Dynamic Clustering in Multi-Factor Copulas with Hidden Markov Models

Abstract

This study proposes a dynamic multi-factor copula model with time-varying, data-driven group assignments. Transitions of firms between groups are modelled using a hidden Markov model, driven by the distance between clusters and the past likelihood of group membership. Using daily returns from S&P 100 stocks between 2015 and 2024, the model is evaluated against two static benchmarks: k-means clustering and industry-based classifications. It was found that the dynamic clustering approach consistently outperforms the static alternatives. Notably, a model with 15 dynamic groups yields better forecasts than an otherwise identical model with 21 static groups. The results show that time-varying group assignments enable the model to adapt to changes in firm characteristics while preserving sufficient persistence in cluster assignments.

1 Introduction

Correlations between asset returns tend to increase significantly during periods of market stress or economic shocks (Chesnay & Jondeau, 2001). This was evident during the 2007–2008 global financial crisis, where models that ignored joint extreme events contributed to the collapse of the housing market (Coval et al., 2009; Zimmer, 2012). This crisis emphasized the need for models that can capture dynamic dependencies in unstable economic conditions. However, financial markets often involve more than 50 variables, resulting in high model complexity (Manner & Reznikova, 2012). This estimation difficulty is reduced by copulas, which separate each variable's marginal behaviour from their dependence structure (Smith, 2015).

Previous research on copulas focused on modelling time-varying dependence through dynamic factor loadings. These loadings represent how strongly each variable is influenced by one or more underlying latent factors, which capture common sources of variation in returns. Oh & Patton (2017) extended this approach by introducing multi-factor copulas with pre-specified static clusters based on SIC industry codes. More recently, Oh & Patton (2023) derived clusters directly from the data using k-means clustering. Assuming stable group assignments over time, they showed that a model with just five data-driven clusters outperforms a comparable model using 21 industry-based clusters in out-of-sample forecasts.

However, no study yet has combined multi-factor copulas with time-varying, data-driven group assignments. Allowing estimated clusters to change over time may better reflect real-world dynamics, such as firms shifting industries, changing strategies or making acquisitions. This idea is supported by João et al. (2023), who developed a linear panel model incorporating a hidden Markov process that allows firms to switch clusters. Their results show that enabling these switches leads to improved model fit. This study investigates whether incorporating time-varying cluster assignments within multi-factor copula models similarly improves predictive performance. Therefore, the research question is: "How does incorporating time-varying cluster assignments in a high-dimensional multi-factor copula model affect predictive performance relative to static clusters based on industry classifications or k-means clustering?"

To answer this research question, initial clusters were estimated using the k-means algorithm, with group transitions over time modelled via a hidden Markov model. This dynamic clustering approach was compared to two benchmark methods: clustering based on Standard Industrial Classification codes and static k-means clustering. In an empirical analysis, the proposed model was applied to daily returns from stocks in the S&P 100 index over the period 2015–2024 to evaluate its real-world performance. Several copula types were considered, including Gaussian, t, and skewed t, along with both static and dynamic factor loadings.

2 Data and Methodology

2.1 Data

The empirical analysis investigates the daily returns of constituents in the S&P 100 index. The sample period is January 2, 2015, to December 31, 2024, including $T = 2515$ trading days. The dataset contains $N = 98$ stocks that were included in the index as of December 31, 2024, and that were continuously traded throughout the full sample period. A list of all included firms can be found in Appendix A.

2.2 A Dynamic Multi-Factor Copula Model

This study builds on the skewed-t copula model proposed by Oh & Patton (2023). Their approach uses a multi-factor copula model with G clusters of variables, each with a market and group-specific factor loading. In addition, it includes a skewness parameter to capture the asymmetric dependence patterns frequently observed in asset returns. The model is estimated in two stages. First, univariate marginal distributions are fitted to each time series using an AR(1) process for the conditional mean and a GJR-GARCH(1,1) model of Glosten et al. (1993) for the conditional variance. In the second stage of the estimation, conditional on these marginals from the previous stage, the joint dependence among the transformed variables is modelled using a skewed t copula. Time variation in the copula is captured by modelling the factor loadings with Generalized Autoregressive Score (GAS) dynamics (Creal et al., 2013), which updates parameters based on the gradient of the conditional copula log-likelihood.

2.3 Time-Varying Clusters with Markov-switching

The multi-factor copula model assumes G clusters of firms. Clusters can be pre-assigned using industry classifications. Alternatively, Oh & Patton (2023) estimated them from the data using k-means clustering, obtaining static groups fixed over the entire sample period. This study extends their method by modelling cluster memberships as Markov states. Firms can then switch clusters over time, while temporal dependence is preserved (Frühwirth-Schnatter, 2011). The implementation of this method is based on João et al. (2023), who combine a hidden Markov model (HMM) with a linear panel model. Here, their approach is adapted for compatibility with a copula model.

Specifically, the cluster membership of firm i is described by the latent process γ_{it}, where $\gamma_{it} = g$ if firm i belongs to cluster g at time t. The initial cluster assignments are set equal to those obtained by static k-means clustering, as described by Oh & Patton (2023). Let $\pi_{gkt} := \mathbb{P}\{\gamma_{i,t+1} = k \mid \gamma_{it} = g\}$ denote the probability of transitioning from state g to state k at time t. These transition probabilities are assumed to be homogeneous across firms and are collected in the transition matrix Π_t. Assuming transitions are more likely between nearby clusters, transition probabilities π_{gkt} are modelled as a function of distance between clusters:

$$\pi_{gkt} = \frac{\exp(-\delta d_{gk,t-1})}{\sum_{q=1}^{G} \exp(-\delta d_{gq,t-1})} \quad g, k = 1, \ldots, G, \tag{1}$$

where $d_{gk,t}$ denotes the distance between clusters g and k at time t and $\delta \geq 0$ is a parameter that controls how fast the transition probability decays as the distance increases. Distances are based on the common market factor:

$$d_{gk,t} = \mid \lambda_{g,t}^{M} - \lambda_{k,t}^{M} \mid . \tag{2}$$

Since cluster memberships are unobserved, they are inferred using filtered probabilities. These probabilities, denoted by $\tau_{ig,t|t} := \mathbb{P}[\gamma_{it} = g \mid \mathcal{F}_t; \theta]$, contain the probability that unit i belongs to cluster g at time t, conditional on the observed data up to time t, $\mathcal{F}_t$. The copula likelihood at time t is then computed by summing over the likelihood of all possible cluster states, weighted by their predicted probabilities $\tau_{ig,t|t-1}$. Note that the copula defines a joint distribution over all units simultaneously and thus does not allow decomposition into individual likelihoods. To address this, the conditional mixture likelihood is defined (DeSarbo & Cron, 1988). This likelihood accounts for uncertainty in the cluster assignment of a single unit i, while keeping the cluster assignments of all other units fixed at their previously estimated values:

$$\mathbf{c}^{(i)}(\mathbf{u}_t \mid \hat{\Gamma}_{t-1}, \mathcal{F}_{t-1}; \boldsymbol{\theta}) = \sum_{g=1}^{G} \tau_{ig,t|t-1} \cdot \mathbf{c}(\mathbf{u}_t \mid \tilde{\Gamma}_{i,g,t-1}, \mathcal{F}_{t-1}; \boldsymbol{\theta}), \tag{3}$$

where $\mathbf{c}(\cdot)$ is the likelihood of the static Gaussian copula, $\hat{\Gamma}_t$ is the vector of estimated cluster assignments at time t and $\tilde{\Gamma}_{i,g,t}$ is identical to $\hat{\Gamma}_t$ except that variable i is reassigned to cluster g.

The predicted cluster probabilities $\tau_{ig,t+1|t}$ are updated recursively using the forward algorithm (Hamilton, 1989). The Markov property implies that the probability that firm i is in cluster g at time $t+1$ depends on the probability that it is currently in cluster k, and on the probability of transitioning from cluster k to g:

$$\tau_{ig,t+1|t} = \mathbb{P}[\gamma_{i,t+1} = g \mid \mathcal{F}_t; \boldsymbol{\theta}] = \sum_{k=1}^{G} \pi_{kgt} \mathbb{P}[\gamma_{it} = k \mid \mathcal{F}_t; \boldsymbol{\theta}] = \sum_{k=1}^{G} \pi_{kgt} \tau_{ik,t|t}. \tag{4}$$

The second step of the forward algorithm is updating the filtered probabilities $\tau_{ig,t|t}$ via Bayes' rule. The predicted probabilities are combined with the conditional likelihood of the observed data under each potential cluster assignment for firm i, while keeping the cluster assignments of all other firms fixed:

$$\begin{aligned} \tau_{ig,t|t} = \mathbb{P}[\gamma_{it} = g \mid \hat{\Gamma}_{t-1}, \mathcal{F}_t; \boldsymbol{\theta}] = \frac{\tau_{ig,t|t-1}\ \mathbf{c}(\mathbf{u}_t \mid \tilde{\Gamma}_{i,g,t-1}, \mathcal{F}_{t-1}; \boldsymbol{\theta})}{\mathbf{c}^{(i)}(\mathbf{u}_t \mid \hat{\Gamma}_{t-1}, \mathcal{F}_{t-1}; \boldsymbol{\theta})} \\ = \frac{\tau_{ig,t|t-1}\ \mathbf{c}(\mathbf{u}_t \mid \tilde{\Gamma}_{i,g,t-1}, \mathcal{F}_{t-1}; \boldsymbol{\theta})}{\tau_{i1,t|t-1}\ \mathbf{c}(\mathbf{u}_t \mid \tilde{\Gamma}_{i,1,t-1}, \mathcal{F}_{t-1}; \boldsymbol{\theta}) + \cdots + \tau_{iG,t|t-1}\mathbf{c}(\mathbf{u}_t \mid \tilde{\Gamma}_{i,G,t-1}, \mathcal{F}_{t-1}; \boldsymbol{\theta})}. \end{aligned} \tag{5}$$

The forward filtering algorithm described in Eqs.(4)-(5) is performed for each variable, after which each unit is assigned to the cluster with the highest posterior probability:

$$\gamma_{i,t} = \arg\max_{g} \tau_{ig,t|t}. \tag{6}$$

For robustness, the forward filtering algorithm is run repeatedly for all firms, updating their cluster assignments until no firm switches clusters between iterations. In practice, convergence is typically achieved after just one iteration.

Regarding the out-of-sample forecasts, each variable is assigned to the cluster with the highest predicted probability from Eq.(4), after which the probabilities are updated using Bayes's rule for use in the next time period.

2.4 Model Estimation and Evaluation

Joint estimation of the dynamic copula parameters and time-varying group assignments is computationally demanding. Therefore, a three-stage estimation procedure is used. A description of this procedure is provided in Appendix B. After estimation, model performance is assessed using the Akaike Information Criterion, the economic relevance of the clusters, and forecast accuracy.

Forecast accuracy is assessed using two scoring rules: the log-likelihood and the conditional likelihood score. The log-likelihood scoring rule of Amisano & Giacomini (2007) measures predictive performance over the full support of copula model M_i, and is defined as

$$S_{l,t}(\hat{\mathbf{u}}_t, M_i) = \log c_t(\hat{\mathbf{u}}_t \mid \hat{\theta}_{C,t}, M_i), \tag{7}$$

where $c_t(\cdot)$ denotes the copula density at time t, and $\hat{\mathbf{u}}_t$ is the vector of estimated uniform marginals. Dependence in the joint lower tail is especially important in risk management and finance, as extreme losses often occur simultaneously. Therefore, the copula is evaluated using the conditional likelihood score proposed by Diks et al. (2014), which focuses on the lower tail:

$$S_{cl,t}(\hat{\mathbf{u}}_t, M_i) = \left(\log c_t(\hat{\mathbf{u}}_t \mid \hat{\theta}_{C,t}, M_i) - \log C_t(\mathbf{q} \mid \hat{\theta}_{C,t}, M_i)\right) \times I[\hat{\mathbf{u}}_t < \mathbf{q}], \tag{8}$$

where $\mathbf{q}$ is an $N \times 1$ threshold vector, $C_t(\cdot \mid \hat{\theta}_{C,t}, M_i)$ is the copula distribution, and $I[\hat{\mathbf{u}}_t < \mathbf{q}] = \prod_{i=1}^{N} I[\hat{u}_{i,t} < q_i]$ indicates joint threshold exceedance. Equation (8) thus measures the log-likelihood of model M_i conditional on the event $\hat{\mathbf{u}}_t < \mathbf{q}$, corresponding to the lower tail region $[0, q_1] \times \cdots \times [0, q_N]$. To allow for time variation, the threshold vector is defined as $\mathbf{q}_t = (\bar{q}_t, \ldots, \bar{q}_t)$, where $\bar{q}_t$ satisfies

$$\frac{1}{1000} \sum_{j=1}^{1000} I[\hat{\mathbf{u}}_{t-j} < \mathbf{q}_t] = q,$$

for a specified tail probability q, such as 0.05. To compare predictive accuracy across models while controlling for multiple testing, the Diebold & Mariano (2002) test is combined with the Model Confidence Set procedure of Hansen et al. (2011), which iteratively eliminates the worst-performing models until a subset of statistically indistinguishable models remains.

3 Results

3.1 Marginal Model

The marginal model is estimated for each individual return series, with detailed results reported Table 3 in Appendix C. Summary statistics show heavy tails and slight negative skewness in the returns (see Figure 1 in Appendix C), while standardized residuals retain mild skewness and substantial excess kurtosis, supporting the skewed t model. Heterogeneity in pairwise correlations further motivates a copula model to capture non-linear and asymmetric dependencies across stock returns.

3.2 In-sample Evaluation of Dynamic and Benchmark Clusters

Using transformed residuals from the marginal model, the copula model is estimated to capture dependence between return series. Time-varying clusters are benchmarked against k-means and industry-based clusters (one- and two-digit Standard Industry Classification (SIC) codes), using a Gaussian copula with static factor loadings. Model fit is evaluated with the Akaike Information Criterion (AIC), where lower values indicate better performance. The optimal value of the transition decay parameter δ (see Eq. (1)) is found by grid search to be 20.

The results show that time-varying clusters achieve significantly better AIC values than static clustering methods, confirming the potential of data-driven dynamic group assignments as suggested by João et al. (2023). Figure 2 in Appendix D shows AIC values across different numbers of clusters G, indicating that 21 is the optimal number of groups for both static and dynamic clusters. Additionally, a model with only six estimated groups outperforms the 21 two-digit SIC groups, consistent with the findings of Oh & Patton (2023).

3.3 Estimated Cluster Assignments

To investigate the economic relevance of cluster transitions, the initial static clusters are compared with the clusters at the end of the sample period. A complete overview of these clusters is provided in Table 4 and Table 5 in Appendix D, respectively. Factor loadings for the model with dynamic cluster assignments, estimated for Gaussian, t, and skewed t copulas with GAS dynamics, are reported in Table 6. The static groups show strong alignment with two-digit SIC classifications and generally show meaningful coherence, although some inconsistencies remain. Allowing for time-varying transitions via the hidden Markov model resolves several of these inconsistencies. For example, General Electric, initially clustered with financial firms, is reassigned to group 4 with industrial companies such as General Dynamics, better reflecting its core business. Improvements are also seen in within-group correlations, with Figure 3 in Appendix D showing that time-varying clusters yield higher correlations for General Electric's group and two other cases.

Across the full sample, 183 transitions are observed, with around 4% of stocks switching each quarter. Figure 4 and Figure 5 in Appendix E report the fraction of stocks switching each quarter and the number of transitions per stock, respectively. Transition rates are low early in the sample but rise from 2020, peaking in early 2021 during the COVID-19 crisis. Elevated rates are also observed in the second and fourth quarters of 2024, possibly linked to geopolitical and economic uncertainty. About 42.9% of firms never change clusters. A small number of stocks switch repeatedly ("flickering"), likely reflecting proximity to multiple cluster boundaries. Some groups do not experience changes, such as groups 12 and 13 (energy and utilities, and oil and gas), which may be related to the steady and clearly defined nature of these industries.

3.4 Out-of-sample Forecast Performance

Next, the dynamic model is evaluated out-of-sample using a rolling estimation window of 1,000 observations, resulting in 1,515 out-of-sample observations, ranging from December 21, 2018 to December 31, 2024. Model parameters are re-estimated every 250 observations, and a one-step-ahead copula density forecast is constructed for each day. For the dynamic clustering model, the transition parameter is set to the optimal in-sample value $\delta = 20$, which also demonstrated robust performance in the out-of-sample evaluation.

Table 1 presents the results of the copula density forecast evaluation across different group sizes for SIC, static k-means, and dynamic clustering. It reports the time-averaged log-likelihood scores $S_{l,t}$, 5% left-tail conditional log-likelihood scores $S_{cl,t}$, and p-values from the Model Confidence Set (MCS). The left panel shows results for static factor loadings, and the right panel for GAS dynamic loadings. The table shows three interesting results. First, dynamic clustering consistently outperforms static clustering across all group sizes under the log scoring rule, in line with in-sample results. In both panels, the model with 21 dynamic groups achieves the highest average log-likelihood, with an average log-likelihood of 31.44 under static loadings and 31.94 with GAS dynamics. In addition, both 21-group models achieve an MCS p-value of 1.00, indicating statistically superior predictive performance relative to all other models. Second, dynamic clustering improves the 5% conditional log score relative to the static clusters, suggesting better performance in capturing joint downside risk. However, the improvement is smaller than for the full log-likelihood, and the MCS includes some static models. Third, allowing for time-varying factor loadings through GAS dynamics leads to further improvements in both the overall and conditional left-tail log-likelihood, even when combined with dynamic clustering.

Table 1
One-step ahead copula density forecasts

	Static loadings		GAS loadings	
	Full	5% tail	Full	5% tail
	$S_{l,t}$(p-val)	$S_{cl,t}$(p-val)	$S_{l,t}$(p-val)	$S_{cl,t}$(p-val)
SIC 1-digit static	24.80(0.00)	2.395(0.00)	24.89(0.00)	2.297(0.00)
SIC 2-digit static	27.48(0.00)	2.424(0.00)	27.62(0.00)	2.289(0.00)
3 static groups	24.14(0.00)	2.401(0.00)	24.47(0.00)	2.363(0.00)
6 static groups	26.80(0.00)	2.417(0.00)	27.24(0.00)	2.413(0.00)
9 static groups	27.86(0.00)	2.417(0.00)	28.14(0.00)	2.415(0.00)
12 static groups	28.88(0.00)	2.422(0.01)	29.29(0.00)	2.414(0.00)
15 static groups	28.96(0.00)	**2.443(0.07)**	29.38(0.00)	2.428(0.00)
18 static groups	30.00(0.00)	**2.447(0.14)**	30.54(0.00)	2.439(0.00)
21 static groups	30.24(0.00)	**2.453(0.31)**	30.59(0.00)	2.456(0.00)
24 static groups	28.62(0.00)	2.334(0.03)	28.97(0.00)	2.438(0.00)
SIC 1-digit dynamic	27.40(0.00)	2.401(0.00)	27.93(0.00)	2.374(0.00)
SIC 2-digit dynamic	29.95(0.00)	2.436(0.00)	30.22(0.00)	2.418(0.00)
3 dynamic groups	24.93(0.00)	2.289(0.00)	25.58(0.00)	2.405(0.00)
6 dynamic groups	27.48(0.00)	2.348(0.00)	27.95(0.00)	2.417(0.00)
9 dynamic groups	28.14(0.00)	2.412(0.00)	28.62(0.00)	2.418(0.00)
12 dynamic groups	29.06(0.00)	2.419(0.00)	29.43(0.00)	2.421(0.00)
15 dynamic groups	30.30(0.00)	**2.451(0.28)**	30.76(0.00)	**4.497(0.08)**
18 dynamic groups	30.56(0.00)	**2.453(0.31)**	31.02(0.00)	**2.509(0.72)**
21 dynamic groups	**31.44(1.00)**	**2.459(0.76)**	**31.94(1.00)**	**2.516(1.00)**
24 dynamic groups	29.68(0.00)	**2.464(1.00)**	29.89(0.00)	2.483(0.00)

Note: This table reports the accuracy of one-step-ahead copula density forecasts for daily returns of S&P 100 stocks, using a multi-factor copula model with Student's t distribution and transition decay $\delta = 20$. The mean log score ($S_{l,t}$) and the 5% conditional log-likelihood score ($S_{cl,t}$) for the lower tail are shown. p-values from the Model Confidence Set (MCS) procedure of Hansen et al. (2011) are reported in parentheses, with bold numbers indicating models that belong to the MCS of their column at a significance level of 5%. The out-of-sample period runs from December 21, 2018, to December 31, 2024, including 1,515 observations. Note that the 1-digit SIC clusters consist of 8 groups and the 2-digit SIC clusters contain 21 groups.

To further evaluate out-of-sample forecasting performance, alternative copula specifications were compared, considering both static versus dynamic (GAS) factor loadings and different copula families (Gaussian, t and skew t). The significance of differences in out-of-sample likelihoods is assessed using the Diebold & Mariano (2002) test, with standard errors computed with the Newey & West (1987) estimator based on 10 lags. All DM-test results can be found in Table 7 and 8 in Appendix F for static and dynamic clustering, respectively.

Three main findings were obtained. First, all test statistics comparing static and GAS versions of the HMM models were found to be positive and highly significant, indicating that the incorporation of GAS dynamics improves the fit across all copula types and group sizes. The effect was observed to be particularly strong for the skewed t copula relative to the Gaussian and Student's t copulas. Second, when copula families were compared under GAS dynamics, the Student's t copula was consistently found to outperform both the Gaussian and skewed t copulas across all group sizes. This outcome is consistent with the findings of Oh & Patton (2023). Third, the skewed t copula was shown to perform substantially worse in out-of-sample forecasts, in some cases performing even worse than the Gaussian copula. The poor performance of the skewed t copula may be due to the inherent penalty that forecast evaluations impose on estimation uncertainty. Unless additional parameters differ significantly from zero and are estimated with sufficient accuracy, the model may achieve better predictive accuracy by excluding them altogether. Finally, Table 7 shows that, in general, the patterns observed for dynamic clusters also hold for static clusters. However, while the skewed t copula still performs poorly, it is not consistently outperformed by the Gaussian copula across all group sizes.

3.5 Economic determinants of Forecast Performance

Forecast performance is further evaluated across economic environments using the Conditional Equal Predictive Ability (CEPA) and Conditional Superior Predictive Ability (CSPA) tests. The CEPA test of Giacomini & White (2006) examines whether predictive accuracy depends on market conditions, while the CSPA test proposed by Li et al. (2022) assesses whether any alternative model systematically outperforms the benchmark. Economic conditions are summarized by market volatility (VIX), cross-sectional dispersion of returns, and the alpha from the Capital Asset Pricing Model (CAPM).

The detailed results of both tests are shown in Table 9 in Appendix G. The results indicate that the dynamic clustering model consistently outperforms static benchmarks. Compared to industry-based clusters, gains are not strongly linked to economic variables. In contrast, when compared to k-means clusters, improvements in forecasting performance are larger in periods of high volatility, high dispersion, and greater CAPM alpha, with alpha being most influential. The CSPA test confirms that dynamic clustering dominates static benchmarks across the entire range of conditions.

4 Conclusion

This paper extends the factor copula model of Oh & Patton (2023) by relaxing the assumption of static group assignments. Firms often adjust their strategy, enter or exit markets, and engage in mergers or acquisitions. As a result, the dependence structures among firms may also change over time, making fixed clusters less realistic. Therefore, this study proposes a copula model that incorporates time-varying clusters using a hidden Markov model, adapting the approach developed for multivariate panel data by João et al. (2023).

In the empirical application to daily returns of S&P 100 stocks from 2015 to 2024, the dynamic clustering model consistently outperforms static benchmarks based on SIC codes and k-means clustering. The improvement appears to be driven by firms that undergo changes in their business activities or were initially misclassified. These gains are evident in both in-sample and out-of-sample evaluations and are particularly strong during periods of high volatility, elevated dispersion, and large CAPM alpha values. However, the model introduces additional computational demands, not all clusters are economically interpretable and some firms switch frequently between clusters.

Future research could address the issue of frequent switching between clusters by incorporating non-Markovian transitions, which prevent switches if a stock has recently switched. Alternatively, flickering could be reduced by using the non-parametric method of João et al. (2024), which uses a modified version of k-means clustering to ensure temporal stability in clusters. A challenge of the latter is its high computational cost, as it requires running the k-means algorithm for each point in time. Another potential extension is to explore whether including additional explanatory variables improves model performance. Variables such as market capitalization or return on equity could be used to inform the transition probabilities between clusters, an approach shown to enhance performance by João et al. (2023). Lastly, future work could focus on scaling the model to accommodate larger datasets containing more firms or other asset classes.

References

Amisano, G. & Giacomini, R. (2007). Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business & Economic Statistics*, *25*(2), 177-190.

Chesnay, F. & Jondeau, E. (2001). Does correlation between stock returns really increase during turbulent periods? *Economic Notes*, *30*(1), 53-80.

Coval, J., Jurek, J. & Stafford, E. (2009). The economics of structured finance. *Journal of Economic Perspectives*, *23*(1), 3-25.

Creal, D. D., Koopman, S. J. & Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, *28*(5), 777-795.

DeSarbo, W. S. & Cron, W. L. (1988). A maximum likelihood methodology for clusterwise linear regression. *Journal of classification*, *5*, 249-282.

Diebold, F. X. & Mariano, R. S. (2002). Comparing predictive accuracy. *Journal of Business & economic statistics*, *20*(1), 134-144.

Diks, C., Panchenko, V., Sokolinskiy, O. & 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.

Frühwirth-Schnatter, S. (2011). Panel data analysis: a survey on model-based clustering of time series. *Advances in Data Analysis and Classification*, *5*, 251-280.

Giacomini, R. & White, H. (2006). Tests of conditional predictive ability. *Econometrica*, *74*(6), 1545-1578.

Glosten, L. R., Jagannathan, R. & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The journal of finance*, *48*(5), 1779-1801.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the econometric society*, 357-384.

Hansen, P. R., Lunde, A. & Nason, J. M. (2011). The model confidence set. *Econometrica*, *79*(2), 453-497.

João, I. C., Lucas, A., Schaumburg, J. & Schwaab, B. (2023). Dynamic clustering of multivariate panel data. *Journal of Econometrics*, *237*(2), 105281.

João, I. C., Schaumburg, J., Lucas, A. & Schwaab, B. (2024). Dynamic nonparametric clustering of multivariate panel data. *Journal of Financial Econometrics*, *22*(2), 335-374.

Li, J., Liao, Z. & Quaedvlieg, R. (2022). Conditional superior predictive ability. *The Review of Economic Studies*, *89*(2), 843-875.

Manner, H. & Reznikova, O. (2012). A survey on time-varying copulas: specification, simulations, and application. *Econometric reviews*, *31*(6), 654-687.

Newey, W. K. & West, K. D. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, 777-787.

Oh, D. H. & Patton, A. J. (2017). Modeling dependence in high dimensions with factor copulas. *Journal of Business & Economic Statistics*, *35*(1), 139-154.

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

Smith, M. S. (2015). Copula modelling of dependence in multivariate time series. *International Journal of Forecasting*, *31*(3), 815-833. (ID: 271676)

Zimmer, D. M. (2012). The role of copulas in the housing crisis. *Review of Economics and Statistics*, *94*(2), 607-620.

Appendix A
List of Firms Used in the Empirical Analysis

Table 2
Summary of firms in the S&P 100

Ticker	Name	SIC	Ticker	Name	SIC	Ticker	Name	SIC
AAPL	Apple	35	DHR	Danaher	50	MS	Morgan Stanley	62
ABBV	Abbvie	28	DIS	Disney	73	MSFT	Microsoft	73
ABT	Abbott Lab.	50	DUK	Duke Energy	49	NEE	Nextera Energy	49
ACN	Accenture	67	EMR	Emerson	35	NFLX	Netflix	78
ADBE	Adobe	73	F	Ford	37	NKE	Nike	30
AIG	Ame Inter	63	FDX	Fedex	45	NVDA	Nvidia	36
AMD	Adv Micro Dev	36	GD	Gen Dynamics	37	ORCL	Oracle	73
AMGN	Amgen	28	GE	Gen Electric	35	PEP	Pepsico	20
AMT	American Tower	48	GILD	Gilead	28	PFE	Pfizer	28
AMZN	Amazon	73	GM	General Motors	37	PG	Procter Gamble	28
AVGO	Broadcom	36	GOOG	Alphabet	73	PM	Philip Morris	21
AXP	Amex	60	GOOGL	Alphabet	73	QCOM	Qualcomm	36
BA	Boeing	37	GS	Goldman Sachs	62	RTX	RTX	37
BAC	Bank of Am	60	HD	Home Depot	52	SBUX	Starbucks	58
BH	Biglari Holdings	58	HON	Honeywell Int	50	SCHW	Schwab Charles	62
BK	Bank of NY	60	IBM	IBM	73	SO	Southern	49
BKNG	Booking	73	INTC	Intel	36	SPG	Simon Property	67
BLK	Blackrock	62	INTU	Intuit	73	T	AT&T	48
BMY	Bristol-Myers	28	JNJ	Johnson&J	28	TGT	Target	53
C	Citigroup	60	JPM	Jpmorgan	60	TMO	Thermo Fisher	38
CAT	Caterpillar	35	KO	Coca Cola	20	TMUS	T-Mobile	48
CHTR	Charter Comm	48	LLY	Lilly Eli	28	TSLA	Tesla	37
CL	Colgate Palmo	28	LMT	Lockheed Mar	37	TXN	Texas Instru	36
CMCSA	Comcast	48	LOW	Lowes	52	UNH	Unitedhealth	63
COF	Capital One	60	MA	Mastercard	73	UNP	Union Pacific	40
COP	Conocophillips	13	MCD	Mcdonalds	58	UPS	United Parcel	45
COST	Costco	53	MDLZ	Mondelez Int	20	USB	US Bancorp	60
CRM	Salesforce	73	MDT	Medtronic	38	V	Visa	73
CSCO	Cisco Sys	36	MET	Metlife	63	VZ	Verizon	48
CVS	C V S Health	59	META	Meta	73	WFC	Wells Fargo	60
CVX	Chevron	13	MMM	3M	50	WMT	Walmart	53
D	Dominion En	49	MO	Altria Group	21	XOM	Exxon Mobil	29
DE	Deere	35	MRK	Merck	28			

SIC	Description	Num	SIC	Description	Num	SIC	Description	Num
1	Mining, construct.	2	4	Transprt, comm's	13	7	Services	15
2	Manuf: food, furn.	16	5	Trade	13	9	Non-classifiable	2
3	Manuf: elec, mach	20	6	Finance, Ins	17	Total		98

Note: This table reports the ticker symbol, company name, and the first two digits of the SIC code for the firms included in the empirical analysis. The sample consists of firms that were constituents of the S&P 100 index as of December 31, 2024, and that were continuously traded throughout the full sample period (2015–2024). Firms that underwent mergers or splits during the sample period are excluded. SIC codes correspond to those assigned as of the midpoint of the sample period (December 31, 2019). For firms that changed their ticker symbol or name, the most recent identifiers are reported.

Appendix B
Estimation procedure

Due to the computational complexity of jointly estimating dynamic copula parameters and time-varying group assignments, a three-stage estimation procedure is adopted. In the first stage, the initial cluster assignments and copula parameters are estimated following Oh & Patton (2023). Static group assignments $\hat{\Gamma}_1$ are obtained by applying k-means clustering to the full sample using a misspecified static Gaussian copula. Conditional on these initial clusters, the copula parameters $\boldsymbol{\psi} = [\omega_1^M, \ldots, \omega_G^M, \omega_1^C, \ldots, \omega_G^C, \alpha^M, \beta^M, \alpha^C, \beta^C, \nu, \zeta]'$ are estimated by maximizing the log-likelihood of the skewed t copula:

$$\hat{\boldsymbol{\psi}} = \arg\max_{\boldsymbol{\psi}} \sum_{t=1}^{T} \log \mathbf{c}_{\text{Skew } t,t}(\mathbf{u}_t; \boldsymbol{\psi} \mid \hat{\Gamma}_1). \tag{9}$$

The skewed t copula nests both the Gaussian and symmetric t copulas as special cases. In both the simulation study and the empirical analysis, all three copulas are estimated for comparison.

In the second stage, time-varying cluster assignments are estimated. For each time period t, the filtered probabilities $\tau_{ig,t|t}$ are updated using the forward algorithm in Eqs. (4)-(5), again based on the misspecified static Gaussian copula. Each stock is then assigned to the cluster with the highest posterior probability, yielding the current group assignment $\hat{\Gamma}_t$.

In the third stage, the dynamic factors loadings are updated using the Generalized Autoregressive Score dynamics, based on the copula parameters obtained in the first stage. GAS adjusts the loadings in response to new information by leveraging the gradient of the log-likelihood function of the conditional copula. The update equations take the form:

$$\lambda_{g,t+1}^M = \omega_g^M + \alpha^M \frac{\partial \log \mathbf{c}_{\text{Skew}t,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \lambda_{g,t}^M} + \beta^M \lambda_{g,t}^M, \quad \text{for } g = 1, \ldots, G \tag{10}$$

$$\lambda_{g,t+1}^C = \omega_g^C + \alpha^C \frac{\partial \log \mathbf{c}_{\text{Skew}t,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)}{\partial \lambda_{g,t}^C} + \beta^C \lambda_{g,t}^C, \quad \text{for } g = 1, \ldots, G \tag{11}$$

where $\mathbf{x}_t = T_{\text{skew}}^{-1}(\mathbf{u}_t; \nu, \zeta)$, and $\mathbf{c}_{\text{Skew}t,t}(\mathbf{x}_t; \mathbf{R}_t, \nu, \zeta)$ denotes the conditional skewed t copula density. Finally, given the current factor loadings and group assignments, the log-likelihood of each observation is computed.

With respect to the remaining parameters, the number of clusters G must be chosen. Although the Akaike Information Criterion (AIC) is commonly used, it may overestimate the number of clusters (Frühwirth-Schnatter, 2011). Therefore, the optimal number of clusters is also validated based on out-of-sample forecasting performance. In theory, the number of clusters could vary over time if all firms were to transition out of a given cluster, implying $G = G_{max}$. However, to preview the results, this has not occurred in practice, likely because the initial static clustering provides a sufficiently accurate starting point. Moreover, maintaining a fixed number of clusters over time is common in literature and appears to be a reasonable assumption (Frühwirth-Schnatter, 2011; João et al., 2023). The decay parameter δ in Eq. (1), which controls the dynamics of cluster transitions, is selected by evaluating model performance over a grid of values.

Appendix C
Summary Statistics and Marginal Model Results

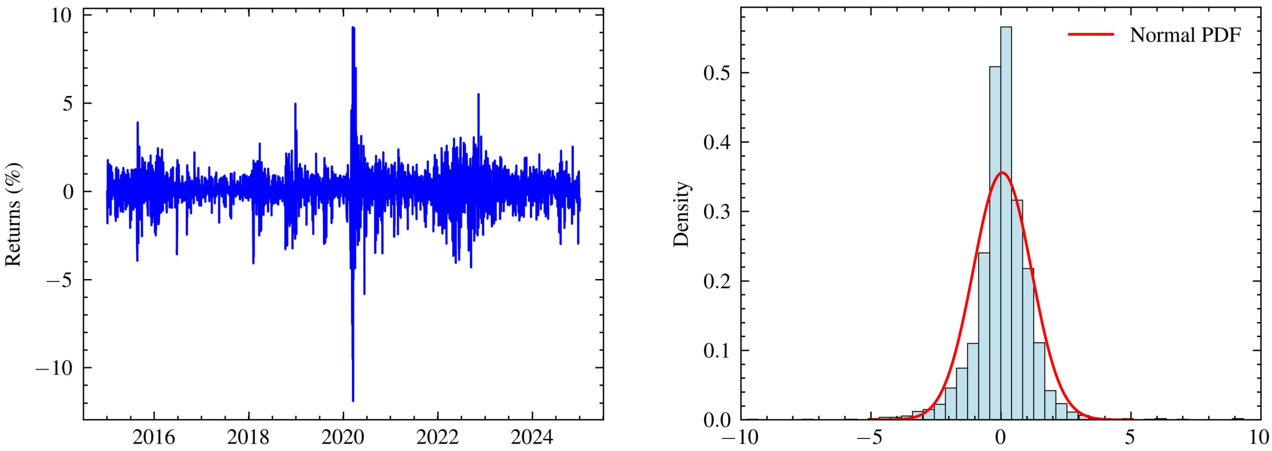

Figure 1. Distribution of S&P 100 returns. The left panel displays the daily value-weighted market returns over the full sample period from 2015 to 2024. The right panel shows the distribution of the returns, along with a fitted Gaussian distribution.

Table 3
Summary statistics for the marginal model

	Cross-sectional distribution					
	Mean	5%	25%	Median	75%	95%
Panel A: Marginal moments						
Mean	0.001	0.000	0.000	0.001	0.001	0.001
Std	0.018	0.012	0.015	0.017	0.020	0.026
Skewness	-0.028	-0.786	-0.228	0.079	0.237	0.819
Kurtosis	15.000	8.201	10.676	13.352	17.703	28.429
Panel B: Marginal model parameters						
Constant	0.001	0.000	0.000	0.001	0.001	0.001
AR(1)	-0.020	-0.052	-0.038	-0.019	-0.004	0.014
$\varpi \times 10^4$	0.009	0.002	0.004	0.006	0.011	0.023
α	0.034	0.001	0.016	0.028	0.043	0.085
κ	0.091	0.016	0.062	0.089	0.119	0.161
β	0.893	0.799	0.871	0.897	0.929	0.957
ξ	4.558	3.555	3.950	4.415	4.941	6.128
ψ	-0.035	-0.092	-0.063	-0.036	-0.010	0.025
Panel C: Correlations of standardized residuals						
Pearson	0.288	0.127	0.214	0.274	0.347	0.484
Spearman	0.330	0.152	0.251	0.316	0.398	0.535

Note: This table reports the cross-sectional distribution of summary statistics from 98 daily return series spanning January 2, 2015 to December 31, 2024. Panel A reports the distribution of the first four moments of the returns, Panel B shows the estimated parameters from the marginal models, and Panel C provides a summary of the pairwise correlations among the standardized residuals.

Appendix D
Comparison between Static and Dynamic Clusters

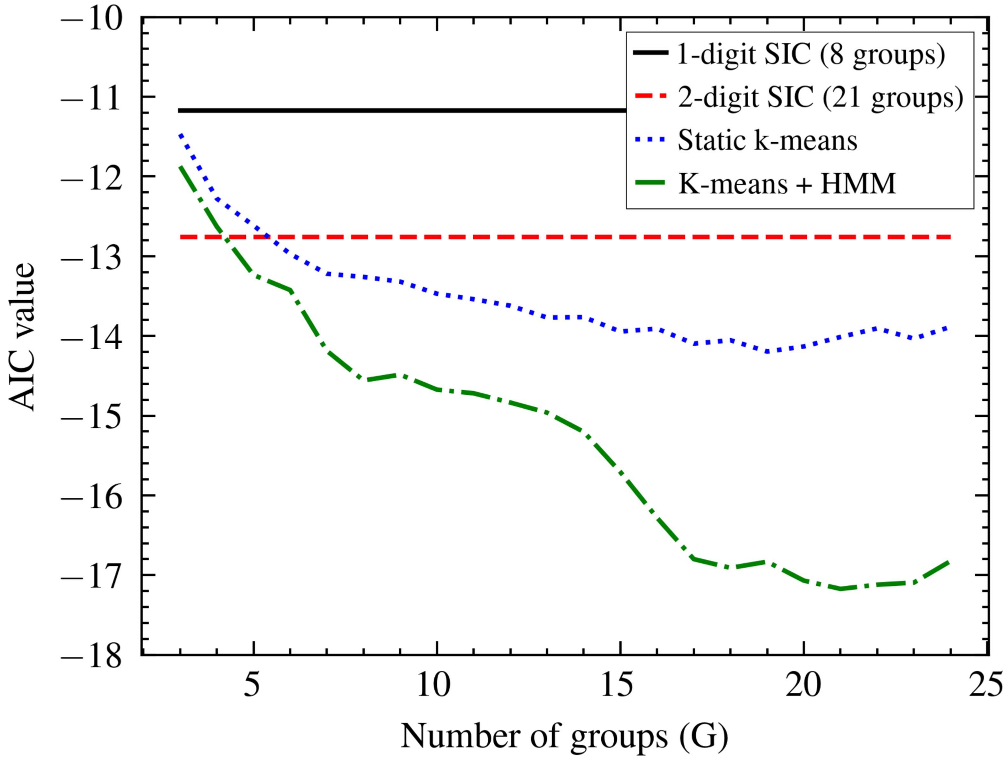

Figure 2. AIC values as a function of the number of groups (G) for static clustering and dynamic clustering combining k-means with a hidden Markov model. For comparison, AIC values corresponding to the 1-digit and 2-digit SIC-based groupings, comprising 8 and 21 groups, respectively, are also shown. Lower AIC values indicate a better model fit. The y-axis is scaled by 10^{-4}.

Table 4
Estimated group assignments with static clustering

Group	Ticker	Name	SIC	Group	Ticker	Name	SIC
1	ABBV	Abbvie	28	8	CAT	Caterpillar	35
	ABT	Abbott Lab.	28		EMR	Emerson	36
	AMGN	Amgen	28		FDX	Fedex	45
	BMY	Bristol-Myers	28		UNP	Union Pacific	40
	GILD	Gilead	28		UPS	United Parcel	45
	JNJ	Johnson&J	28				
	LLY	Lilly Eli	28	9	AMZN	Amazon	73
	MDT	Medtronic	38		BKNG	Booking	73
	MRK	Merck	28		META	Meta	73
	PFE	Pfizer	28		NFLX	Netflix	78
	TMO	Thermo Fisher	38				
	UNH	Unitedhealth	63	10	ACN	Accenture	73
					CSCO	Cisco Sys	36
2	BAC	Bank of Am	62		IBM	IBM	73
	BK	Bank of NY	60		ORCL	Oracle	73
	C	Citigroup	62				
	COF	Capital One	60	11	COST	Costco	53
	GS	Goldman Sachs	62		CVS	C V S Health	59
	JPM	Jpmorgan	60		TGT	Target	53
	MET	Metlife	63		WMT	WalMart	53
	MS	Morgan Stanley	60				
	SCHW	Schwab Charles	62	12	D	Dominion En	49
	USB	US Bancorp	60		DUK	Duke Energy	49
	WFC	Wells Fargo	60		NEE	Nextera Energy	49
					SO	Southern	49
3	AMT	American Tower	67				
	CL	Colgate Palmo	28				
	KO	Coca Cola	20	13	COP	Conocophillips	13
	MDLZ	Mondelez Int	20		CVX	Chevron	29
	MO	Altria Group	21		XOM	Exxon Mobil	29
	PEP	Pepsico	20				
	PG	Procter Gamble	28	14	AIG	Ame Inter	63
	PM	Philip Morris	21		AXP	Amex	61
	SPG	Simon Property	67		GE	Gen Electric	35
4	BA	Boeing	37	15	BH	Biglari Holdings	58
	GD	Gen Dynamics	37		TMUS	T-Mobile	48
	HON	Honeywell Int	37		TSLA	Tesla	37
	LMT	Lockheed Mar	37				
	MMM	3M	38	16	MCD	Mcdonalds	58
	RTX	RTX	37		NKE	Nike	30
					SBUX	Starbucks	58
5	ADBE	Adobe	73				
	CRM	Salesforce	73	17	F	Ford	37
	INTU	Intuit	73		GM	General Motors	37
	MA	Mastercard	73				
	MSFT	Microsoft	73	18	GOOG	Google	73
	V	Visa	73		GOOGL	Google	73
6	AAPL	Apple	35	19	DE	Deere	35
	AMD	Adv Micro Dev	36		INTC	Intel	36
	AVGO	Broadcom	36				
	NVDA	Nvidia	36	20	BLK	Blackrock	62
	QCOM	Qualcomm	36		DHR	Danaher	38
	TXN	Texas Instru	36				
				21	HD	Home Depot	52
7	CHTR	Charter Comm	48		LOW	Lowes	52
	CMCSA	Comcast	48				
	DIS	Disney	48				
	T	AT&T	48				
	VZ	Verizon	48				

Note: This table reports the estimated static group assignments for 21 groups, as determined by the AIC-selected optimal number of clusters. For firms that have changed their ticker symbol or name, the most recent identifiers are shown. Groups are ordered by the number of stocks.

Table 5
Estimated final group assignments with dynamic clusters

Group	Ticker	Name	SIC	Group	Ticker	Name	SIC
1	ABBV	Abbvie	28	9	AAPL	Apple	35
	ABT	Abbott Lab.	28		AMZN	Amazon	73
	AMGN	Amgen	28		COST	Costco	53
	BMY	Bristol-Myers	28		LLY	Lilly Eli	28
	GILD	Gilead	28		META	Meta	73
	JNJ	Johnson&J	28		NFLX	Netflix	78
	KO	Coca Cola	20				
	MRK	Merck	28	10	ACN	Accenture	73
	PEP	Pepsico	20		CSCO	Cisco Sys	36
	PFE	Pfizer	28		HON	Honeywell Int	37
					IBM	IBM	73
2	BAC	Bank of Am	62		MMM	3M	38
	BK	Bank of NY	60				
	C	Citigroup	62	11	CHTR	Charter Comm	48
	COF	Capital One	60		CMCSA	Comcast	48
	GS	Goldman Sachs	62		CVS	C V S Health	59
	JPM	JPMorgan	60		MDLZ	Mondelez Int	20
	MET	Metlife	63		MDT	Medtronic	38
	MS	Morgan Stanley	60		UNH	Unitedhealth	63
	USB	US Bancorp	60				
	WFC	Wells Fargo	60	12	D	Dominion En	49
					DUK	Duke Energy	49
3	AMT	American Tower	62		NEE	Nextera Energy	49
	CL	Colgate Palmo	60		SO	Southern	49
	DIS	Disney	62				
	PG	Proctor Gamble	60	13	COP	Conocophilips	13
	PM	Philip Morris	62		CVX	Chevron	29
	SPG	Simon Property	60		XOM	Exxon Mobil	29
	T	AT&T	63				
	VZ	Verizon	60	14	AIG	Ame Inter	63
	WMT	Walmart	60		AXP	Amex	61
					BA	Boeing	37
4	GD	Gen Dynamics	37				
	GE	Gen Electric	35	15	BH	Biglari Holdings	58
	LMT	Lockheed Mar	37		TMUS	T-Mobile	48
	RTX	RTX	37		TSLA	Tesla	37
5	ADBE	Adobe	73	16	BKNG	Booking	73
	CRM	Salesforce	73		NKE	Nike	30
	INTU	Intuit	73		PCLN	Priceline	73
	MSFT	Microsoft	73		SBUX	Starbucks	58
	ORCL	Oracle	73				
				17	F	Ford	37
6	AMD	Adv Micro Dev	36		GM	General Motors	37
	AVGO	Broadcom	36				
	INTC	Intel	36	18	GOOG	Google	73
	NVDA	Nvidia	36		GOOGL	Google	73
	QCOM	Qualcomm	36				
	TXN	Texas Instru	36	19	DE	Deere	35
					MCD	Mcdonalds	58
7	DHR	Danaher	38				
	TGT	Target	53	20	BLK	Blackrock	62
	TMO	Thermo Fisher	38		MA	Mastercard	73
					SCHW	Schwab Charles	62
8	CAT	Caterpillar	35		UNP	Union Pacific	40
	EMR	Emerson	36		V	Visa	73
	FDX	Fedex	45				
	UPS	United Parcel	45	21	HD	Home Depot	52
					LOW	Lowes	52

Note: This table reports the estimated dynamic group assignments for 21 groups at the end of the sample period. For firms that have changed their ticker symbol or name, the most recent identifiers are shown. Cluster numbers are consistent with the original numbering assigned in the static clustering.

Table 6
Estimation results for the optimal 21 group model

Panel A: Parameter estimation accuracy						
	Gaussian		t		Skew t	
	est.	s.e.	est.	s.e.	est.	s.e.
ω_1^M	0.052	0.003	0.007	0.001	0.006	0.001
ω_2^M	0.105	0.007	0.016	0.003	0.013	0.001
ω_3^M	0.053	0.003	0.007	0.001	0.006	0.001
ω_4^M	0.074	0.005	0.011	0.001	0.009	0.001
ω_5^M	0.083	0.005	0.012	0.002	0.008	0.001
ω_6^M	0.072	0.005	0.010	0.000	0.009	0.001
ω_7^M	0.055	0.004	0.008	0.001	0.007	0.001
ω_8^M	0.081	0.005	0.012	0.001	0.010	0.001
ω_9^M	0.061	0.004	0.009	0.001	0.007	0.001
ω_{10}^M	0.078	0.005	0.011	0.001	0.009	0.001
ω_{11}^M	0.053	0.003	0.008	0.001	0.006	0.001
ω_{12}^M	0.050	0.003	0.007	0.002	0.006	0.001
ω_{13}^M	0.084	0.005	0.012	0.001	0.010	0.001
ω_{14}^M	0.073	0.005	0.011	0.001	0.009	0.001
ω_{15}^M	0.039	0.002	0.005	0.001	0.005	0.001
ω_{16}^M	0.064	0.004	0.009	0.002	0.008	0.001
ω_{17}^M	0.105	0.007	0.016	0.001	0.013	0.002
ω_{18}^M	0.520	0.032	0.076	0.002	0.063	0.008
ω_{19}^M	0.063	0.004	0.011	0.002	0.009	0.001
ω_{20}^M	0.078	0.005	0.009	0.001	0.014	0.002
ω_{21}^M	0.118	0.007	0.017	0.003	0.008	0.001
ω_1^C	0.003	0.001	0.002	0.001	0.002	0.000
ω_2^C	0.005	0.001	0.002	0.001	0.004	0.001
ω_3^C	0.003	0.001	0.003	0.001	0.003	0.000
ω_4^C	0.003	0.001	0.002	0.001	0.002	0.000
ω_5^C	0.003	0.001	0.003	0.001	0.003	0.000
ω_6^C	0.003	0.001	0.003	0.001	0.003	0.000
ω_7^C	0.002	0.001	0.002	0.001	0.002	0.000
ω_8^C	0.002	0.001	0.002	0.001	0.002	0.000
ω_9^C	0.002	0.001	0.003	0.001	0.003	0.000
ω_{10}^C	0.001	0.000	0.002	0.001	0.002	0.000

Table 6
Estimation results for the optimal 21 group model, continued

Panel A: Parameter estimation accuracy

	Gaussian		t		Skew t	
	est.	s.e.	est.	s.e.	est.	s.e.
ω_{11}^C	0.002	0.001	0.002	0.001	0.002	0.000
ω_{12}^C	0.007	0.001	0.006	0.002	0.006	0.001
ω_{13}^C	0.007	0.001	0.007	0.002	0.007	0.001
ω_{14}^C	0.002	0.001	0.002	0.001	0.002	0.000
ω_{15}^C	0.001	0.001	0.001	0.001	0.000	0.000
ω_{16}^C	0.002	0.001	0.002	0.001	0.002	0.000
ω_{17}^C	0.006	0.001	0.005	0.001	0.001	0.001
ω_{18}^C	0.037	0.006	0.034	0.007	0.034	0.004
ω_{19}^C	0.002	0.001	0.001	0.001	0.001	0.001
ω_{20}^C	0.000	0.000	0.000	0.000	0.001	0.001
ω_{21}^C	0.007	0.001	0.006	0.002	0.006	0.002
α^M	0.033	0.001	0.011	0.003	0.010	0.001
β^M	0.917	0.005	0.987	0.008	0.990	0.002
α^C	0.006	0.001	0.008	0.002	0.008	0.001
β^C	0.995	0.005	0.995	0.008	0.996	0.001
ν			0.034	0.003	0.034	0.001
ζ					-0.398	0.001

Panel B: Estimation details

	Gaussian	t	Skew t
$\log \mathcal{L}$	86643.05	89266.23	87968.92
AIC	-173194	-178438	-175841
BIC	-172926	-178164	-175562
Time (clustering) (hrs)	1.23	1.23	1.23
Time (copula) (hrs)	2.71	2.98	2.68
EM iterations	95.54	95.54	95.54

Note: This table reports the estimated parameters and standard errors for the multi-factor copula model with dynamic group assignments, estimated via a Hidden Markov Model. Results are presented for the Gaussian, t, and skew-t copulas. The model was estimated on the full sample using 21 groups, selected as optimal based on the Akaike Information Criterion (AIC). Panel B reports model fit measures, computational time, and the number of EM iterations. All estimations were performed on a machine with an Apple M1 processor (8 cores).

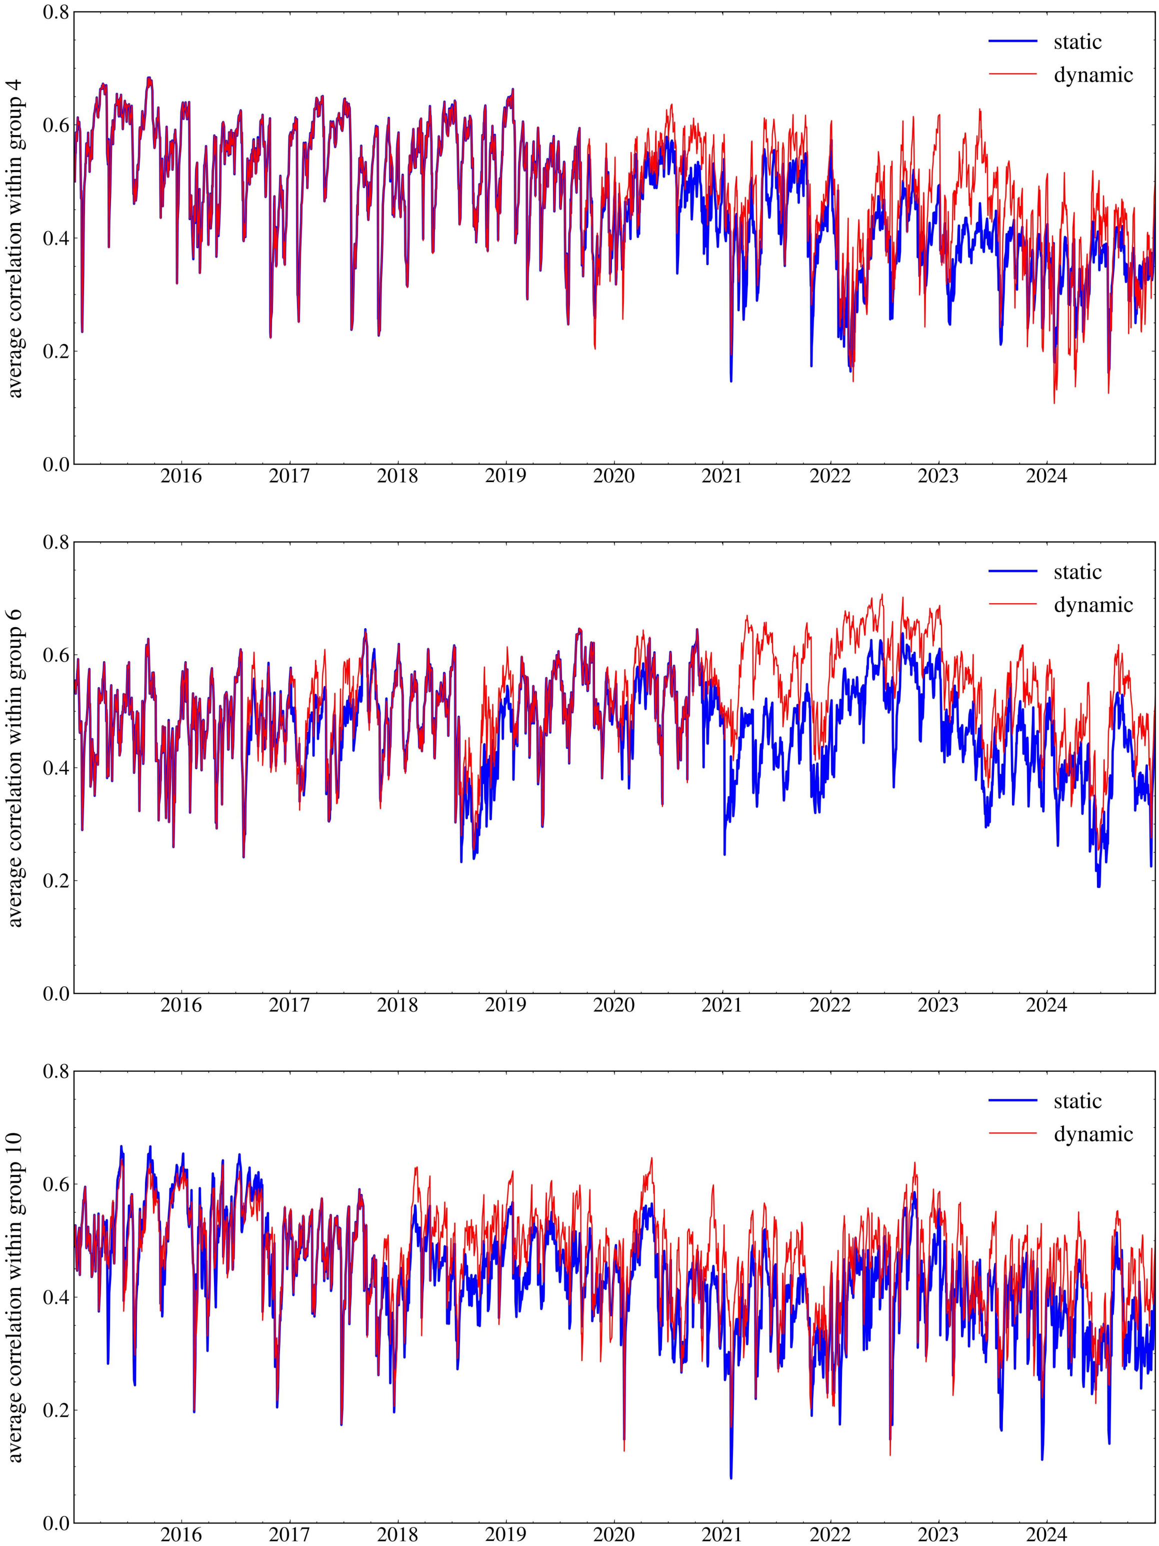

Figure 3. Model-implied rank correlations over the full sample period. The upper panel displays correlations for group 4 (from Table 4), comparing static clusters with dynamic clusters from the Markov-switching model. The middle panel and lower shows the same comparison for group 6 and 10, respectively. The results are obtained from the model with GAS dynamics and the Gaussian copula.

Appendix E
Transition Dynamics

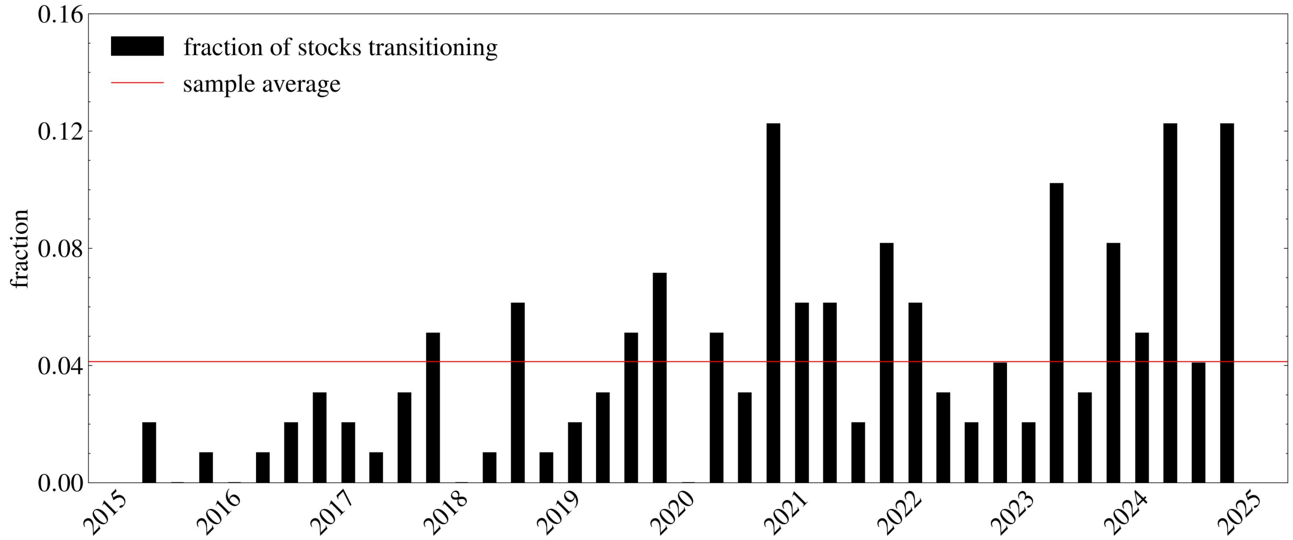

Figure 4. Timing of cluster transitions. The black bars indicate the fraction of stocks that are estimated to change groups for each quarter between 2015Q1 and 2024Q4. The red line reports the average transition frequency over the full sample.

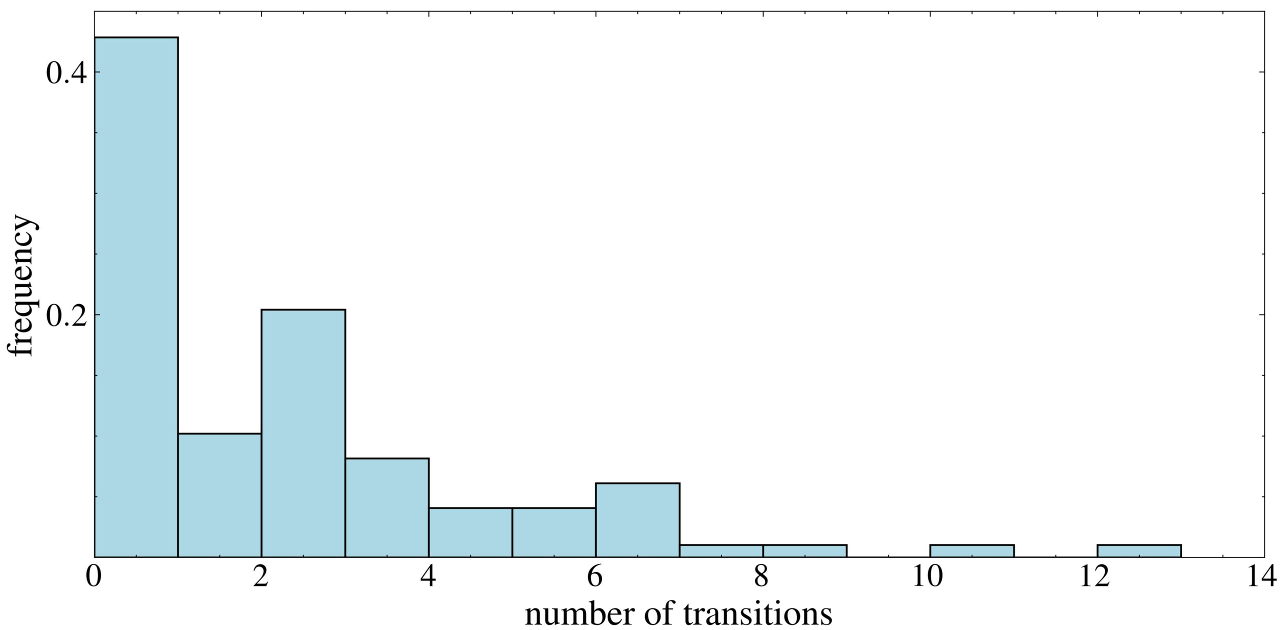

Figure 5. Histogram of cluster transitions counts per stock.

Appendix F
Diebold-Mariano test results

Table 7
Comparison of different copulas for static clustering

	Static vs. GAS			Copula shape		
	Gaussian	t	skew$t$	G vs. t	G vs. skew$t$	t vs. skew$t$
SIC 1-digit	3.170	3.704	27.771	12.223	-3.025	-18.945
SIC 2-digit	1.905	3.728	31.834	12.748	10.320	-8.435
3 groups	0.327	1.632	24.721	11.212	9.305	-9.1612
6 groups	3.017	4.294	30.483	12.466	11.088	-6.035
9 groups	3.807	5.172	7.481	12.741	6.647	-17.572
12 groups	3.574	5.117	11.443	13.057	7.085	-17.954
15 groups	3.058	4.528	22.474	12.276	-5.912	-17.475
18 groups	5.412	6.968	34.708	12.238	10.625	-7.451
21 groups	4.218	6.680	35.023	12.582	9.942	-7.051
24 groups	3.532	6.140	20.083	12.996	-5.367	-20.002

Note: This table reports Diebold-Mariano t statistics for pairwise comparisons of models with static clusters using their out-of-sample log-likelihood. The left panel compares models assuming static factor loadings with those using GAS dynamics, for a Gaussian, t and skew-t copula 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 labelled "A vs. B," a positive t-statistic implies that model B outperforms model A, whereas a negative t-statistic suggests that model A performs better than model B.

Table 8
Comparison of different copula's for dynamic clustering

	Static vs. GAS			Copula shape		
	Gaussian	t	skew$t$	G vs. t	G vs. skew$t$	t vs. skew$t$
SIC 1-digit	11.497	11.404	30.732	11.858	-8.489	-19.712
SIC 2-digit	5.601	5.605	21.014	12.843	-7.388	-18.094
3 groups	5.129	4.852	21.643	11.879	-5.861	-18.594
6 groups	5.833	5.249	18.869	12.299	-7.487	-19.124
9 groups	6.058	5.424	13.251	12.900	-9.863	-14.459
12 groups	5.121	4.140	20.386	13.246	-6.182	-18.543
15 groups	4.974	4.618	25.893	12.363	-6.521	-17.847
18 groups	5.841	6.795	32.697	12.290	-7.342	-17.624
21 groups	6.767	7.261	31.105	12.906	-6.699	-21.958
24 groups	5.489	6.434	29.453	11.358	-7.231	-21.762

Note: This table reports Diebold-Mariano t statistics for pairwise comparisons of models with dynamic HMM clusters using their out-of-sample log-likelihood. The left panel compares models assuming static factor loadings with those using GAS dynamics, for a Gaussian, t and skew-t copula 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 implies that model B outperforms model A, whereas a negative t-statistic suggests that model A performs better than model B.

Appendix G
Giocomini & White and CSPA test results

Table 9
Economic determinants of forecast performance

	Dynamic vs. SIC 2-digit				Dynamic vs. static k-means			
Intercept	1.716	1.716	1.716	1.716	1.349	1.349	1.349	1.349
(s.e.)	(0.107)	(0.107)	(0.107)	(0.106)	(0.112)	(0.111)	(0.103)	(0.104)
[t-stat]	[16.106]	[16.106]	[16.058]	[16.181]	[12.094]	[12.182]	[13.103]	[12.925]
VIX	-0.056			-0.142	0.226			-0.016
(s.e.)	(0.082)			(0.099)	(0.104)			(0.124)
[t-stat]	[-0.683]			[-1.437]	[2.168]			[-1.039]
Dispersion		0.057		0.235		0.657		0.328
(s.e.)		0.090		(0.126)		(0.155)		(0.130)
[t-stat]		[0.631]		[1.872]		[4.243]		[2.528]
Abs. alpha			-0.096	-0.188			0.895	0.750
(s.e.)			(0.092)	(0.105)			(0.118)	(0.124)
[t-stat]			[-1.038]	[-1.791]			[7.589]	[6.043]
R^2 (%)	0.030	0.030	0.086	0.368	0.393	3.325	6.164	6.599
GW p-value$_{\text{ALL}}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
GW p-value$_{\text{SLOPES}}$	0.495	0.528	0.299	0.403	0.000	0.000	0.000	0.000
CSPA p-value	0.000	0.000	0.000	-	0.000	0.000	0.000	-

Note: This table reports the results of the Giacomini & White (2006) tests and the Li et al. (2022) (CSPA) tests. The benchmark model is the model with 21 time-varying clusters with GAS dynamics and a student t copula. The conditioning variables are the VIX index, the dispersion (cross-sectional standard deviation of returns) and the absolute value of the cross-sectional average CAPM alpha. For the GW tests, conditioning variables are standardized to ensure the comparability of the test statistics.