	0.010 (0.001)	0.904 (0.010)	60,955	0.003 (0.001)	0.966 (0.011)	28.58 (0.92)	64,953
Momentum	0.009 (0.001)	0.963 (0.003)	61,368	0.003 (0.001)	0.992 (0.002)	29.33 (1.01)	65,144
Industry	0.010 (0.001)	0.964 (0.006)	76,531	0.005 (0.000)	0.989 (0.002)	34.49 (1.22)	79,458

NOTE: This table reports maximum likelihood parameter estimates of the multi-factor copula model, applied to daily equity returns of 90 assets listed at the S&P 500 index. The 10 groups associated with the models are formed based on different risk-factors, such as the book-to-market ratio, size and momentum. In addition, we consider the model based on industry groups. Standard errors are provided in parenthesis and based on the (sandwich) robust covariance matrix estimator. We report only the A and B of all estimated parameters (hence omitting the intercepts) and the copula log-likelihood for all models. The sample comprises daily returns from January 2, 2001 until December 31, 2014 (3521 observations).

proposed by Diks et al. (2014),

$$S_{cl,t}(\hat{\mathbf{u}}_t, M_j) = \left(\log c_t(\hat{\mathbf{u}}_t | \hat{\boldsymbol{\theta}}_{C,t}, M_j) - \log C_t(\mathbf{q} | \hat{\boldsymbol{\theta}}_{C,t}, M_j) \right) \times I[\hat{\mathbf{u}}_t < \mathbf{q}], \quad (16)$$

where \mathbf{q} is an $N \times 1$ vector and $C_t(\cdot | \hat{\boldsymbol{\theta}}_{C,t}, M_j)$ is the conditional copula function, and $I[\hat{\mathbf{u}}_t < \mathbf{q}] = \prod_{i=1}^N I[\hat{u}_{i,t} < q_i]$ with $q_i \in [0, 1]$, $i = 1, \dots, N$. Hence, (16) is the log-likelihood of model M_j conditional on $\mathbf{u}_t < \mathbf{q}$ (element-wise). For any $\mathbf{q} = (q_1, \dots, q_N)$ this boils down to the joint lower region $[0, q_1] \times \dots \times [0, q_N]$. Obviously, when $q_i = 1$ for all i , we recover the log scoring rule. We use a time-varying threshold vector $\mathbf{q}_t = (\bar{q}_t, \dots, \bar{q}_t)$, where \bar{q}_t is such that $\frac{1}{1000} \sum_{j=1}^{1000} I[\hat{\mathbf{u}}_{t-j} < \mathbf{q}_t] = q$, with $q = 0.01$ or 0.05 , for each t in the (rolling) estimation sample. We thus compare the copula density forecasts in the joint empirical lower 1% or 5% tail.

For both scoring rules, models that deliver higher values are preferred. We can test whether differences in the scoring rule values for models M_i and M_j are significant by defining the score differential

$$d_{x,ij,t} = S_{x,t}(\hat{\mathbf{u}}_t, M_i) - S_{x,t}(\hat{\mathbf{u}}_t, M_j), \quad \text{with } x = l, cl. \quad (17)$$

The null hypothesis of equal predictive ability is equivalent to $H_0 : \mathbb{E}[d_{x,ij,t}] = 0$, which can be tested using a standard Diebold and Mariano (1995) test statistic. Since we deal with a substantial number of different models and factor structures and hence many different copula density forecasts, we consider the MCS of Hansen, Lunde, and Nason (2011). The MCS automatically accounts for the dependence between model outcomes given that all models are based on the same data.

Table 6 shows the results of the copula density forecast evaluation. We report the mean of the log scores and the conditional likelihood scores, as well as the p -values of the MCS.

The table shows three interesting results. First, in line with our full-sample results, the MF-LT t model performs best in terms of predictive ability when evaluated over the full copula support using the log scoring rule. The same pattern emerges from the MCS. The MCS p -value equals 1 for the MF-LT t , whereas that of all other models is below 0.01. Second, similar to the in-sample results, most of the gain for the factor copulas is obtained by allowing for industry-specific factors. For example, changing the equifactor from fixed (1F-Equi t) to industry-specific loadings (1F-Group t) increases the average log-score by only 0.75 points (from 21.08 to 21.83). Allowing for different industry factors (MF t), however, implies an additional increase of almost 4 points to an average log-score of 25.60. Allowing for cross-exposures in the MF-LT specification results in yet a further increase by 0.5 points. Third, when we consider density forecasts in the joint lower tail, the MF-LT model is always part of the MCS. In that case, however, also the MGARCH specifications perform well and are included in the MCS, in particular the block DECO- t model. The differences in the conditional likelihood scores are small in these cases, however, and below 0.015 points.

Overall, we conclude that the flexibility provided by the new MF-LT t model is also important out of sample using density forecast criteria. The more flexible parameterization allows for a larger class of dependence matrices than more restrictive one-factor models. This extension appears to be empirically important in high dimensions.

Table 6. One-step ahead copula density forecasts.

Model	Full	1% tail	5% tail
	$S_{ls,t}(p\text{-val})$	$S_{cl,t}(p\text{-val})$	$S_{cl,t}(p\text{-val})$
1F-Equi	20.07 (0.00)	1.401 (0.00)	4.022 (0.00)
1F-Equi t	21.08 (0.00)	1.443 (0.00)	4.142 (0.00)
1F-Group N	20.73 (0.00)	1.411 (0.00)	4.045 (0.00)
1F-Group t	21.83 (0.00)	1.445 (0.00)	4.177 (0.00)
2F N	22.52 (0.00)	1.436 (0.00)	4.138 (0.00)
2F t	23.53 (0.00)	1.469 (0.01)	4.267 (0.00)
MF N	24.95 (0.00)	1.466 (0.00)	4.284 (0.00)
MF t	25.60 (0.00)	1.494 (0.22)	4.373 (0.04)
MF-LT N	25.32 (0.00)	1.469 (0.00)	4.291 (0.00)
MF-LT t	26.10 (1.00)	1.500 (0.40)	4.400 (0.36)
cDCC N	22.42 (0.00)	1.501 (0.50)	4.369 (0.36)
cDCC t	24.37 (0.00)	1.509 (1.00)	4.384 (0.36)
DECO N	19.76 (0.00)	1.415 (0.00)	4.020 (0.00)
DECO t	21.01 (0.00)	1.447 (0.00)	4.143 (0.00)
Block DECO N	25.24 (0.00)	1.478 (0.03)	4.314 (0.00)
Block DECO t	26.02 (0.01)	1.505 (0.76)	4.415 (1.00)

NOTE: This table evaluates the accuracy of one-step ahead copula density forecasts (in the left tail) of daily return series for 100 stocks from the S&P500 index, obtained by various factor copula and copula MGARCH models, assuming a Gaussian or Student's t distribution (denoted by N or t). We consider a 1-Factor model with equi-loadings (1F-Equi), a 1-Factor model with group-specific loadings (1F-Group), a 2-Factor model with one equifactor and an additional factor with group-specific loadings (2F), a multi-factor copula model with one equi-factor plus G group-specific factors (MF), and the lower triangular multi-factor model (MF-LT). In addition, we show the results of the cDCC model of Engle (2002) and the (block) DECO model of Engle and Kelly (2012). The table presents the mean of the log score (S_{ls}) and the conditional (tail) likelihood score (S_{cl}) for the lower joint 1% and 5% tail. We present the p -value associated with the model confidence set of Hansen, Lunde, and Nason (2011) in parentheses. Bold numbers in this row represent models that belong to the model confidence set at a significance level of 5%. The out-of-sample period covers December 28, 2004 until December 31, 2014 and contains 2521 observations.

4.5. Economic Out-of-Sample Performance

Finally, we assess the forecasting performance of the different models from an economic perspective. We do so by considering the ex-post variance of the ex-ante GMVP; compare Chiriac and Voev (2011) and Engle and Kelly (2012), among others. The best forecasting model should provide portfolios with the lowest ex-post variance.

Assuming that an investor aims to minimize the 1-step ahead portfolio volatility at time t subject to being fully invested, the resulting GMVP weights $\mathbf{w}_{t+1|t}$ are obtained as the solution of the quadratic programming problem

$$\min \mathbf{w}_{t+1|t}^\top (\mathbf{H}_{t+1|t} \mathbf{R}_{t+1|t}^* \mathbf{H}_{t+1|t}) \mathbf{w}_{t+1|t}, \quad \text{s.t. } \mathbf{w}_{t+1|t}^\top \mathbf{1} = 1, \quad (18)$$

with $\mathbf{H}_{t+1|t}$ the 1-step forecasts of the variances based on the marginal models, and $\mathbf{R}_{t+1|t}^*$ the one-step ahead forecast of the correlation matrix. As the forecast of the correlation matrix $\mathbf{R}_{t+1|t}^*$ is not the same as the forecast of the copula dependence matrix $\mathbf{R}_{t+1|t}$, we obtain the former by simulating 20,000 returns from the joint distribution of returns as constructed from the marginals and the conditional copula. Following Chiriac and Voev (2011), we assess the predictive ability of the different models by comparing the results to the *ex-post* realizations of the conditional standard deviation $\sigma_{p,t}$, given by $\sigma_{p,t} = \sqrt{\mathbf{w}_{t+1|t}^\top \mathbf{RC}_{t+1} \mathbf{w}_{t+1|t}}$, with \mathbf{RC}_{t+1} the realized covariance matrix obtained using 5-min returns. We decompose this matrix into

realized variances and a realized correlation matrix, where the latter is by definition ill-conditioned and not positive definite. We use the “eigenvalue cleaning” method used by Hautsch, Kyj, and Oomen (2012) to get a positive definite correlation matrix. Having constructed the ex-post conditional portfolio standard deviation, we test the model performance by means of the MCS approach with a significance level of 5%.

Alongside the GMVP’s volatility, we also calculate a number of other relevant quantities, such as portfolio turnover (TO_t), concentration (CO_t), and the total short position (SP_t) for each competing model at time t . Turnover at time t is defined as

$$TO_t = \sum_{i=1}^N \left| w_{i,t+1|t} - w_{i,t|t-1} \frac{1 + y_{i,t}}{1 + \mathbf{w}_{t|t-1}^T \mathbf{y}_t} \right|, \quad (19)$$

where $w_{i,t|t-1}$ is the i th element of the weight vector $\mathbf{w}_{t|t-1}$. It measures the value of the portfolio that is bought/sold when rebalancing the portfolio to its new optimal position from time t to $t + 1$. A model that produces more stable correlation matrix forecasts implies in general less turnover and hence, less transaction costs. Portfolio concentration and total portfolio short position both measure the amount of extreme portfolio allocations. Again, more stable forecasts of $\mathbf{R}_{t+1|t}^*$ should result in less extreme portfolio weights. Portfolio concentration is defined as

$$CO_t = \left(\sum_{i=1}^N w_{i,t|t-1}^2 \right)^{1/2}, \quad (20)$$

while the total portfolio short position SP_t is given by

$$SP_t = \sum_{i=1}^N w_{i,t|t-1} \cdot I[w_{i,t|t-1} < 0], \quad (21)$$

with $I[\cdot]$ an indicator function that takes the value 1 if the i th element of the weight vector is lower than zero.

Table 7 reports the economic out-of-sample performance results. As for the density forecast results, the factor copulas also perform best in terms of economic performance. Again, the MF-LT model is included in the MCS. There are, however, also a number of remarkable differences. In terms of the ex-post variance of the GMVP, the 1-Factor copulas with industry specific loadings now perform best. This contrasts with the density forecast setting, where the MF-LT t model performed best in-sample and out-of-sample. The multi-factor models also now still (marginally) outperform the block-DECO model in terms of ex-post variance of the GMVP. The 1-Factor models, however, have the best performance, both in terms of ex-post variance, turnover, concentration, and total short positions. Unlike the density forecast setting, we also note that the choice of the distribution plays a less important role in Table 7.

To reconcile the findings in terms of economic performance with those of the density forecast evaluation from the previous subsection, it is important to note that the GMVP evaluation takes a very specific perspective. The GMVP focuses on an area of the forecast distribution where differences are more concentrated by design: all models focus on a portfolio with ex-ante *minimum* variance. If the different models are any good, differences in this concentrated performance measure are harder to

Table 7. Minimum variance portfolio results.

Model	$\bar{\sigma}_p$ (p-val)	TO	CO	SP
1F-Equi N	0.527 (0.00)	0.359	0.336	−0.595
1F-Equi t	0.528 (0.00)	0.351	0.339	−0.604
1F-Group N	0.513 (0.97)	0.332	0.307	−0.530
1F-Group t	0.513 (1.00)	0.319	0.310	−0.541
2F N	0.537 (0.00)	0.395	0.339	−0.604
2F t	0.542 (0.00)	0.377	0.343	−0.615
MF N	0.537 (0.00)	0.418	0.357	−0.636
MF t	0.538 (0.00)	0.407	0.359	−0.642
MF-LT N	0.523 (0.07)	0.407	0.353	−0.622
MF-LT t	0.522 (0.11)	0.397	0.353	−0.625
cDCC N	0.556 (0.00)	0.763	0.470	−1.032
cDCC t	0.564 (0.00)	0.774	0.477	−1.054
DECO N	0.525 (0.01)	0.345	0.335	−0.596
DECO t	0.525 (0.00)	0.346	0.335	−0.598
Block DECO N	0.531 (0.00)	0.398	0.352	−0.622
Block DECO t	0.532 (0.00)	0.399	0.352	−0.623

NOTE: This table reports results on a global minimum variance portfolio strategy, based on 1-step ahead predictions of the daily covariance matrix, according to four different type of factor copulas, the cDCC model of Engle (2002) and the (block) DECO model of Engle and Kelly (2012). The columns represent two types of one-factor copulas (one equi-factor or one factor with group-specific loadings, denoted by 1F-eq and 1F-gr), one 2-Factor model (one equi-factor plus an additional factor with group-specific loadings) and two types of multi-factor copula models (one equi-factor plus G group-specific factors and the MF-LT model). Each type of model is further discriminated across distribution (Gaussian vs. a Student’s t) For each model, we show the mean of the ex-post portfolio standard deviation, the p -value corresponding with the model confidence set of Hansen, Lunde, and Nason (2011), using a significance level of 5%, and the mean of the portfolio turnover (TO), concentration (CO) and the total portfolio short positions (SP). Bold numbers indicate the models that stay within the MCS, or the lowest (absolute) portfolio turnover, concentration and total portfolio short positions. The out-of-sample period goes from December 28, 2004 until December 31, 2014 and contains 2521 observations.

obtain. This is corroborated by the results in Table 7: although the results are sometimes statistically significantly different, they are all quite close in economic terms (with the possible exception of the cDCC). Using such a performance measure therefore benefits more parsimonious models over richly parameterized models with their associated estimation risk. This may explain why the simpler 1-Factor models do better here. By contrast, if the full density or complete tail area is taken into account as in the previous subsection, the additional flexibility of the more complex factor models has a beneficial effect on performance, particularly in the current high-dimensional setting.

5. Conclusions

We introduced factor structures within the class of closed-form factor copula models for high dimensions. The new factor copula model is computationally tractable with score-driven dynamics, implying a closed-form conditional copula density. Parameters can be estimated in a straightforward way by ML and/or a fast two-step approach that combines a moment-based estimator and the ML approach.

Our factor structures are based on group-specific characteristics such as industry classification. In addition, an important feature of our model is that it allows for more than one factor. Extensions to the model are also easily possible, such as the inclusion of covariates to describe the factor loadings dynamics,

and different time-varying group structures, such as the risk-based groups (size, value, momentum) also considered in this article. This can be done without causing any difficulty for the positive definiteness of the implied dependence matrix.

Empirically, we modeled the dependence across 100 equity returns from the S&P 500 index over the period 2001–2014. We found that our factor copula models outperform multivariate GARCH (MGARCH) based counterparts, such as the (c)DCC and (block) DECO. In-sample, the multi-factor copula model has a better fit than one-factor models and benchmarks such as the cDCC and (block-)DECO. Out-of-sample, the good performance of multi-factor copula models persists. A simple static industry-based group structure for the copula appears better statistically than risk-based groups based on size, value, or momentum. Measured in terms of density forecasts, the multi-factor models perform best, whereas in terms of the GMVP variance simpler 1-Factor models outperform other models. In all settings, we thus find score-driven factor copulas to describe the dynamics of the data well. Given their computational ease and closed-form likelihood expression, they thus provide a useful tool for modeling high-dimensional dynamic dependence structures.

Supplementary Materials

The supplementary materials contain a web appendix with derivations of the score of the proposed multi-factor copula models, and further implementation details. Second, Matlab code and an associated README file are provided with testdata in order to estimate the parameters of all factor-copula models.

Acknowledgments

We thank Andrew Patton, the associate editor, two anonymous referees, David Blaauw, Tijn Wijdogen, and participants at the 10th Annual SoFiE conference and seminar participants at Tinbergen Institute Amsterdam, Lund University, Heidelberg University, Maastricht University, and Vrije Universiteit Amsterdam for helpful comments.

References

- Aielli, G. (2013), “Dynamic Conditional Correlation: On Properties and Estimation,” *Journal of Business and Economic Statistics*, 31, 282–299. [2,5]
- Amisano, G., and Giacomini, R. (2007), “Comparing Density Forecasts via Weighted Likelihood Ratio Tests,” *Journal of Business and Economic Statistics*, 25, 177–190. [10]
- Blasques, F., Koopman, S., and Lucas, A. (2015), “Information-Theoretic Optimality of Observation-Driven Time Series Models for Continuous Responses,” *Biometrika*, 102, 325–343. [5]
- Carhart, M. M. (1997), “On Persistence in Mutual Fund Performance,” *The Journal of Finance*, 52, 57–82. [10]
- Cherubini, U., Mulinacci, S., Gobbi, F., and Romagnoli, S. (2011), *Dynamic Copula Methods in Finance*, Chichester: Wiley. [1]
- Chiriac, R., and Voev, V. (2011), “Modelling and Forecasting Multivariate Realized Volatility,” *Journal of Applied Econometrics*, 26, 922–947. [11]
- Christoffersen, P., Errunza, V., Jacobs, K., and Jin, X. (2014), “Correlation Dynamics and International Diversification Benefits,” *International Journal of Forecasting*, 30, 807–824. [2]
- Christoffersen, P., Errunza, V., Jacobs, K., and Langlois, H. (2012), “Is the Potential for International Diversification Disappearing? A Dynamic Copula Approach,” *Review of Financial Studies*, 25, 3711–3751. [2]
- Creal, D., Koopman, S., and Lucas, A. (2011), “A Dynamic Multivariate Heavy-Tailed Model for Time-Varying Volatilities and Correlations,” *Journal of Business Economic Statistics*, 29, 552–563. [5]
- (2013), “Generalized Autoregressive Score Models With Applications,” *Journal of Applied Econometrics*, 28, 777–795. [1,2,5]
- Creal, D., Koopman, S., Lucas, A., and Zamojski, M. (2018), “Generalized Autoregressive Method of Moments,” Tinbergen Institute Discussion Paper, TI 14-029/III (updated). [5]
- Creal, D., and Tsay, R. (2015), “High Dimensional Dynamic Stochastic Copula Models,” *Journal of Econometrics*, 189, 335–345. [1,2,3,4,5,7]
- Diebold, F., and Mariano, R. (1995), “Comparing Predictive Accuracy,” *Journal of Business and Economic Statistics*, 13, 253–263. [11]
- Diks, C., Panchenko, V., Sokolinskiy, O., and van Dijk, D. (2014), “Comparing the Accuracy of Multivariate Density Forecasts in Selected Regions of the Copula Support,” *Journal of Economic Dynamics and Control*, 48, 79–94. [11]
- Engle, R. (2002), “Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models,” *Journal of Business and Economic Statistics*, 20, 339–350. [2,5,8,11,12]
- Engle, R., and Kelly, B. (2012), “Dynamic Equicorrelation,” *Journal of Business Economic Statistics*, 30, 212–228. [2,3,4,5,6,8,11,12]
- Engle, R., Ng, V., and Rothschild, M. (1990), “Asset Pricing With a Factor-Arch Covariance Structure: Empirical Estimates for Treasury Bills,” *Journal of Econometrics*, 45, 213–237. [2]
- Fama, E., and French, K. (1993), “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics*, 33, 3–56. [1,10]
- (2016), “Dissecting Anomalies With a Five-Factor Model,” *Review of Financial Studies*, 29, 69–103. [1]
- Fan, J., Fan, Y., and Lv, J. (2008), “High Dimensional Covariance Matrix Estimation Using a Factor Model,” *Journal of Econometrics*, 147, 186–197. [2]
- Fan, J., Liao, Y., and Liu, H. (2016), “An Overview of the Estimation of Large Covariance and Precision Matrices,” *The Econometrics Journal*, 19, C1–C32. [2]
- Fan, J., Liao, Y., and Mincheva, M. (2011), “High Dimensional Covariance Matrix Estimation in Approximate Factor models,” *Annals of Statistics*, 39, 3320. [2]
- Fan, Y., and Patton, A. (2014), “Copulas in Econometrics,” *Annual Review of Economics*, 6, 179–200. [1,3]
- Hafner, C., and Manner, H. (2012), “Dynamic Stochastic Copula Models: Estimation, Inference and Applications,” *Journal of Applied Econometrics*, 27, 269–295. [1,5]
- Hansen, B. (1994), “Autoregressive Conditional Density Estimation,” *International Economic Review*, 35, 705–730. [8]
- Hansen, P., Lunde, A., and Nason, J. (2011), “The Model Confidence Set,” *Econometrica*, 79, 453–497. [2,11,12]
- Harvey, A. (2013), *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*, Cambridge: Cambridge University Press. [1,5]
- Hautsch, N., Kyj, L., and Oomen, R. (2012), “A Blocking and Regularization Approach to High-Dimensional Realized Covariance Estimation,” *Journal of Applied Econometrics*, 27, 625–645. [12]
- Krupskii, P., and Joe, H. (2013), “Factor Copula Models for Multivariate Data,” *Journal of Multivariate Analysis*, 120, 85–101. [2]
- Laurent, S., Rombouts, J., and Violante, F. (2013), “On Loss Functions and Ranking Forecasting Performances of Multivariate Volatility Models,” *Journal of Econometrics*, 173, 1–10. [7]
- Lucas, A., Schwaab, B., and Zhang, X. (2014), “Conditional Euro Area Sovereign Default Risk,” *Journal of Business and Economic Statistics*, 32, 271–284. [1,5]
- (2017), “Modeling Financial Sector Joint Tail Risk in the Euro Area,” *Journal of Applied Econometrics*, 32, 171–191. [1,2,4,5]
- Manner, H., and Reznikova, O. (2012), “A Survey on Time-Varying Copulas: Specification, Simulations, and Application,” *Econometric Reviews*, 31, 654–687. [1]
- McNeil, A., Frey, R., and Embrechts, P. (2015), *Quantitative Risk Management: Concepts, Techniques and Tools*, Princeton, NJ: Princeton University Press. [1,5]
- Mitchell, J., and Hall, S. (2005), “Evaluating, Comparing and Combining Density Forecasts Using the KLIC With an Application to the Bank of

- England and NIESR Fan-Charts of Inflation,” *Oxford Bulletin of Economics and Statistics*, 67, 995–1033. [10]
- Oh, D., and Patton, A. (2017), “Modeling Dependence in High Dimensions With Factor Copulas,” *Journal of Business and Economic Statistics*, 35, 139–154. [1,2,5]
- (2018), “Time-Varying Systemic Risk: Evidence From a Dynamic Copula Model of CDS Spreads,” *Journal of Business and Economic Statistics*, 36, 181–195. [1,2,3,4,5,9]
- Pakel, C., Shephard, N., Sheppard, K., and Engle, R. F. (2020), “Fitting Vast Dimensional Time-Varying Covariance Models,” *Journal of Business & Economic Statistics* (just-accepted). [6]
- Patton, A. (2006), “Modelling Asymmetric Exchange Rate Dependence,” *International Economic Review*, 47, 527–556. [1,2]
- (2009), “Copula-Based Models for Financial Time Series,” in *Handbook of Financial Time Series*, eds. T. Mikosch, J. P. Kreiß, R. Davis, and T. Andersen, Berlin, Heidelberg: Springer, pp. 767–785. [1,3]
- (2013), “Copula Methods for Forecasting Multivariate Time Series,” in *Handbook of Economic Forecasting* (Vol. 2), eds. G. Elliott and A. Timmermann, Oxford: Elsevier, pp. 899–960. [1,3,6]
- Salvatierra, I., and Patton, A. (2015), “Dynamic Copula Models and High Frequency Data,” *Journal of Empirical Finance*, 30, 120–135. [10]