RESEARCH DIVISION

Working Paper Series

International Stock Comovements with Endogenous Clusters

Laura Coroneo,
Laura E. Jackson
and
Michael T. Owyang

Working Paper 2018-038B
https://doi.org/10.20955/wp.2018.038

December 2018

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
P.O. Box 442
St. Louis, MO 63166

International Stock Comovements with Endogenous Clusters*

Laura Coroneo — University of York
Laura E. Jackson — Bentley University
Michael T. Owyang — Federal Reserve Bank of St. Louis

October 24, 2018

Abstract

We use an endogenous cluster factor model to examine international stock return comovements of country-industry portfolios. Our model allows country-industry portfolio comovements to be driven by a global and a cluster component, with the cluster membership endogenously determined. Results indicate that country-industry portfolios tend to cluster mainly within geographical areas that can include one or more countries. The cluster component was the main driver of country-industry portfolio returns for most of the sample, except from mid-2000 to mid-2010s when the global component had a more prominent role. At the end of the sample, a large cluster among European countries emerges.

Keywords: diversification, risk, international financial markets, clustered factor model.

JEL classification: C38, G15

*Dongna Zhang and Hannah G. Shell provided research assistance. The authors benefited from conversations with Chris Otrok. The views expressed here do not represent those of the Federal Reserve Bank of St. Louis or the Federal Reserve System. Laura Coroneo gratefully acknowledges the support of the ESRC grant ES/K001345/1.

1 Introduction

Understanding the determinants of international stock returns has important implications for the design of portfolio diversification strategies. A large literature in international finance focusses on understanding the gains from international portfolio diversification—especially the role of country and industry factors. While the classical result is that it is better to diversify across countries rather than across industries (see Lessard 1974, Heston and Rouwenhorst 1994, Griffin and Karolyi 1998), more recent evidence suggests that industry factors are gaining importance (see Cavaglia, Brightman and Aked 2000, Baele and Inghelbrecht 2009). In addition, due to global financial market integration, international stock returns are increasingly driven by global, rather than local factors (see Pukthuanthong and Roll 2009, Brooks and Del Negro 2006, Eiling, Gerard, Hillion and de Roon 2012).

Following Heston and Rouwenhorst (1994), the standard approach to international portfolio design assumes that the structure of comovement across international equity returns is known. This assumption raises the issue of selecting the level of granularity of the underlying factors. Roll (1992) suggests that industries should be grouped into a relatively small number of broad categories, and Brooks and Del Negro (2005) find that regional effects are stronger than country effects. However, there is no clear consensus about how regional or sectoral factors should be specified. In addition, there is growing evidence that the factor structure of international equity returns has been changing over time, see Eiling et al. (2012), Bekaert, Hodrick and Zhang (2009) and Brooks and Del Negro (2004). An alternative approach that does not require specifying a particular factor structure is the Asset Pricing Theory model of Conner and Korajczyk (1986). However, this approach does not easily allow to construct replicating portfolios as all the factors load on all the country-industry portfolios.

In this paper, we examine international stock return comovements of country-industry portfolios using a factor model with endogenously-determined groupings or *clusters*. Our model allows fluctuations in country-industry portfolio excess returns to be decomposed into three components:

a pervasive component driven by a global factor that loads on all country-industry portfolios, a less pervasive component driven by a cluster factor that loads on a subset of country-industry portfolios, and an idiosyncratic component specific to each country-industry portfolio. Following the Arbitrage Pricing Theory literature, we assume that the global and cluster factors are latent and that the idiosyncratic comovement is not priced. This implies that, for portfolios not in the same cluster, comovement is driven only by the global factor.

A crucial feature of our model is that the cluster membership is endogenously determined. We allow a cluster factor to be common to portfolios in one or more countries and/or one or more industries. If a cluster factor drives all the portfolios in only one country (industry), then it coincides with the country (industry) factor used in the literature. Similarly, if a cluster factor drives all the portfolios in a number of close countries (industries) than it is a regional (sectoral) factor. Cluster membership is determined by cluster indicator that, following Fruhwirth-Schnatter and Kaufmann (2008) and Francis, Owyang and Savascin (2017b), can be estimated using a multinomial hierarchical prior that takes into account covariates that could lead to comovements.

Using monthly excess returns on country-industry portfolios for 23 countries and 25 industries from January 1980 to December 2016, we estimate our factor model with endogenous clusters using Bayesian techniques. To allow for time-variation in the factor structure, following Bekaert et al. (2009), we re-estimate the model every 2.5 years using a window of 5 years of data. Results indicate that country-industry portfolios tend to cluster mainly within geographical areas that can include one or more countries. For the full-sample, most clusters include a diverse group of industries, thus suggesting that within-country comovement is more prominent than within-industry, across-country, comovement. This indicates greater potential benefits from diversifying across geographical areas rather than across sectors.

The rolling-window results highlight the effects of globalization throughout out our sample. The cluster component was the main driver of country-industry portfolio returns for most of the sample, except from mid-2000 to mid-2010s when the global component had a more prominent role. At the end of the sample, a large cluster composed of European countries emerges, cluster membership appears to be broadly more influential once again, and the importance of the global factor is diminished. Additionally, our endogenous cluster model is more successful overall at explaining the features of cross-portfolio comovement than alternatives previously considered in the literature.

Our result on the emergence of a large European cluster in the last part of the sample is related to the growing literature on the Euro Area financial market integration that suggests that in the Euro Area diversification over industries yields more efficient portfolios than diversification over countries, see Hardouvelis, Malliaropulos and Priestley (2007), Cappiello, Kadareja and Manganelli (2010) and Moerman (2008). This paper is also related to the growing literature on endogenous clustering recently used in time series models to identify state and national recessions in the U.S. (Hamilton and Owyang 2012); state and national housing contractions (Hernández-Murillo, Owyang and Rubio 2017); and country and global downturns (Francis, Owyang and Soques 2017a).

The paper is organised as follows. Section 2 describes the endogenous cluster factor model for return comovements of country-industry portfolios. Section 3 describes the data, the estimation procedure and the rolling window estimation. Section 4 describes the estimation results both on the full sample and using the rolling window estimation. Section 5 contains results comparing our endogenous cluster model with two alternatives considered in the literature, both in terms of in-sample variance decompositions and out-of-sample minimum variance portfolio allocation. Section 6 discusses our choice for the number of clusters and considers some alternatives. Finally, Section 7 concludes.

2 A Model of Portfolio Comovements

Consider a panel of value-weighted portfolios constructed for C countries and I industries. Our objective is to model the common movements of the excess returns of these portfolios both across countries and across industries. Let R_{cit} represent the period$-t$ excess return of the portfolio for industry i in country c. We assume that fluctuations in R_{cit} can be decomposed into three components: a pervasive component driven by a global factor G_t that loads on all country-industry portfolios, a less pervasive component driven by a cluster factor F_{kt} that loads on a subset of country-industry portfolios, and an idiosyncratic component ϵ_{cit}. Following the Arbitrage Pricing Theory literature, we assume that the global and cluster factors are latent. We also assume that there exists K unique cluster factors, with each country-industry portfolio belonging to a single cluster, $k \in 1, \ldots, K$. Define $\gamma_{ci}^{k} \in \{0, 1\}$, a cluster indicator that takes on a value of 1 when the country c, industry i portfolio belongs to cluster k and 0 otherwise. The assumption that a country-industry portfolio is uniquely associated with a single cluster implies that $\sum_{k=1}^{K} \gamma_{ci}^{k} = 1$.[1]

The excess return of the country-c–industry-i portfolio can then be written as

$$R_{cit} = \overline{R}_{ci} + b_{ci}^{G} G_t + \sum_{k=1}^{K} \gamma_{ci}^{k} b_{ci}^{k} F_{kt} + \epsilon_{cit} \tag{1}$$

where $\overline{R}_{ci}$ is the expected excess return for the country-c–industry-i portfolio; $\epsilon_{cit} \sim N(0, \sigma_{ci}^2)$; $E[\epsilon_{cit}' \epsilon_{dit}] = 0$ for $c \neq d$; and $E[\epsilon_{cit}' \epsilon_{cjt}] = 0$ for $i \neq j$. We further assume that $\sum_{c=1}^{C} \sum_{i=1}^{I} \gamma_{ci}^{k} > 1$ for all k, which requires all the cluster factors to load on at least two series.

We assume that all the factors, both global and cluster factors, evolve as independent AR(1) processes. Collecting the factors $\mathbf{F}_t = [G_t, F_{1t}, \ldots, F_{Kt}]'$, we have

$$\mathbf{F}_t = \Phi \mathbf{F}_{t-1} + \mathbf{e}_t \tag{2}$$

where Φ is diagonal with elements given by $[\phi_G, \phi_1, \ldots, \phi_K]$; the innovations to the factor processes are $\mathbf{e}_t \sim iidN(\mathbf{0}, \mathbf{I}_{K+1})$; and $E\left[e_{mt}' \epsilon_{cit}\right] = 0$ for all m.

The model in (1)-(2) has a number of implications for the comovements between country-industry portfolios. First, $E[\epsilon_{cit}' \epsilon_{dit}] = 0$ and $E[\epsilon_{cit}' \epsilon_{cjt}] = 0$ imply that comovements across portfolios are a product of the factor structure as idiosyncratic comovement is not priced. Second, for portfolios not in the same cluster, comovement is driven only by the global factor. Third, the assumption $\sum_{c=1}^{C} \sum_{i=1}^{I} \gamma_{ci}^{k} > 1$ implies that no portfolio is subject to purely idiosyncratic and global

[1] We make this assumption to easily allow to construct replicating portfolios and to compare with standard approaches, but it is straightforward to relax.

fluctuations. Thus, no cluster can contain only one portfolio, otherwise the cluster factor would not be identified separately from idiosyncratic fluctuations unique to that portfolio.

2.1 Relation to the Current Literature

The model in (1)-(2) has some similarity to other models used in the literature. One can interpret our model as a more flexible version of Kose, Otrok and Whiteman (2003), who use a hierarchical model with global, regional, and country factors. In their model and other models that followed, the regions (and countries) are defined ex ante. This is equivalent to placing a point prior on the cluster indicator, γ_{ci}^{k}—in effect, pre-allocating portfolios to particular clusters.

For example, suppose that we believe that all of the portfolio correlation is generated within-industry, across countries. In our model, this occurs when, for a given $k = 1, ..., I$, we have that $\gamma_{ci}^{k} = 1$ for all c and for $k = i$. That is, all of the portfolios associate with industry i are collected into the same cluster $k = i$, regardless of country. Thus, F_{kt} behaves as an industry factor, inducing comovements for industry i across all countries. Of course, the loadings for each country may differ, affecting the share of the variance of R_{cit} explained by its industry factor.

In the same way, our endogenous cluster model allows all the portfolios in a country to be grouped across industries. This happens when, for a given $k = 1, ..., C$, we have that $\gamma_{ci}^{k} = 1$ for all i and for $k = c$, where F_{kt} behaves as a country factor, inducing comovements for country c across all industries. In general, our model also allows for regional factors, for example, if, for a given k, we have that $\gamma_{ci}^{k} = \gamma_{di}^{k} = 1$ for all i and some $c \neq d$. In this case, the cluster factor F_{kt} induces comovements in all industries across countries c and d. It is straightforward to extend this logic to more than two countries.

The preceding discussion makes it obvious that in our model the degree of cross-country, cross-industry comovements depends critically on the value of γ_{ci}^{k}. Starting from Lessard (1974) and Heston and Rouwenhorst (1994), the standard approach is to set ex ante the value of γ_{ci}^{k} based on country or industry classification. Roll (1992) suggests that industries should be grouped into a relatively small number of "sufficiently informative industry measurements". More recent studies use regional classification, see Brooks and Del Negro (2005) and Bekaert et al. (2009). However, there is no clear consensus about how regional or sectoral factors should be specified. Francis et al. (2017b) argue that using predetermined clusters can lead to large misspecification about their composition. They propose an algorithm which can estimate the value of γ_{ci}^{k} using a multinomial hierarchical prior that takes into account covariates that could lead to comovements. Alternatively, Ando and Bai (2017) analyze a large number of financial industry stock returns and allow for endogenous clustering based on similar sensitivities to both observable and unobservable factors. The resulting clusters suggest that ex ante classifications based on country, region, industry, or market-specific characteristics are insufficient for explaining heterogeneous behavior of financial markets around the world.

2.2 Endogenous Clusters

In principle, flexible allocation of the county-industry portfolios to different clusters could be implemented as a model selection problem. One could posit alternative cluster memberships, estimate the models, and then choose the model that best fits the data. However, achieving true flexibility across a number of alternatives could be computationally burdensome. Here, we will allow the cluster grouping to be endogenously determined and we allow a cluster factor to be common to portfolios in one or more countries and/or one or more industries. If a cluster factor drives all the portfolios in only one country (industry), then it coincides with the country (industry) factor used in the literature. Similarly, if a cluster factor drives all the portfolios in a number of close countries (industries) than it is a regional (sectoral) factor.

Suppose there exists a vector, $\mathbf{z}_{ci}$, of variables that could influence whether a portfolio for industry i in country c belongs to cluster k. We assess the prior probability that portfolio for industry i in country c belongs to cluster k as

$$Pr[\gamma_{ci}^{k} = 1|\mathbf{z}_{ci}] = \begin{cases} \exp\left(\mathbf{z}_{ci}'\boldsymbol{\alpha}_k\right) / \left[1 + \sum_k \exp\left(\mathbf{z}_{ci}'\boldsymbol{\alpha}_k\right)\right], & k = 1, \ldots, K-1 \\ 1/\left[1 + \sum_k \exp\left(\mathbf{z}_{ci}'\boldsymbol{\alpha}_k\right)\right], & k = K \end{cases} \tag{3}$$

for $c = 1, \ldots, C$ and $i = 1, \ldots, I$, and where we have normalized $\boldsymbol{\alpha}_K = \mathbf{0}$. In this multinomial framework, the country-c–industry-i portfolio cannot be affiliated with more than one cluster.

At this point, we should highlight some features of the multinomial prior. First, the vector, $\mathbf{z}_{ci}$, need not be composed of the same variables for each cluster k. This allows different characteristics to influence the composition of the clusters. For example, portfolios of countries that speak English as a primary language may be more likely to be included in cluster 1, while portfolios of countries with common currency may be more likely to be included in cluster 2. Second, note that the covariate vector does not have a time subscript, implying that the composition of the regions (and sectors) does not vary over time. Hamilton and Owyang (2012) argue that the prior hyperparameters can be viewed as population parameters signifying the relationships of the countries within a region.

2.3 Variance Decomposition

Given the model in (1)-(2), we can decompose the covariance between the country-c–industry-i portfolio's excess returns and country-d–industry-j portfolio's excess returns as follows

$$cov(R_{ci}, R_{dj}) = b_{ci}^G b_{dj}^G var\left(G\right) + \sum_{k=1}^{K} \gamma_{ci}^k \gamma_{dj}^k b_{ci}^k b_{dj}^k var\left(F_k\right) + cov\left(\epsilon_{cit}, \epsilon_{djt}\right). \tag{4}$$

Because we have assumed that the cross-portfolio residual correlation is zero (i.e., $cov\left(\epsilon_{cit}, \epsilon_{djt}\right) = 0$ for $i \neq j$ and $c \neq d$), the global factor and a potential common cluster factor are the only two possible sources of comovements between the (c, i) −portfolio's excess return with the (d, j) −portfolio's excess return. The component of the covariance attributable to each of these sources is determined by the variance of the factor and the product of the two portfolio's loadings. The component of the covariance attributable to the cluster factors is also determined by the product of the portfolio's membership indicators. Specifically, if there is a factor F_k for which $\gamma_{ci}^k \gamma_{dj}^k = 1$, the contribution of the common cluster factor to the covariance between the two portfolios is given by the factor variance weighted by their exposure to the common cluster factor, $b_{ci}^k b_{dj}^k var\left(F_k\right)$.

Having obtained an avenue for decomposing the covariance between two country-industry portfolios, we can now compute the components of the covariance between the value-weighted portfolios of countries c and d. Let w_{ci} and w_{dj} reflect the individual portfolio weights based on the average market capitalization and $W_{cd} = \sum_{i=1}^{I} \sum_{j=i+1}^{I} w_{ci} w_{dj}$ be a scalar that normalizes the weights to sum to one. The covariance between value-weighted portfolios of countries c and d can then be obtained from a weighted sum of the covariance between each of the individual country-industry covariances:

$$cov(R_c, R_d) = \frac{1}{W_{cd}} \sum_{i=1}^{I} \sum_{j=i+1}^{I} w_{ci} w_{dj} cov(R_{ci}, R_{dj}), \tag{5}$$

where, as argued above, $cov(R_{ci}, R_{dj})$ is determined by the global factor and the cluster factor, provided the country-industry portfolios belong to the same cluster. Thus, the covariance between country portfolios is decomposed into two components: the first is due to global integration and the second is due to cluster integration. The covariance between value-weighted industry portfolios can be obtained similarly, by instead integrating over countries for a single industry.

3 Implementation

In this section, we describe the data and methods used to obtain our results based on our endogenous cluster model in (1)-(2).

3.1 Data

We use monthly excess returns on country-industry portfolios for 23 countries and 25 industries (N = 575) from January 1980 to December 2016 (T = 444). All data are downloaded from Datastream using the Level 1 industry classification and total returns, which include reinvested dividends. Country-industry portfolio returns are constructed by calculating a value-weighted return for the portfolio for each period. We convert local currency returns into U.S. dollars with the DataStream exchange rate conversion facility and compute excess returns using the 3-month T-bill rate.

Tables 1 and 2 list, respectively, the countries and industries in our sample along with the earliest

Table 1: Countries list

	Country	Code	Earliest Start	Latest Start	#na
1	US	US	01/1980	02/1998	0
2	UK	UK	01/1980	12/2007	0
3	Germany	BD	01/1980	06/2006	0
4	France	FR	01/1980	08/2000	0
5	Italy	IT	01/1980	12/1995	1
6	Australia	AU	01/1980	12/2000	0
7	Austria	OE	01/1980	06/2008	4
8	Belgium	BG	01/1980	05/2005	2
9	Denmark	DK	01/1980	07/2013	3
10	Finland	FN	04/1988	05/2005	1
11	Norway	NW	02/1980	11/2014	2
12	Sweden	SD	02/1982	07/2008	1
13	Netherlands	NL	01/1980	06/2004	0
14	New Zealand	NZ	02/1988	06/2010	0
15	Portugal	PT	02/1988	02/2008	2
16	Spain	ES	04/1987	07/2014	0
17	Ireland	IR	01/1980	01/2014	3
18	Switzerland	SW	01/1980	06/2011	1
19	Greece	GR	02/1988	09/2003	2
20	Canada	CN	01/1980	10/1993	0
21	Hong Kong	HK	01/1980	07/2014	1
22	Japan	JP	01/1980	07/2005	0
23	Singapore	SP	01/1980	05/2011	1

Note: This table lists the countries in our sample along with their code (third column), the earliest and latest start dates for portfolios in each country (fourth and fifth columns) and the number of industries for which a portfolio in each country is not available (last column).

Table 2: Industries list

	Industry	Code	Earliest Start	Latest Start	#na
1	Oil and Gas	OILGS	01/1980	12/2006	0
2	Basic Materials	BMATR	01/1980	02/1990	0
3	Chemicals	CHMCL	01/1980	07/2014	1
4	Basic Resources	BRESR	01/1980	08/2007	1
5	Industrials	INDUS	01/1980	04/1994	0
6	Construction and Materials	CNSTM	01/1980	04/2001	0
7	Industrial Goods and Services	INDGS	01/1980	02/1990	0
8	Diversified Real Estate Inv. Trusts	RITDV	01/1980	07/2014	11
9	Consumer Goods	CNSMG	01/1980	11/1996	0
10	Auto and Parts	AUTMB	01/1980	06/2011	6
11	Food and Beverages	FDBEV	01/1980	10/2005	0
12	Personal and Household Goods	PERHH	01/1980	05/2003	0
13	Health Care	HLTHC	01/1980	11/2007	0
14	Consumer Services	CNSMS	01/1980	12/1997	0
15	Retail	RTAIL	01/1980	11/2014	1
16	Media	MEDIA	01/1980	07/2006	1
17	Travel and Leisure	TRLES	01/1980	01/2004	0
18	Telecommunications	TELCM	01/1980	03/2007	0
19	Utilities	UTILS	01/1980	02/2001	1
20	Financials	FINAN	01/1980	06/1998	0
21	Banks	BANKS	01/1980	06/2010	0
22	Insurance	INSUR	01/1980	07/2000	1
23	Real Estate	RLEST	01/1980	08/2013	0
24	Financial Services	FINSV	01/1980	06/2008	0
25	Technology	TECNO	01/1980	07/2007	1

Note: This table lists the industries in our sample along with their code (third column), the earliest and latest start dates for portfolios in each industry (fourth and fifth columns) and the number of countries for which a portfolio in each industry is not available (last column).

and the latest start dates for portfolios in each industry or country, and the number of portfolios in each industry or country that are not available. A total of 24 country-industry portfolios—mostly 'Auto and Parts' and 'Diversified Real Estate Investment Trusts' in small countries—are not available. In addition, some country-industry portfolios have a very short time series, with some starting only in 2014. While the estimation algorithm can handle missing observations, we only use portfolios for which we observe at least 50% of the observations within the sample under consideration. Therefore, for the full-sample version, we end up with $N = 482$. We treat the unbalanced panel as containing missing observations which is easily dealt with in the Kalman filter algorithm for extracting the common factors. Because we use value-weighted measures throughout the paper, the use of an unbalanced panel does not affect our results.

3.2 Estimation

The model outlined in the preceding section can be estimated using Bayesian techniques (see Gelfand and Smith 1990, Casella and George 1992, Carter and Kohn 1994). Bayesian methods allow us to estimate the cluster membership parameters directly using reversible jump Metropolis-Hastings steps in the Gibbs sampler.[2]

The sampler is an MCMC algorithm which draws from the conditional distributions of each parameter block conditional on the previous draws from the remaining parameters. The sequence of draws from the conditional distributions converges to the joint posterior. Let $\mathbf{Y}$ represent the data, $\boldsymbol{\Theta}$ represent the full set of model parameters, and $\mathbf{F}$ represent the full set of factors. Conditional on the number of clusters K, the model parameters and factors can be drawn in four blocks: (1) the membership indicators, γ, the factor loadings, $\mathbf{b}$, and the innovation variances, $\sigma^{\mathbf{2}}$; (2) the factors, $\mathbf{F}$; (3) the set of factor autoregressive parameters, ϕ; and (4) the multinomial prior hyperparameters, α. In the last block, we sample two additional sets of values: a vector of augmented data ξ and the logistic variance, χ.

[2] Ando and Bai (2017) present a frequentist alternative to our methodology. We have the advantage of being able to incorporate prior information into the estimation and to parameterize the prior to account for the observable characteristics of the data. This helps provide information to determine on what basis the clusters originate.

The prior for the parameters of each series slope coefficients is normal, $b_{ci} = [b^G_{ci}, b^k_{ci}]' \sim N(\beta_0; \mathbf{B}_0)$, and the innovation variances are inverse gamma, $\sigma^{-2}_{ci} \sim \Gamma(\nu_0, \Upsilon_0)$. The factor AR parameters have normal priors, $\phi \sim N(\mathbf{v_0}, \mathbf{V_0^{-1}})$. The multinomial prior hyperparameters also have normal priors, $\alpha \sim N(\mathbf{a_0}, \mathbf{A_0^{-1}})$.

While the factors in hierarchical models such as Kose et al. (2003) can be drawn from faster procedures outlined in Otrok and Whiteman (1998), the model posited here is not necessarily hierarchical. Thus, we draw the factor from smoothed Kalman filter posterior distributions. Fortunately, conditional on the model parameters and the cluster memberships, the state space is linear and the Kalman filter posteriors are straightforward to obtain.[3]

The main issue in the estimation of the model is that the cluster memberships can change across Gibbs iterations. To solve this problem, we draw the memberships and the loadings jointly. We first propose moving a portfolio to a different cluster. We can then compute the ratio of the posterior likelihoods between the new and old cluster memberships. We accept the new composition with a probability equal to this ratio of posterior likelihoods and, if accepted, draw a new set of factor loadings.[4] For our proposal, we choose the alternate cluster with equal probability assigned to all alternatives.

Estimation of the hyperparameters of the multinomial logistic prior is similar to an empirical Bayes strategy. The Metropolis algorithm described above allocates portfolios (with higher probability) to the clusters that have higher likelihood. In the case of a "tie", one can think of the prior as allocating the portfolio to the cluster that is most similar in the z_{ci} sense. The weights placed on the various elements of z_{ci} are determined to maximize the overall likelihood. The standard approach to estimating the multinomial logistic prior follows the data augmentation technique of Tanner and Wong (1987) and introduces a set of latent vectors, ξ. Each vector, $\xi^k = \left(\xi^k_{11}, \ldots, \xi^k_{CI}\right)'$, consists of latent variables for all countries $c = 1, \ldots, C$ and industries $i = 1, \ldots, I$ with values

[3] Because the sign of the factor and its loading are not separately identified, we impose restrictions on the signs of the factors as outlined in Francis et al. (2017b).

[4] Troughton and Godsill (1997) show that the ratio of the posterior likelihoods does not depend on the draw of the slope coefficients (in our case, the factor loadings). Thus, we only draw the loadings if the proposal is accepted.

such that

$$\xi_{ci}^{k} \geq 0, \text{ if } \gamma_{ci}^{k} = 1$$
$$\xi_{ci}^{k} < 0, \text{ otherwise.}$$

Each ξ_{ci}^{k} element can be drawn from a truncated logistic distribution with associated variance, χ_{ci}^{k}. We apply the methodology described in Francis et al. (2017b) to draw these prior hyperparameters.[5]

After initializing the sampler with 30,000 draws to allow for convergence, we execute 20,000 iterations to form the joint posterior distribution. Notice that while the factors are assumed to be uncorrelated, the small-sample results may produce posterior estimates of the factors with some non-zero correlation. When constructing covariances and correlations, we orthogonalize the cluster factors from the global factor. We then estimate the factor loadings based on the original global factor and the orthogonalized cluster factor with the posterior mode cluster membership for each excess return portfolio. The variance terms in equation (4) are computed based on the sample characteristics of the posterior mean factor estimates.

3.3 Time-Variation

To identify time-variations in the factor structure, following Bekaert et al. (2009), we re-estimate the model every 2.5 years using a window of 5 years of data, essentially assuming that within the 5-year period the cluster indicators, factor loadings and volatilities are constant.[6] We then compute the empirical covariance matrix of our portfolios for each window, $cov_{\tau}(R_{ci}, R_{dj})$, using the appropriate subsample of data. The covariance between two portfolios can change over time through five channels: (i) changes in their exposures to the global factor $b_{ci\tau}^{G} b_{dj\tau}^{G}$, (ii) changes in cluster memberships $\gamma_{ci\tau}^{k} \gamma_{dj\tau}^{k}$, (iii) changes in exposure to the cluster factors $b_{ci\tau}^{k} b_{dj\tau}^{k}$, (iv) changes in the volatility of the global factor $var_{\tau}(G)$, and (v) changes in the volatility of the common cluster

[5] We refer the reader to Francis et al. (2017b) for a thorough discussion of the econometric details of the estimation.

[6] We exclude portfolio excess return series for which we are missing more than 50% of the observations within each 5-year window. Therefore, the full set of observable series used to estimate the model at each point in time changes as we gain new information on series that appear later in the sample.

factor $var_\tau(F_k)$. If an increase in the covariance between two portfolios is due to an increase in their exposure to the global factor, then it indicates an increase in global integration. An increase in the covariance between two portfolios due to changes in cluster membership and in the exposure to the cluster factors indicates a increase in cluster integration.

The time-varying covariances between the value-weighted portfolios of countries c and d can be computed similarly as their full sample analogue using the individual portfolio weights given by the average market capitalization within the subsample. The time-varying covariance between value-weighted portfolios of industries i and j can be computed accordingly.

4 Results

In this section, we present estimation results from our cluster model, based both on the full sample of data and on rolling windows, with a focus on the cluster composition and its implication for country-industry portfolios comovements.

4.1 Full-Sample

We start by estimating the model on the full sample of data assuming $K = 15$. To construct the hierarchical prior for cluster membership, we include both country and industry dummies. Because we estimate the prior hyperparameters, including country and industry dummies will allow us to determine which of these factors—if either—are relevant for cluster formation. Figure 1 summarizes the cluster composition for the full sample as a heat map. The heat map shows the mode cluster inclusion probability—the highest probability cluster—which is computed as the maximum across k of the sum of γ_{ci}^k over the saved Gibbs iterations. The x-axis lists the 23 countries in our sample and the y-axis shows the 25 industries. The shaded boxes indicate to which cluster each country-industry portfolio is assigned.

The propensity for vertical shading in the figure suggests the presence of clusters more associated with countries rather than industries. We see come clear country-level cluster definitions: Cluster

Figure 1: Full-Sample Cluster Heat Map

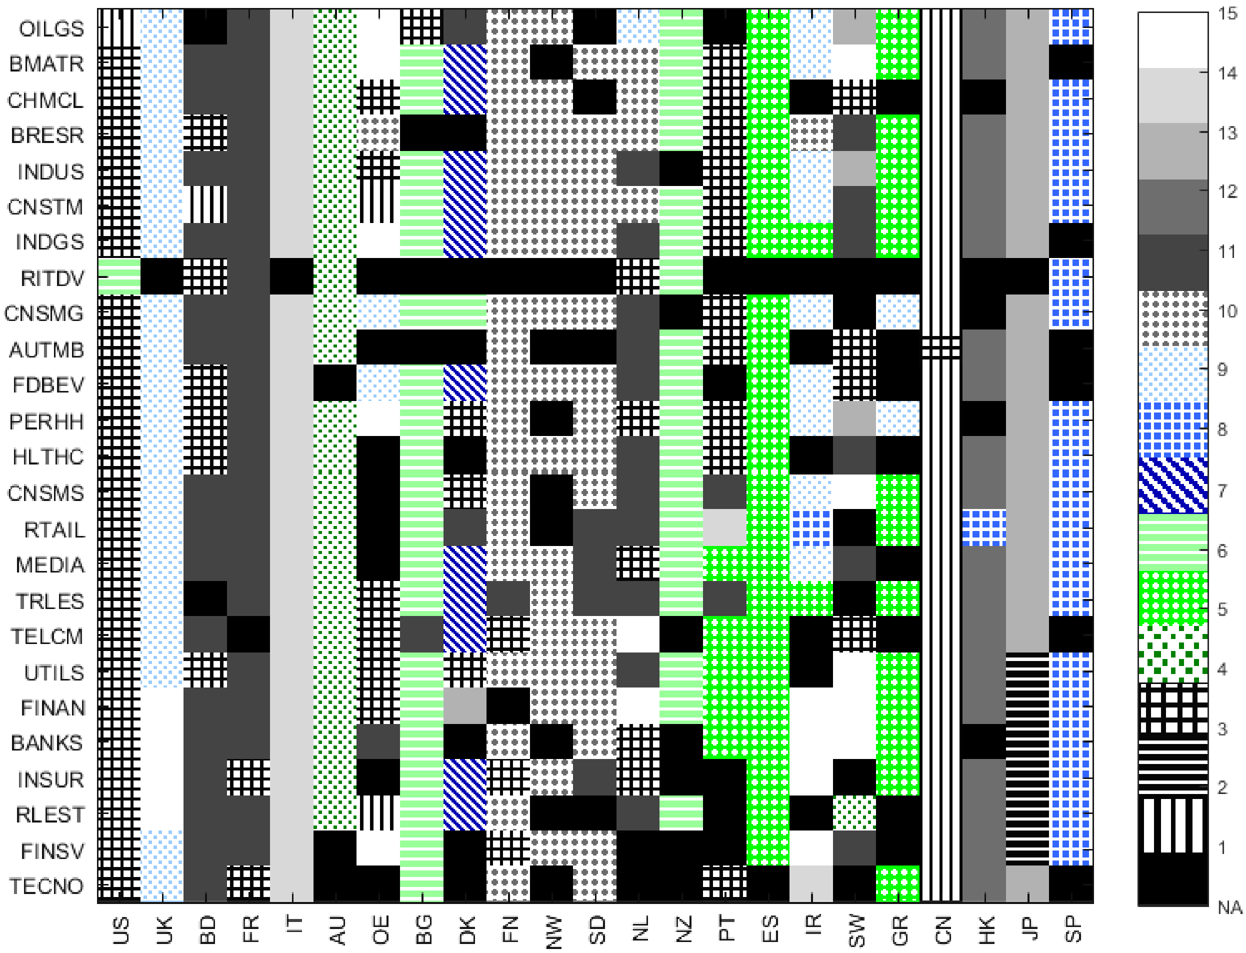

The shaded boxes indicate to which cluster each country-industry portfolio is assigned based on the mode cluster inclusion probability, computed as the maximum across k of the sum of γ_{ci}^{k} over the saved Gibbs iterations. The x-axis lists the 23 countries in our sample and the y-axis shows the 25 industries. Country-industry portfolios that were not included in the estimation because of missing data are denoted by NA.

3-US, Cluster 9-UK/Ireland, Cluster 11-France/Germany, Cluster 14-Italy, Cluster 4-Australia, Cluster 6-Belgium/New Zealand, Cluster 7-Denmark, Cluster 10-Norway/Finland/Sweden, Cluster 5-Spain/Greece/Portugal, Cluster 1-Canada, Cluster 12-Hong Kong Cluster 13-Japan and Cluster 8-Singapore. While most clusters are strongly identified with one country, some include larger geographical areas. In particular, Clusters 10 and 11 appear to be European clusters with varied membership across northern and western European countries. The heat map is characterized by a number of distinct vertical patterns but few horizontal—cross-industry—patterns, suggesting that within-country comovement dominates within-industry comovement. Thus, diversification can be achieved by investing across geographical areas rather than across sectors.

To assess the role of prior information about country and industry in the determination of cluster membership, we estimate an alternative version of the model with a uniform prior probability of membership across clusters. Thus, we impose that each portfolio is equally likely to belong to any of the 15 possible clusters. Figure 2 reports the mode posterior cluster membership indicators for all country-industry portfolios in the full sample with this uniform prior. Similar to the results with a hierarchical prior, we find comparable cases of country-level cluster definitions. For example, in this version of the model we also find clusters strongly associated with US (Cluster 3), UK (Cluster 14), Germany (Cluster 9), France (Cluster 15), Italy (Cluster 7), Australia/New Zealand (Cluster 6), Portugal/Spain (Cluster 1), Hong Kong (Cluster12), Japan (Cluster 8) and Singapore (Cluster 10). The remaining clusters show minor differences regarding broader clusters of multiple countries. In accordance with the previous results, we find less evidence of clustering across industries. Since the results with a uniform prior are generally consistent with those using the hierarchical prior, we prefer to utilize the slightly more informative specification to provide added context to explaining how and why country-industry portfolios may comove.

These results are based on full-sample excess returns, likely spanning a very diverse timeline of economic and financial conditions. It will likely be more realistic to consider time-variation in the level of comovement, allowing for cluster membership and volatilities to change over time. We address this issue below.

Figure 2: Full-Sample Cluster Heat Map with Uniform Prior

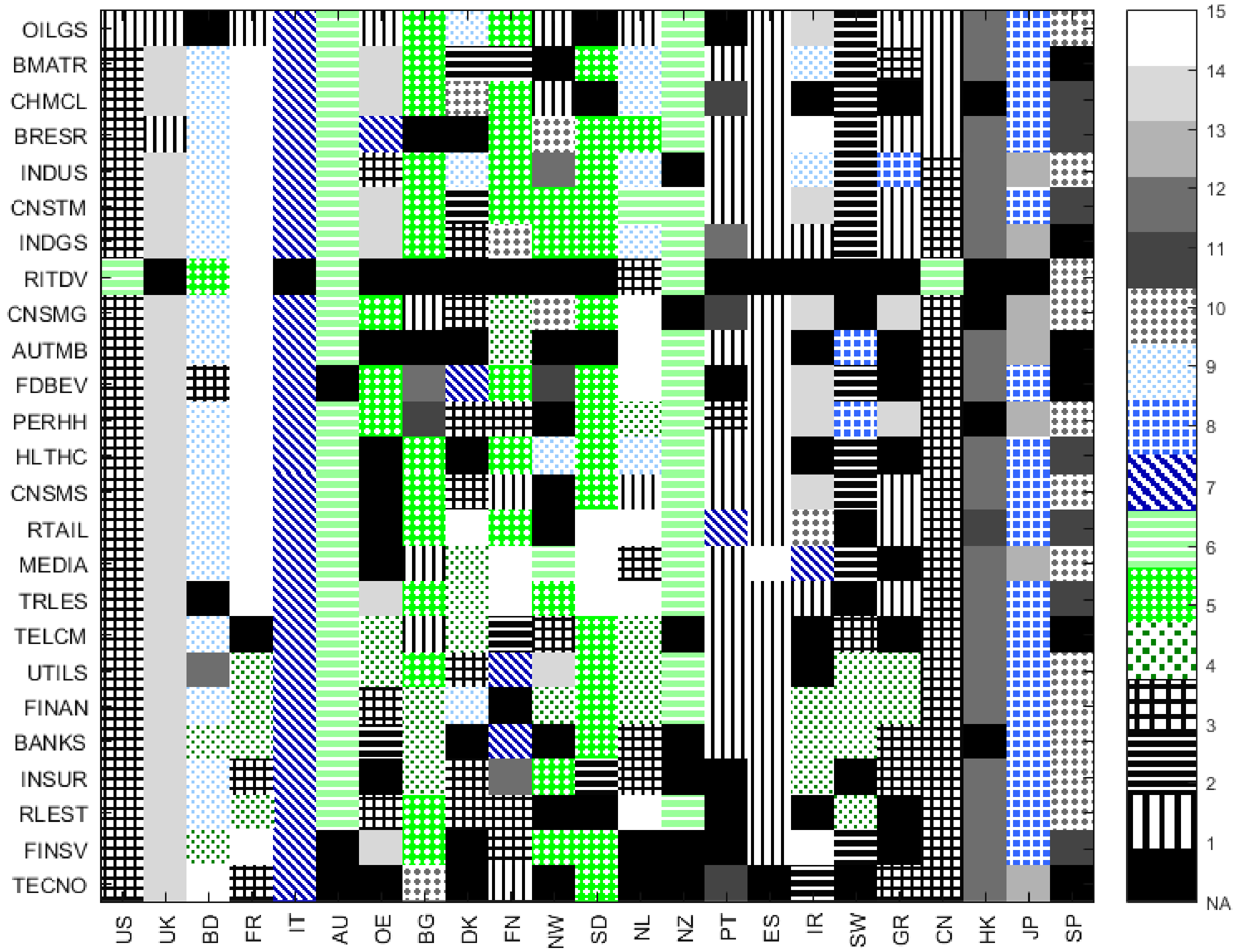

The shaded boxes indicate to which cluster each country-industry portfolio is assigned based on the mode cluster inclusion probability, computed as the maximum across k of the sum of γ_{ci}^{k} over the saved Gibbs iterations. In this specification, we employ a uniform prior making it equally likely that a portfolio belongs to any of the 15 possible clusters. The x-axis lists the 23 countries in our sample and the y-axis shows the 25 industries. Country-industry portfolios that were not included in the estimation because of missing data are denoted by NA.

Figure 3: Global Factor

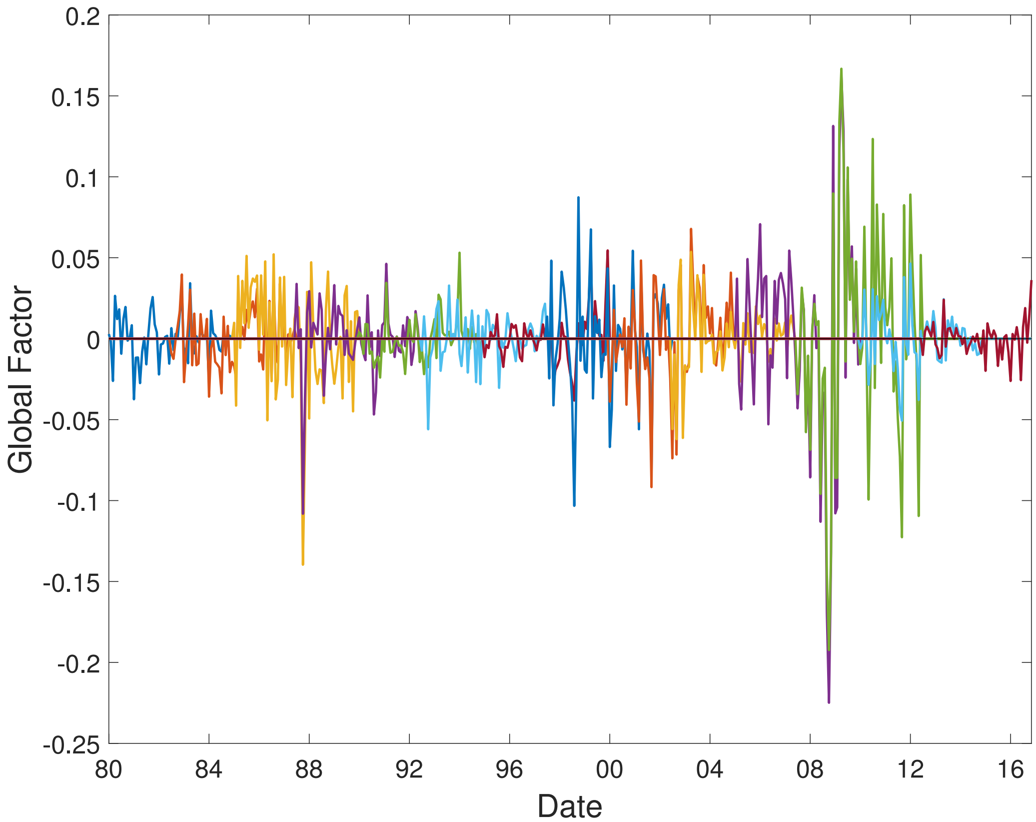

Posterior mean estimates of the global factor estimated from each of the 5-year rolling windows.

4.2 Rolling Windows

Next, we estimate the model in (1)-(2) on 14 rolling-window subsamples, each using 5-years of data and rolling the window forward 2.5 years at each iteration. The first subsample uses data from January 1980 through December 1984.[7] Figure 3 shows the time series of posterior mean estimates of the global factor throughout the subsamples. For the periods in which the windows overlap, we show the global factor estimated for each separate subsample. The volatility of the

[7] Due to data availability, when we roll forward 2.5 years to estimate the last window, this produces a window with only 4.5 years of data remaining. This window spans the months from July 2012 through December 2016.

global component increases dramatically in the latter portion of the sample, exhibiting large swings right around the global financial crisis in 2008 and 2009 that affected most financial markets.

In order to assess the relative importance of the global and cluster factors, we apply the variance decomposition in (4) to all the country-industry portfolio returns and compute a value-weighted average of the three components (global, cluster and idiosyncratic) using subsample average market capitalizations. A clear pattern in Figure 4 emerges: the cluster factors explain the majority of the variance up until the mid-2000's when the global factor starts to provide more explanatory power. This indicates that, as the effects of globalization permeated across the countries and industries in our sample, the global factor becomes more important in explaining the variance of excess returns, reaching a peak after 2009. However, at the end of the sample, the global factor loses importance and cluster membership appears to be more influential in explaining comovement among portfolios.

We also compute a value-weighted variance decompositions for a subset of countries (US, UK, Germany, Italy, Australia, and Japan) and a subset of industries (Oil&Gas, Financials, Utilities, Consumer Goods, Basic Materials, and Technology) by using (4). Figures 5 and 6 illustrate variation in the variance decomposition across countries and industries, respectively. The world factor explains the largest share of the variance for all six countries and all six industries around 2008-2009 and declines in importance near the end of the sample. More prominent differences are seen with respect to the importance of the cluster factors. Early in the sample, the relative importance of the global, cluster, and idiosyncratic factors fluctuates for the US, UK, and Germany. However, for Italy and Australia, the cluster factor is consistently more important until the mid- to late-2000's but then reemerges again at the end of the sample. The idiosyncratic component explains the largest share of the variance for Japan throughout most of the sample, comparable in magnitude to the cluster factor, and both explaining considerably more than the global factor, with the exception of 2009. The influence of cluster membership for industries tells a slightly different story in Figure 6. While for all industries considered here the cluster factor does seem to matter somewhat for explaining the variance of observed excess returns, the relative contribution is considerably smaller. At industry level, the idiosyncratic component dominates both the global and cluster factors for

Figure 4: Rolling Window Variance Decompositions

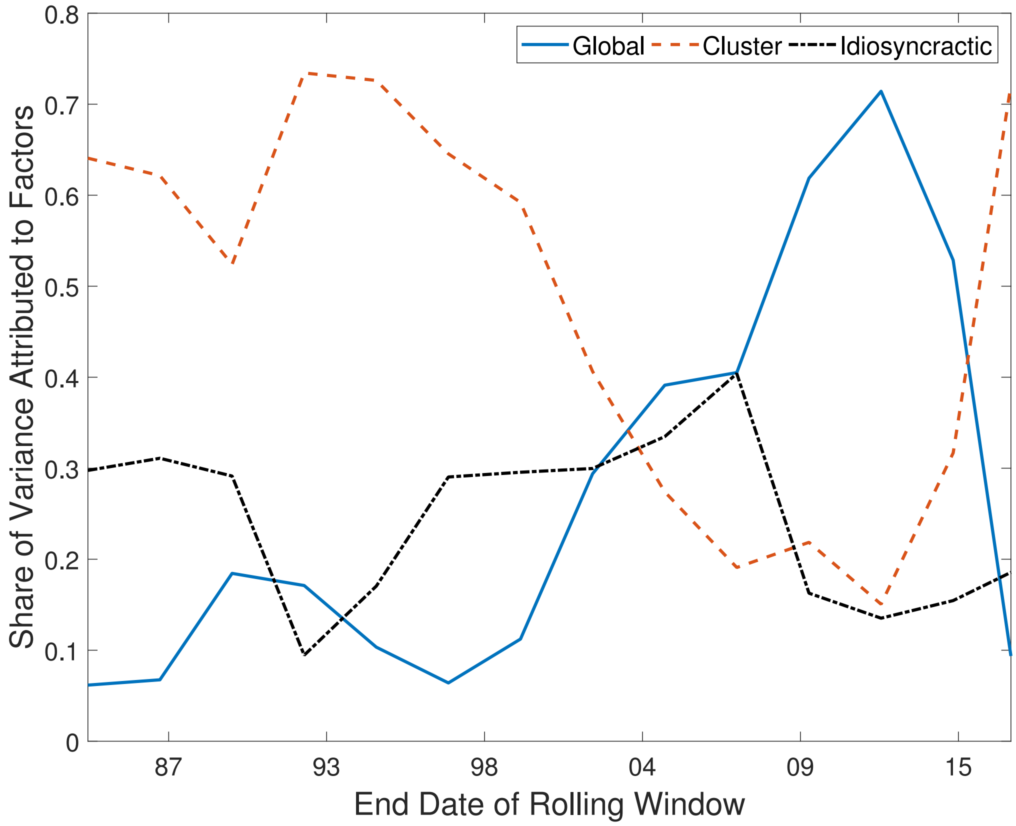

The value-weighted excess return variance in each subsample is computed as

$$Var_\tau(R) = \frac{1}{W_\tau}\sum_{c=1}^{C}\sum_{i=1}^{I} w_{ci\tau} Var_\tau(R_{ic})$$

where $Var_\tau(R_{ic})$ is decomposed as in (4), and W_τ is the total market capitalization in subsample τ, i.e. $W_\tau = \sum_{c=1}^{C}\sum_{i=1}^{I} w_{ci\tau}$. Value reported are in percentage over the total variance.

Figure 5: Rolling Window Variance Decompositions - By Country

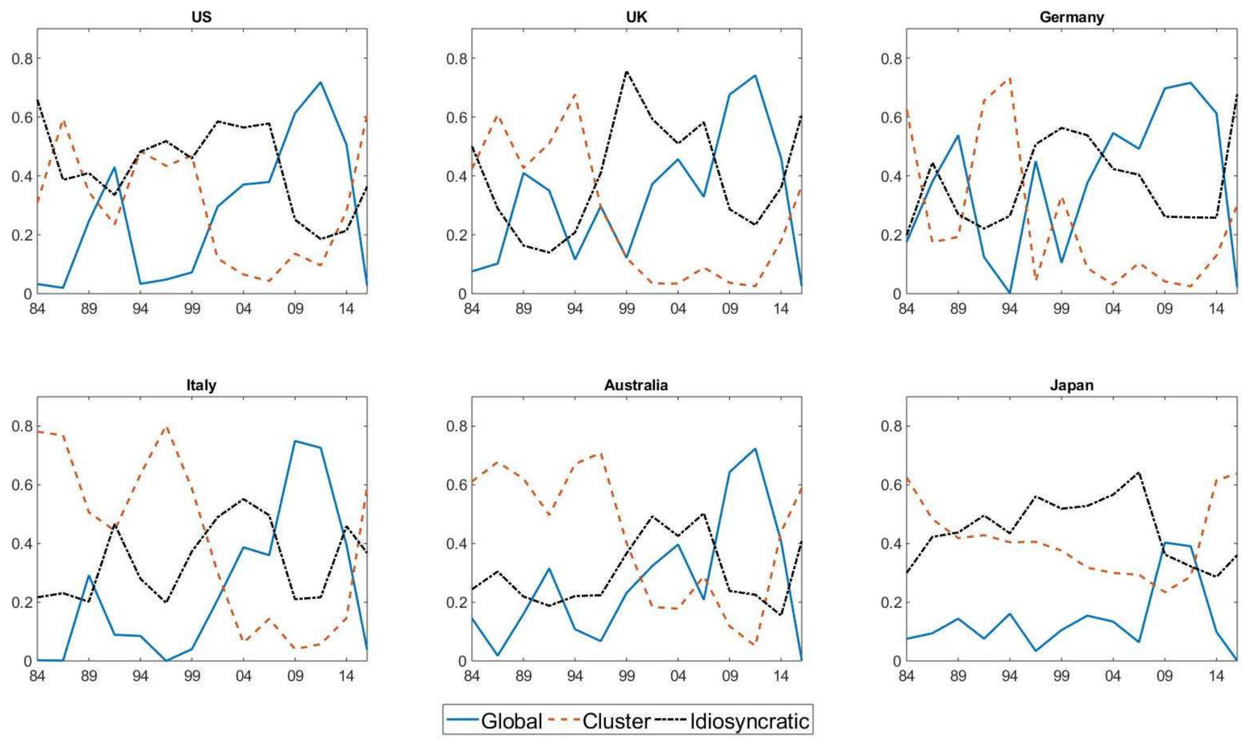

The value-weighted excess return variance in each subsample is computed as

$$Var_{\tau}(R_c) = \frac{1}{W_{c\tau}} \sum_{i=1}^{I} w_{ci\tau} Var_{\tau}(R_{ic})$$

where $Var_{\tau}(R_{ic})$ is decomposed as in (4), and $W_{c\tau}$ is the total market capitalization in country c in subsample τ, i.e. $W_{c\tau} = \sum_{i=1}^{I} w_{ci\tau}$. Value reported are in percentage over the total variance.

Figure 6: Rolling Window Variance Decompositions - By Industry

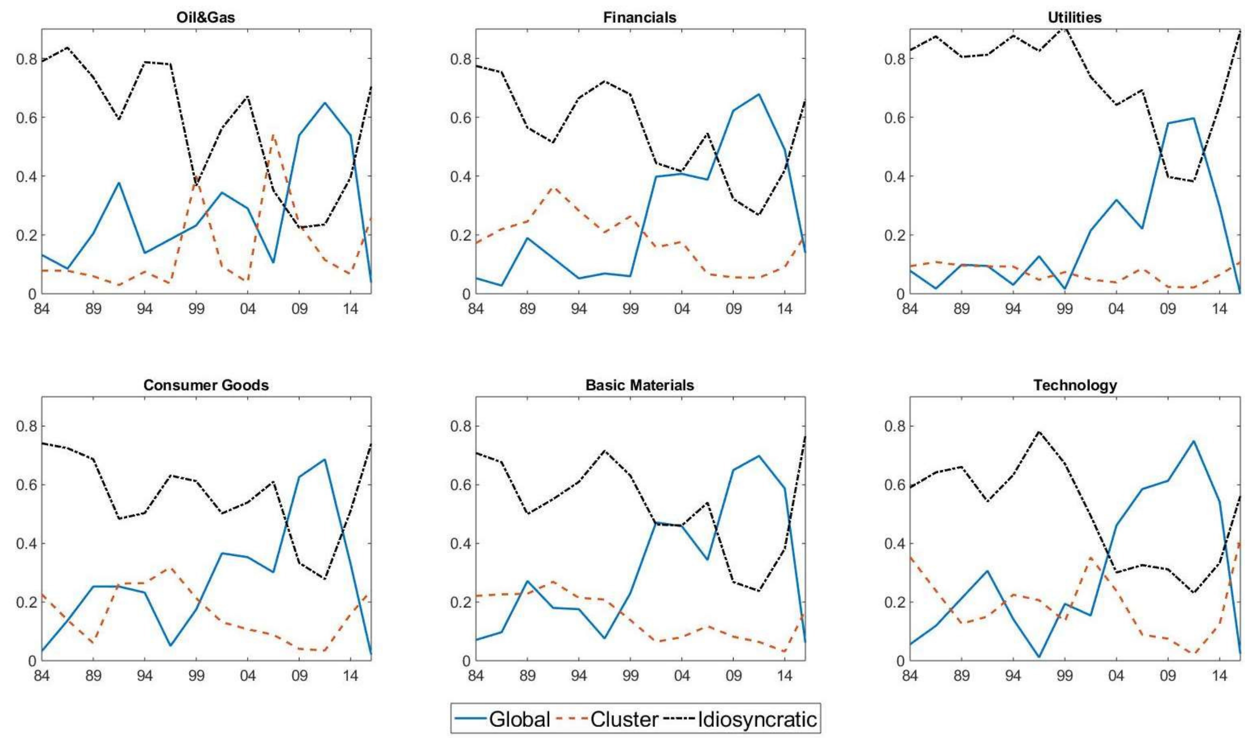

The value-weighted excess return variance in each subsample is computed as

$$Var_\tau(R_i) = \frac{1}{W_{i\tau}} \sum_{c=1}^{C} w_{ci\tau} Var_\tau(R_{ic})$$

where $Var_\tau(R_{ic})$ is decomposed as in (4), and $W_{i\tau}$ is the total market capitalization in industry i in subsample τ, i.e. $W_{i\tau} = \sum_{c=1}^{C} w_{ci\tau}$. Value reported are in percentage over the total variance.

most of the pre-Great Recession sample. This suggests that cluster factors are more related to country rather than to industry variations.

For each window, we construct the correlation between value-weighted portfolios of countries and industries. Rather than showing the full table of correlations for all 14 windows, we instead show the time-variation in the correlation decomposition between selected country and industry pairings.[8] Figure 7 illustrates how the total correlation has changed over time between the following country pairs: US-Canada, Hong Kong-Singapore, Germany-Spain, Italy-Germany, Australia-New Zealand, and Belgium-New Zealand. The total correlation is decomposed into a global component and a cluster component. Data from Spain first appear in the window ending in 1989 and data from New Zealand in the window ending in 1992. Prior to these dates, any correlations with these countries are considered to be missing. An initial observation highlights that the total correlation between all country pairs increased substantially in the late 1980's and again leading up to, and for some time after, the global financial crisis in 2008-2009. During each of these episodes, the correlation due to global integration was considerably larger in each pair.

Most of the correlations between countries are mainly driven by the global component, with the cluster component only playing a marginal and often temporary role. This is the case for the correlations between the US and Canada, Hong Kong and Singapore, Germany and Spain, and Italy and Spain. In addition, the correlation between Belgium and New Zealand is entirely explained by the global component, as the share of the total correlation attributed to the cluster component is zero (or virtually zero) in all subsamples. This happens because all the country-industry portfolio returns in these two countries are driven by separate cluster factors in all subsamples, so the only source of comovement is the global component.[9] On the contrary, the correlation between Australia and New Zealand displays a different pattern, as the cluster component explains a large share of the

[8] While we opt to show only six potential country pairs and four potential industry pairs; all correlations are available upon request.

[9] Taking a closer look back at the full-sample results in Figure 1 reveals one drawback to assuming constant cluster membership and volatility over time, as using all data from 1980-2016 produces a common cluster between Belgium and New Zealand. Therefore, the full-sample analysis may provide suggestive evidence only of common clusters that are widely persistent and misses the more realistic comovement that results when allowing for time-variation in both cluster memberships and volatilities.

Figure 7: Correlation between selected country portfolios within rolling-window subsamples

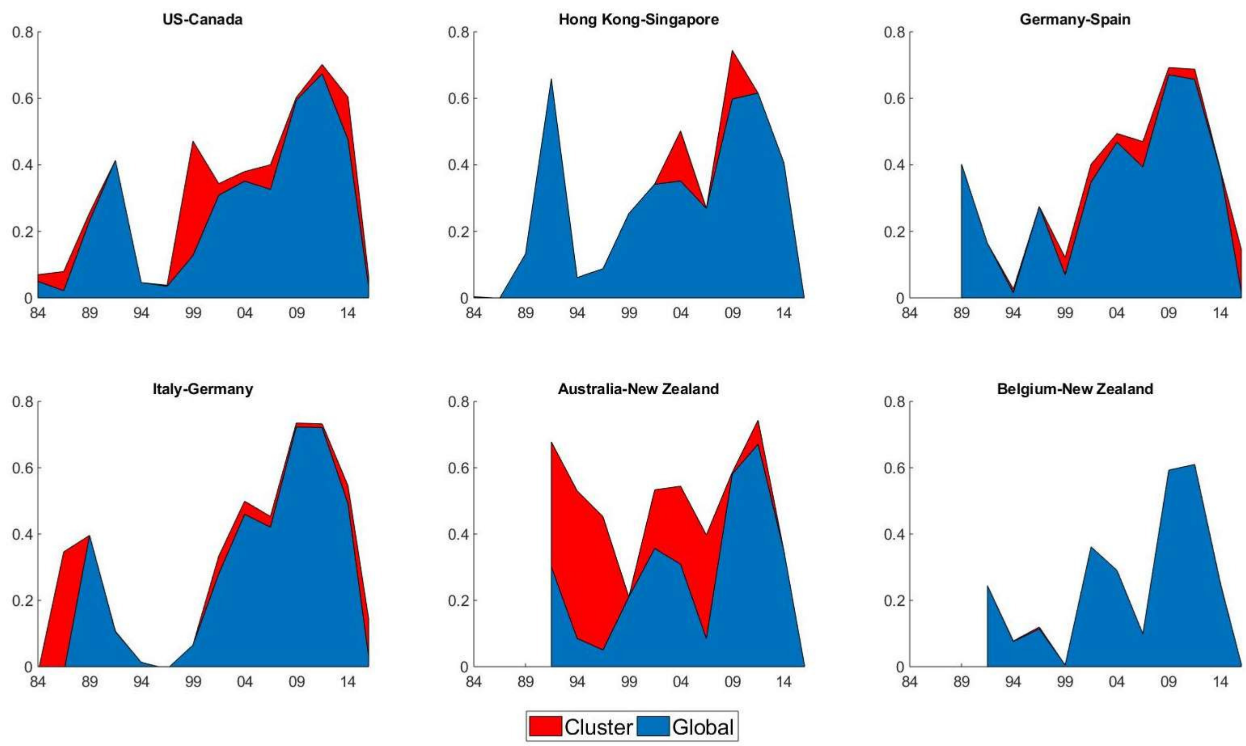

The correlation between country c and d is computed as

$$corr_\tau(R_c, R_d) = \frac{1}{W_{cd\tau}} \sum_{i=1}^{I} \sum_{j=i+1}^{I} w_{ci\tau} w_{dj\tau} \frac{cov_\tau(R_{ci}, R_{dj})}{SD_\tau(R_{ci}) \times SD_\tau(R_{dj})}$$

where $cov_\tau(R_{ci}, R_{dj})$ is decomposed as in (4), $W_{cd\tau} = \sum_{i=1}^{I} \sum_{j=i+1}^{I} w_{ci\tau} w_{dj\tau}$, and $SD_\tau(R_{ci})$ and $SD_\tau(R_{dj})$ are the sample standard deviations of the (c, i) and (d, j) portfolio's excess return in subsample τ.

overall correlation for most of the subsamples. This indicates that some country-industry portfolio returns in these two countries are driven by the same cluster factors, implying that the comovement between these two countries is driven by both the global and the cluster component.

Figure 8 shows time variation in the total correlation between the following industry pairs: Oil and Gas-Basic Materials, Banks-Financials, Telecom-Technology, and Auto and Parts-Consumer Goods, Utilities-Consumer Goods, and Travel and Leisure-Retail. Similarly to the country correlations, the total correlation between all industry pairs under consideration increases steadily and reaches a peak around the global financial crisis. The global component is the main driver of all the correlations, especially around the financial crisis. The cluster component is also important for all the industry pairs (except Utilities and Consumer Goods) in most of the subsamples. This indicates that some country-industry portfolios returns in these industries are driven by the same cluster factors, implying that the comovement between the two industries is driven by both the global and the cluster component.

Comparing results in Figures 7–8, it is clear that the cluster component has a more prominent role in explaining the comovement across industries rather than across countries. This means that diversifying a portfolio across two industries leads to a larger exposure to cluster risk than diversifying across two countries. This result is stable over time and further corroborates our evidence that country-industry portfolio returns tend to cluster within geographical areas. As a consequence, larger benefits in diversification can be achieved by investing across geographical areas rather than across sectors.

While it appears in both the country and industry correlation as if the relevance of the global component and overall level of correlation diminishes in the very end of the sample, further examination of the results reveals an interesting dynamic developing when extending the data through 2016. To provide some insight into cluster composition, we identify the clusters via the analogous, time-varying mode cluster membership from the full-sample analysis. Figure 9 depicts a heat map showing the posterior mode inclusion probabilities from the final subsample which ends in 2016. A number of country-specific clusters seems to emerge: Cluster 2-US, Cluster 8-Australia, Cluster

Figure 8: Correlation between select industry portfolios within rolling-window subsamples

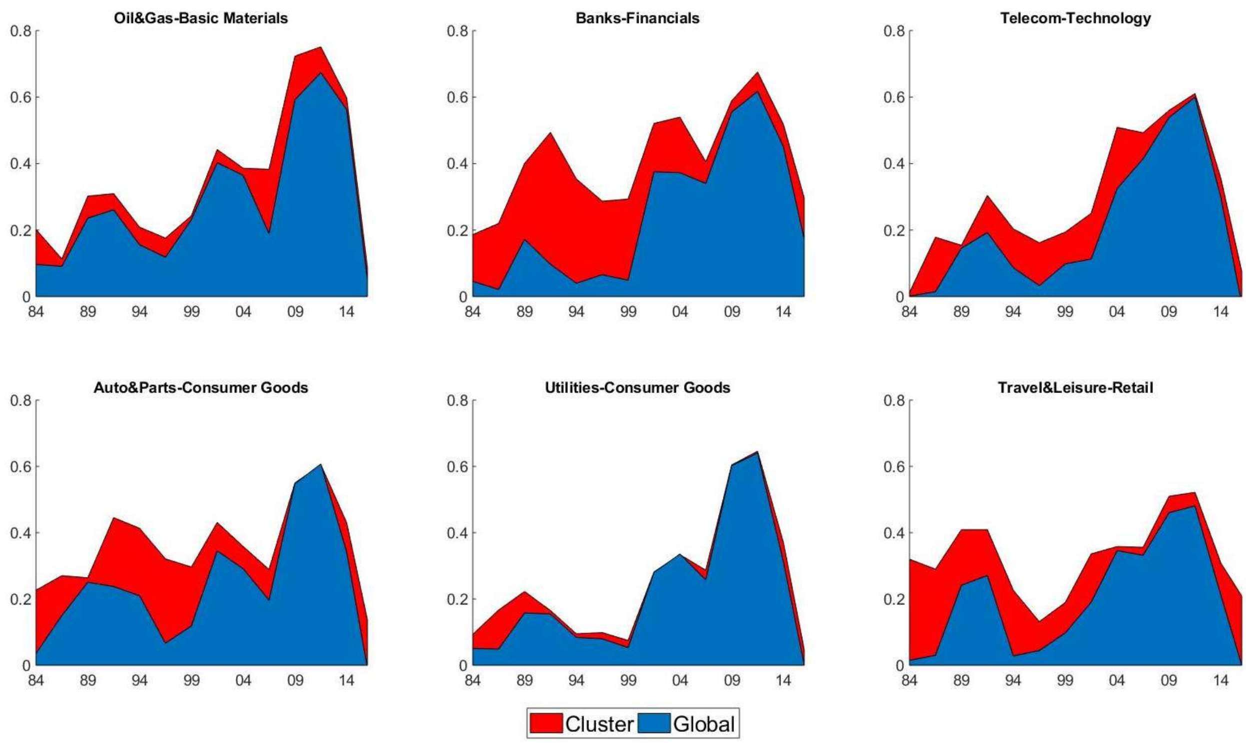

The correlation between industry i and j is computed as

$$corr_\tau(R_i, R_j) = \frac{1}{W_{ij\tau}} \sum_{c=1}^{C} \sum_{d=c+1}^{C} w_{ci\tau} w_{dj\tau} \frac{cov_\tau(R_{ci}, R_{dj})}{SD_\tau(R_{ci}) \times SD_\tau(R_{dj})}$$

where $cov_\tau(R_{ci}, R_{dj})$ is decomposed as in (4), $W_{ij\tau} = \sum_{c=1}^{C} \sum_{d=c+1}^{C} w_{ci\tau} w_{dj\tau}$, and $SD_\tau(R_{ci})$ and $SD_\tau(R_{dj})$ are the sample standard deviations of the (c, i) and (d, j) portfolio's excess return in subsample τ.

15-New Zealand, Cluster 1-Greece, Cluster 10-Canada, Cluster 7-Hong Kong, Cluster 11-Japan and Cluster 9-Singapore. However the most evident result in Figure 9 is that two clusters, Clusters 6 and 12, include a large number of European countries, indicating stronger European integration than in the full sample. Finally, consistent with the full-sample results, most clusters include a more diverse collection of industries and thus few specific industries stand out in their own independent clusters, as for example Basic Resources and Basic Material that are often in Cluster 13.

5 Comparison with Benchmarks

Results in the previous section indicate that country-industry portfolios tend to cluster mainly within geographical areas. Therefore, to understand the advantages of using our endogenous cluster model, we compare our model with a world-country factor model defined as follows:

$$R_{cit} = \overline{R}_{ci} + b_{ci}^{G} G_t + \sum_{d=1}^{C} \mathbf{1}_{d=c} b_{ci}^{d} F_{dt} + \epsilon_{cit}. \tag{6}$$

where $\mathbf{1}_{d=c}$ is an indicator that takes value 1 if the portfolio is in country d, i.e., $d = c$, and zero otherwise. This model differs from our endogenous cluster model in two dimensions. First, in (6) country-industry portfolio excess returns in each country are driven by a common country-specific factor, while in our endogenous cluster model in (1) the cluster indicator coefficients γ determine which portfolios share a common factor. Second, the model in (6) has a fixed number of factors given by $C + 1$, while in our endogenous cluster model we have that the total number of factors is $K + 1$, where we have imposed that $K < C$. More generally, the model in (6) can be seen as a country-only Heston and Rouwenhorst (1994) model that allows for non-unitary factor loadings. If geography is the unique driver of the country-industry portfolio clustering, this benchmark should outperform our endogenous cluster model, at least in terms of goodness of fit.

In addition to the world-country factor model, we also consider an arbitrage pricing theory (APT) model with both global and regional factors. Rather than including C country factors as in the previous benchmark, perhaps broader aggregation at the regional level is more appropriate

Figure 9: Cluster Heat Map for the Final 2016 Sub-Sample

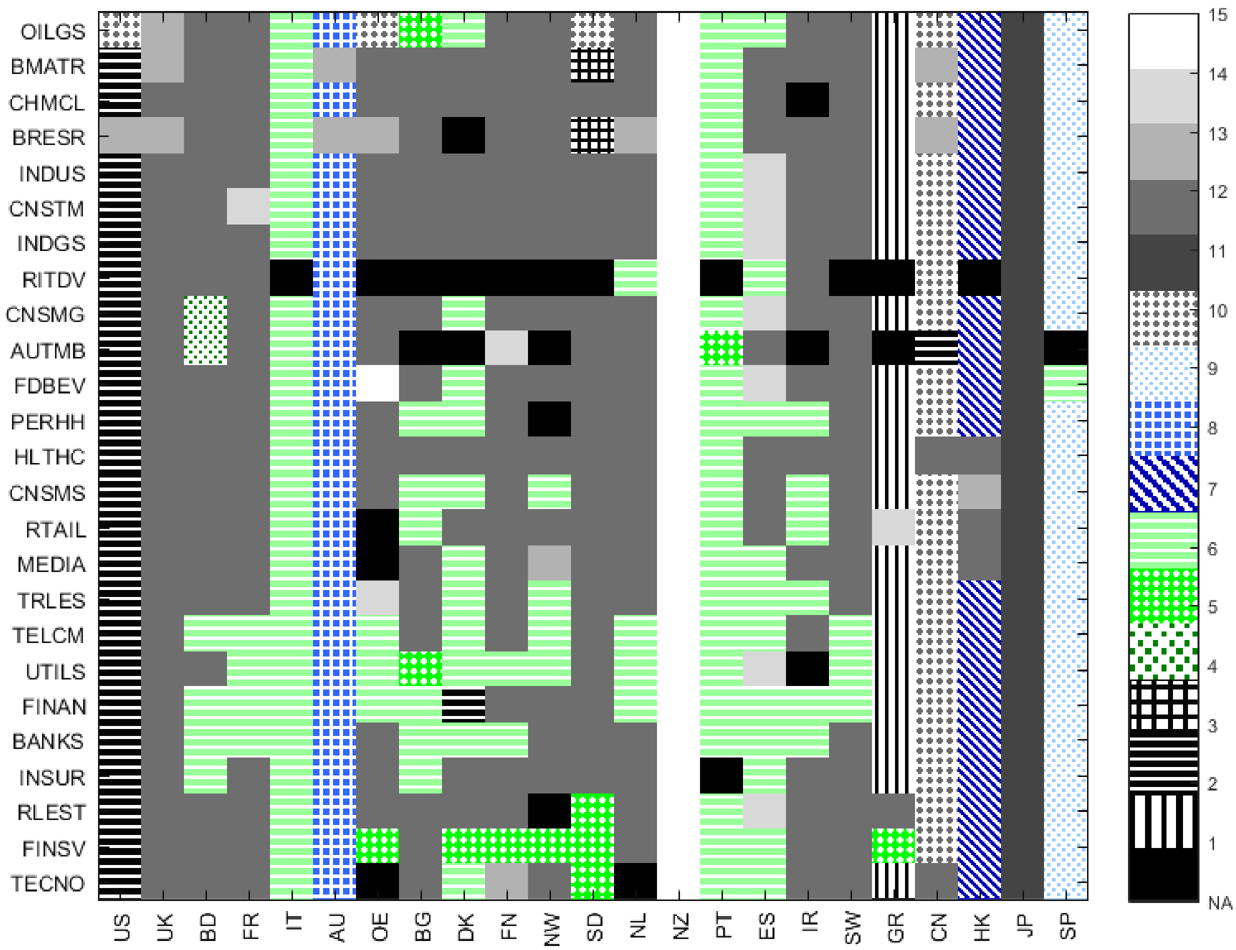

The shaded boxes indicate to which cluster each country-industry portfolio is assigned based on the mode cluster inclusion probability, computed as the maximum across k of the sum of γ_{ci}^{k} over the saved Gibbs iterations. The model is estimated using the final subsample in the rolling-window analysis, with data ending in 2016. The x-axis lists the 23 countries in our sample and the y-axis shows the 25 industries. Country-industry portfolios that were not included in the estimation because of missing data are denoted by NA.

since some countries tend to move together, especially near the end of our sample. We follow a similar approach to that of Bekaert et al. (2009) and include one global factor and separate the portfolios into three regions representing North America, Europe, and the Far East, estimating three factors for each region.[10] Thus, the world-region factor model we consider can be defined as follows:

$$R_{cit} = \overline{R}_{ci} + b_{ci}^{G} G_t + \sum_{r=1}^{3} \mathbf{1}_{c \in r} \sum_{j=1}^{3} b_{ci}^{rj} F_{rt}^{j} + \epsilon_{cit}, \tag{7}$$

where $\mathbf{1}_{c\in r}$ is an indicator that takes value 1 if country c is in region r, and zero otherwise. In this model, country-industry portfolio excess returns are driven by one global factor and three common, region-specific geographical factors; in contrast, our endogenous cluster model does not impose ex ante any specific geographical clustering. The world-region model in (7) has 10 factors, considerably fewer than the world-country model. This allows to group multiple countries together by assuming some type of international comovement within a given region. Bekaert et al. (2009) find that a world-region APT model outperforms a variety of alternatives both in- and out-of-sample.

To assess the relative in-sample performance of the three models, in Figure 10, we report their variance decompositions obtained estimating the three models over rolling windows. The values for the endogenous cluster model are the same as in Figure 4. The figure reveals the variance decompositions of our endogenous cluster and the world-country models follow the same general pattern but with two important differences. First, even if the world-country model has a larger number of factors (C+1=24) than our endogenous cluster model (K+1=16), it does not fit the observed country-industry portfolios better than the more parsimonious endogenous cluster model. On the contrary, the variance of the idiosyncratic component of the endogenous cluster model is lower than the one of the world-country model for all the rolling windows except the ones ending in December 1989 and in June 2007. The average variance of the idiosyncratic component of the endogenous cluster model is 0.245, while for the world-country model it is 0.265. The world-

[10]Bekaert et al. (2009) include three factors for each region plus three global factors while we only consider a single global factor. We have also estimated the model with three global factors and find that this results in overfitting the data with poor out-of-sample performance. Thus, we opt to present the more parsimonious results with a single global factor. The alternative results are available upon request.

Figure 10: Rolling Window Variance Decompositions: Comparison

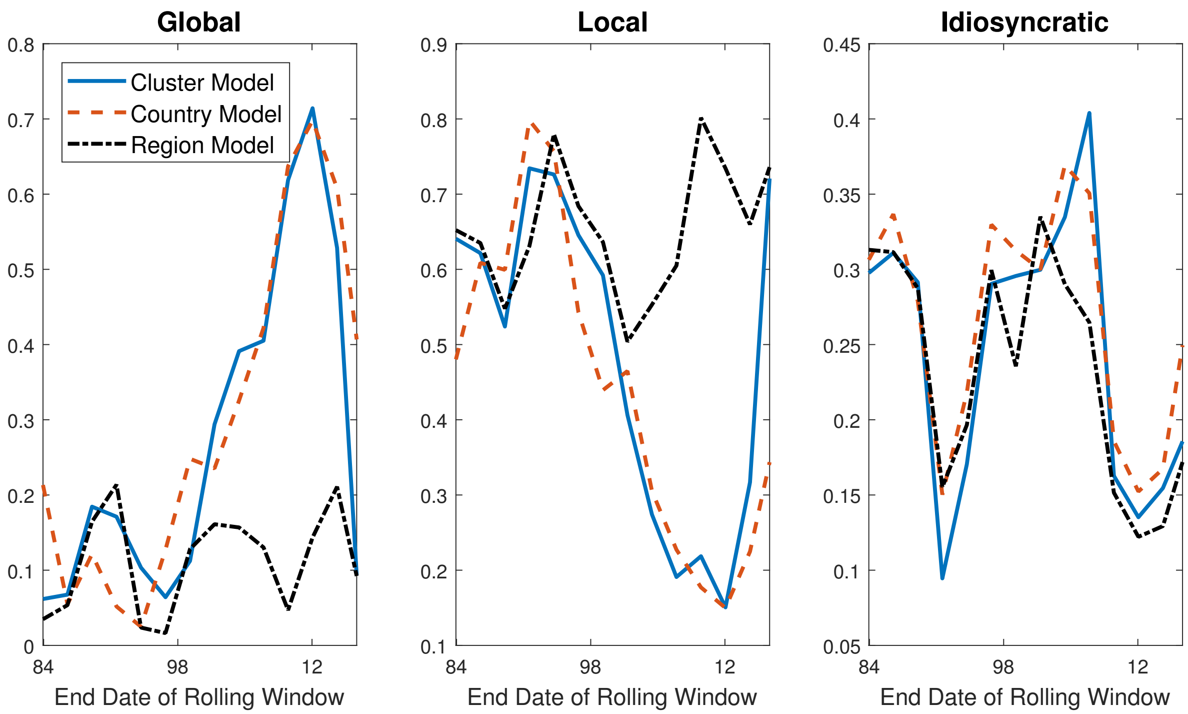

Rolling window variance decompositions for the endogenous cluster model (blue continuous line), the world-country model (red dashed line) and the world-region model (black dash-dotted line). The value-weighted excess return variance in each subsample is computed as

$$Var_\tau(R) = \frac{1}{W_{ci}} \sum_{c=1}^{C} \sum_{i=1}^{I} w_{ci} Var_\tau(R_{ic})$$

where $Var_\tau(R_{ic})$ is decomposed as in (4), and $W_{ci} = \sum_{c=1}^{C} \sum_{i=1}^{I} w_{ci}$. Value reported are in percentage over the total variance.

region model appears to fit slightly better as the average variance attributed to the idiosyncratic component is 0.233. Note that, while the variance decomposition for the local component in the world-region model appears to be larger than that of either of the other two models, this represents the cumulative variance for all three regional factors (within each respective region) in the former specification. The endogenous cluster and world-country models explain comparable shares of the variance of global fluctuations in the latter part of the sample while the world-region model explains very little, attributing much more comovement to the multiple regional factors. In the later part of the sample, including three factors per region seems to capture most of the cross-country commonality within a given region, leaving little to be explained at the global level. Second, the largest difference between the models takes place at the end of the sample when the world-country model gives a higher weight to the global factor while the endogenous cluster model detects a large European cluster (as noted in the previous section). Notice that the fit of the endogenous cluster model is better than the one of the world-country model in the last part of the sample, indicating that our model is better capturing the features of the data. This illustrates the advantage of using our endogenous cluster model that allows the cluster composition to change across rolling windows.

To further corroborate these insights, we perform an out-of-sample portfolio optimization exercise, as in Bekaert et al. (2009). Every two and half years, we estimate the three models using a window of five years of data and compute the model-implied variance-covariance matrix $\widehat{V}_\tau$. Following Connor and Korajczyk (1986), the factor model structure facilitates a simple computation for the inverse of the model-implied variance-covariance matrix as follows:

$$\widehat{V}_\tau^{-1} = D_\tau^{-1} - D_\tau^{-1} B_\tau V_\tau^F \left(V_\tau^F + V_\tau^F B_\tau' D_\tau^{-1} B_\tau V_\tau^F \right)^{-1} V_\tau^F B_\tau' D_\tau^{-1},$$

where D_τ is an $n \times n$ matrix with σ_{ci}^2 variances along the diagonal and zeros elsewhere, B_τ is the matrix of factor loadings, and $V_\tau^F = cov\left(\mathbf{F}\right)$. We then use the estimated variance-covariance matrix to compute the portfolio weights for the global minimum variance portfolio for the next period, as

Table 3: Out of sample performance

	Cluster	WC	WR	EW	VW
1985.1-2016.12	10.70	10.85	13.04	15.89	15.48
1985.1-1999.12	11.70	11.12	13.59	13.92	15.11
2000.1-2016.12	9.84	10.61	12.57	17.59	15.80

Note: This table reports the average ex-post volatility of five portfolios. The first three columns refer to the minimum variance portfolio that uses the estimated variance-covariance matrix from the endogenous cluster model (cluster), the world-country model (WC) and the world-region model (WR). EW denotes the equally-weighted portfolio of all the country-industry portfolios. VW denotes the value-weighted portfolio of all the country industry portfolios. All values are in annualized percentage points.

follows

$$w_{\tau+1} = \frac{\widehat{V}_{\tau}^{-1\prime} \iota}{\iota' \widehat{V}_{\tau}^{-1} \iota}$$

where ι is a vector of ones. We hold this portfolio for two years and a half and compute the ex-post volatility of the portfolio using sample excess returns. At the end of the two year and a half period, we update the portfolio weights $w_{\tau+2}$ using the estimate of the model-implied variance-covariance matrix with the updated sample, $\widehat{V}_{\tau+1}$. We repeat this process until the end of the sample and average the ex-post portfolio volatilities over the subsamples. The model that best captures the variance-covariance structure of the country-industry portfolios should minimize the ex-post volatility.

In Table 3, we report the ex-post volatility of the minimum variance portfolio over the full sample, and also over two half-subsamples. The Table also reports the ex-post volatility of two naive benchmark portfolios. The EW portfolio is constructed equally weighting all the country-industry portfolios and the VW portfolio is constructed value-weighting all the country industry portfolios. Results in Table 3 indicate that the minimum variance portfolio constructed using our endogenous cluster model, the world-country model, and the world-region model generate much lower volatility than naive strategies, both on the full sample and in the two subsamples. As for the relative performance of the three models considered here, the table shows that our endogenous cluster model is overall better able to capture the variance-covariance structure of the country-industry portfolio excess returns.

6 Robustness

Our preferred specification allows for each country-industry portfolio to belong to one of $K = 15$ potential endogenous clusters. In theory, one could search for the optimal K given some criterion such as BIC. However, as explained in Pamminger, Frühwirth-Schnatter et al. (2010), determining the appropriate penalty factor for computing the BIC in these types of models is problematic and BIC often overfits the number of clusters when working with large panels of data. Alternatively, Ando and Bai (2017) propose a new panel information criterion to select the number of factors in a similar model framework but in a frequentist setting. However, given that we examine the time-varying nature of cross-country and cross-industry comovement in stock returns, it is likely the case that the appropriate number of clusters changes over time. Given preliminary analysis over the full sample, we elect to include 15 clusters as this produces reasonable groupings that are fairly straightforward to interpret and consistent across time. In order to assess the robustness of our modeling choice, we considered a variety of clustering possibilities with $K = \{12, 18, 21\}$ across the rolling-window subsamples. The cluster relationships are comparable across models with slight variation, which is to be expected. Similarly to the model comparison presented in Section 5, Table 4 reports the ex-post volatility of the minimum variance portfolio over the full sample and two half-subsamples for various values of K.

All specifications generate similar volatilities when computed across the full sample. However, we find clear evidence of time variation in the number of clusters necessary to accurately capture the comovement across portfolios with fewer clusters, and thus more integration, in the latter sub-period. We are more interested in accurately describing this comovement for the most recent period and thus place greater consideration on model performance in the 2000-2016 sub-period. While $K = 18$ produces a slightly lower volatility than $K = 15$, the suggested clusters are less clearly

Table 4: Out of sample performance - Comparing the number of clusters

	15 Clusters	12 Clusters	18 Clusters	21 Clusters
1985.1-2016.12	10.70	10.41	10.47	10.70
1985.1-1999.12	11.70	10.86	11.25	11.27
2000.1-2016.12	9.84	10.02	9.80	10.21

Note: This table reports the average ex-post volatility of five portfolios. The first three columns refer to the minimum variance portfolio that uses the estimated variance-covariance matrix from the endogenous cluster model (cluster), the world-country model (WC) and the world-region model (WR). EW denotes the equally-weighted portfolio of all the country-industry portfolios. VW denotes the value-weighted portfolio of all the country industry portfolios. All values are in annualized percentage points.

defined. Thus, given only marginal differences in relative performance, our choice to estimate $K = 15$ clusters does not appear to be too restrictive or arbitrary.

7 Conclusion

Much of the evidence on the construction of international portfolios has suggested that diversifying across countries is a better strategy than diversifying across industries. In this paper, we utilize a factor model with endogenous clustering to examine international stock return comovements of country-industry portfolios. We find that country-industry portfolios tend to cluster mainly within geographical areas and that the optimal portfolios are not simply country-level aggregates but may also be continental or sub-continental. This suggests greater potential benefits from diversifying across geographical areas rather than across sectors.

Our rolling window results show that the cluster component was the main driver of country-industry portfolio returns for most of the sample, except from mid-2000 to mid-2010s when the global component had a more prominent role. At the end of the sample, a large cluster among European countries emerges. Comparison with benchmark models highlights the importance of allowing for endogenous clusters that can change over time.

References

Ando, Tomohiro, and Jushan Bai (2017) 'Clustering huge number of financial time series: A panel data approach with high-dimensional predictors and factor structures.' *Journal of the American Statistical Association* 112(519), 1182–1198

Baele, L., and K. Inghelbrecht (2009) 'Time-varying integration and international diversification strategies.' *Journal of Empirical Finance* 16(3), 368–387

Bekaert, G., R.J. Hodrick, and X. Zhang (2009) 'International stock return comovements.' *The Journal of Finance* 64(6), 2591–2626

Brooks, R., and M. Del Negro (2004) 'The rise in comovement across national stock markets: market integration or it bubble?' *Journal of Empirical Finance* 11(5), 659–680

——— (2005) 'Country versus region effects in international stock returns.' *The Journal of Portfolio Management* 31(4), 67–72

——— (2006) 'Firm-level evidence on international stock market comovement.' *Review of Finance* 10(1), 69–98

Cappiello, Lorenzo, Arjan Kadareja, and Simone Manganelli (2010) 'The impact of the euro on equity markets.' *Journal of Financial and Quantitative Analysis* 45(2), 473–502

Carter, Chris K, and Robert Kohn (1994) 'On gibbs sampling for state space models.' *Biometrika* 81(3), 541–553

Casella, George, and Edward I George (1992) 'Explaining the gibbs sampler.' *The American Statistician* 46(3), 167–174

Cavaglia, S., C. Brightman, and M. Aked (2000) 'The increasing importance of industry factors.' *Financial Analysts Journal* 56(5), 41–54

Conner, Gregory, and Robert Korajczyk (1986) 'Performance measurement with the arbitrage pricing theory.' *Journal of Financial Economics* 15, 373–394

Connor, Gregory, and Robert A Korajczyk (1986) 'Performance measurement with the arbitrage pricing theory: A new framework for analysis.' *Journal of financial economics* 15(3), 373–394

Eiling, E., B. Gerard, P. Hillion, and F.A. de Roon (2012) 'International portfolio diversification: Currency, industry and country effects revisited.' *Journal of International Money and Finance* 31, 1249–1278

Francis, Neville, Michael T. Owyang, and Daniel Soques (2017a) 'Business cycles across space and time.' *manuscript*

Francis, Neville, Michael T Owyang, and Ozge Savascin (2017b) 'An endogenously clustered factor approach to international business cycles.' *Journal of Applied Econometrics* 32(7), 1261–1276

Fruhwirth-Schnatter, Sylvia, and Sylvia Kaufmann (2008) 'Model-based clustering of multiple time series.' *Journal of Business and Economic Statistics* 26(1), 78–89

Gelfand, Alan E, and Adrian FM Smith (1990) 'Sampling-based approaches to calculating marginal densities.' *Journal of the American Statistical Association* 85(410), 398–409

Griffin, J.M., and A.G. Karolyi (1998) 'Another look at the role of the industrial structure of markets for international diversification strategies.' *Journal of Financial Economics* 50(3), 351–373

Hamilton, James D, and Michael T Owyang (2012) 'The propagation of regional recessions.' *Review of Economics and Statistics* 94(4), 935–947

Hardouvelis, Gikas A, Dimitrios Malliaropulos, and Richard Priestley (2007) 'The impact of emu on the equity cost of capital.' *Journal of International Money and Finance* 26(2), 305–327

Hernández-Murillo, Rubén, Michael T. Owyang, and Margarita Rubio (2017) 'Clustered housing cycles.' *Regional Science and Urban Economics* 66, 185–197

Heston, S.L., and K.G. Rouwenhorst (1994) 'Does industrial structure explain the benefits of international diversification?' *Journal of Financial Economics* 36(1), 3–27

Kose, M Ayhan, Christopher Otrok, and Charles H Whiteman (2003) 'International business cycles: World, region, and country-specific factors.' *American Economic Review* 93(4), 1216–1239

Lessard, D.R. (1974) 'World, national, and industry factors in equity returns.' *The Journal of Finance* 29(2), 379–391

Moerman, Gerard A (2008) 'Diversification in euro area stock markets: Country versus industry.' *Journal of International Money and Finance* 27(7), 1122–1134

Otrok, Christopher, and Charles H Whiteman (1998) 'Bayesian leading indicators: measuring and predicting economic conditions in iowa.' *International Economic Review* pp. 997–1014

Pamminger, Christoph, Sylvia Frühwirth-Schnatter et al. (2010) 'Model-based clustering of categorical time series.' *Bayesian Analysis* 5(2), 345–368

Pukthuanthong, K., and R. Roll (2009) 'Global market integration: An alternative measure and its application.' *Journal of Financial Economics* 94(2), 214–232

Roll, R. (1992) 'Industrial structure and the comparative behavior of international stock market indices.' *The Journal of Finance* 47(1), 3–41

Tanner, Martin A., and Wing Hung Wong (1987) 'The calculation of posterior distributions by data augmentation.' *Journal of the American Statistical Association* 82(398), 528–540

Troughton, Paul T, and Simon J Godsill (1997) 'Bayesian model selection for time series using markov chain monte carlo.' In 'Acoustics, Speech, and Signal Processing, 1997. ICASSP-97., 1997 IEEE International Conference on,' vol. 5 IEEE pp. 3733–3736