Available online at www.sciencedirect.com
ScienceDirect

Journal of Empirical Finance 13 (2006) 396 – 416

Journal of EMPIRICAL FINANCE

www.elsevier.com/locate/jempfin

Geographic versus industry diversification: Constraints matter ☆

Paul Ehling [a,*], Sofia B. Ramos [b]

[a] *Department of Financial Economics, BI Norwegian School of Management, Nydalsveien 37, 0442 Oslo, Norway*
[b] *Department of Finance and Accounting, ISCTE-Business School, Portugal*

Accepted 21 April 2006
Available online 14 July 2006

Abstract

This research addresses whether geographic diversification provides benefits over industry diversification in the Eurozone. Our contribution is to show that in the absence of constraints, no empirical evidence is found to support the argument that geographic diversification dominates industry diversification, except in the euro subperiod. With short-selling constraints, however, the tangency portfolio of geographic diversification is not attainable by industry diversification. In out-of-sample geographic minimum variance portfolios outperform industry portfolios in economic terms, although we cannot establish statistical significance.
© 2006 Elsevier B.V. All rights reserved.

JEL classification: G11; G15
Keywords: Block-bootstrap tests; Diversification gains; EMU; Geographic diversification; Industry diversification

1. Introduction

The goal of this paper is to examine the performance of geographic and industry allocation in the Eurozone. Many professionals and researchers assumed the advent of the Euro will decrease the dominance of geographic diversification. We evaluate this view using a new approach that allows us direct comparison of the diversification strategies from a pure performance perspective.

The traditional choice of geographic diversification has come into question from two directions. First, events such as the deregulation of markets or the elimination of international barriers to capital movements are considered as catalysts of market integration that may, therefore, affect

☆ This paper supersedes “Geographical versus Industrial Diversification: A Mean Variance Spanning Approach”.

* Corresponding author. Tel.: +47 464 10 505; fax: +47 464 10 801.
E-mail address: paul.ehling@bi.no (P. Ehling).

0927-5398/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.jempfin.2006.04.001

the dominance of country factors. The prime example is the European Monetary Union or EMU. It is often argued that country-specific policy shocks will be dampened down as a result of a single monetary policy and coordinated fiscal policies constrained by the Stability and Growth Pact. Also Carrieri et al. (2004) demonstrate empirically that country integration does not rule out industry segmentation. Therefore, according to the market integration perception, the advent of the *Euro* marks the end of the geographic diversification success story and the start of the industry diversification dominance.

Second, a number of empirical studies have found mixed results as to country and industry effects in equity returns. Some of these studies provide evidence that industry factors are already more important than country factors (Cavaglia et al., 2000; Galati and Tsatsaronis, 2001). The results of others suggest only that industry effects are becoming increasingly important while countries are losing explanatory power (Baca et al., 2000; Brooks and Del Negro, 2004a; Isakov and Sonney, 2004). Most of the literature relies on the Heston and Rouwenhorst (1994) methodology[1], which itself has come under criticism because of the severe restrictions on the underlying factor model (see Brooks and Del Negro, 2004b). In a new approach, Carrieri et al. (2005) demonstrate that changes in the benefits of geographic and industry allocation indicate changes on the real economy side. Hence, any EMU impact on the benefits of geographic versus industry diversification should be reflected in our tests.

We use a mean-variance efficiency test proposed by Basak et al. (2002) (hereafter BJS (2002)) to address the performance of both diversification strategies. The test measures the difference between the variance of a benchmark, for instance tangency and minimum variance portfolios, and a mimicking portfolio of the corresponding diversification strategy with identical returns. Our approach allows us direct comparison of the two diversification strategies without relying on unobservable country or industry factors. It is also possible to compare either strategy with any other benchmark of interest. Further, the test is flexible enough to incorporate short sale constraints or bounds on portfolio positions, which are not analyzed in the literature. Finally, block-bootstrapping the *p*-values show that the BJS (2002) test is robust in small samples and never reverses our conclusions.

We find unconstrained geographic and industry diversification at the tangency portfolios to be statistically equivalent. However, the signs of the BJS (2002) test indicate that the industry efficient frontier dominates the country frontier. These results are reinforced at the corresponding minimum variance portfolios. Interestingly, it is the very last subperiod associated with the advent of the *Euro* currency where the country frontier lies outside of the industry frontier. Thus, in the case of unconstrained diversification, we find, contrary to the market integration perception, only in the *euro* subperiod empirical evidence to support the argument that geographic diversification outperforms industry diversification.

Introducing short-selling constraints shrinks the industry frontier dramatically[2]. In particular, industry portfolios cannot attain the country tangency portfolio return. Nevertheless, the tiny industry frontier lies partially again outside of the country frontier and significantly so. This small part of the industry frontier includes both the tangency and the minimum variance portfolio. That is, the industry efficient frontier lies almost entirely inside the country frontier, spans a much

[1] See also Errunza and Padmanablan (1988), Grinold et al. (1989), Beckers et al. (1992), Drummen and Zimmermann (1992), Roll (1992), Arshanapalli et al. (1997), Griffin and Karolyi (1998), Rouwenhorst (1999) and Heckman et al. (1998).

[2] Gerard et al. (2002) and Eiling et al. (2005) use a variant of mean-variance spanning tests to addresses the relative performance of geographic and industry Sharpe ratios. Their results are overall consistent with our conclusions. They argue that without short sales restrictions, country and industry diversification are a redundant strategy relative to one another. Additionally, their industry portfolios are also more affected by short sales constraints than country portfolios.

smaller range of returns, but "peaks" outside of the country frontier. Again the *euro* subperiod is distinct from the full sample as well as other subsamples. Now, however, the country efficient frontier is more affected by the short-selling constraint and, hence, spans a smaller range of returns than the industry frontier.

An explanation for the underperformance of country allocation in the absence of constraints is that industry indexes are more exposed to the dominant single factor (the EMU index) in equity returns than country indexes. This suggests, according to Green and Hollifield (1992), that industry portfolios are better suited than country portfolios to eliminate factor risk. It also implies that industry diversification must be affected to a much greater extent than geographic diversification when short-selling constraints are imposed. In the *euro* subperiod the exposure of the country and industry indexes to the EMU factor is reversed.

As for comparison with other benchmarks, the efficiency of a passive diversification strategy like the EMU index is always rejected. In other words, both geographic and industry diversification have a significantly lower standard deviation than the EMU index[3]. A simultaneous cross-country and cross-industry diversification based on the ten Datastream level three sector indexes within each of the eleven EMU countries, up to 110 indexes, clearly outperforms both industry and geographic diversification in an unconstrained or constrained optimization.

Additional results based on block-bootstrapping or on subsample analysis suggest that it is very difficult to discriminate between geographic diversification and industry diversification. This paper proceeds as follows: in Section 2, we present the data and the BJS (2002) test. Section 3 reports the empirical results. Section 4 contains an extensive robustness analysis. Concluding remarks follow in Section 5. Finally, we describe our block-bootstrap analysis in Appendix A.

2. Data and methodology

2.1. Data

We use Datastream country and sector indexes for the eleven EMU constituents: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, The Netherlands, Portugal, and Spain. The data provided by Datastream include weekly US dollar-denominated returns from January 12, 1990 to July 29, 2005 (representing 812 observations)[4].

The sample is divided into different subperiods not only to capture time patterns but also to incorporate effects attributable to actions of the new monetary authority. The subperiods are labeled *pre-convergence*, *convergence*, and *euro*. The starting date of the *convergence* subperiod is associated with the signing of the Maastricht treaty, and the end of the subperiod (December 31, 1998) is associated with the fixing of the conversion rates.

We focus the analysis on level three of the Datastream sector classification, which represents ten sectors per country: Basic Industries, Cyclical Goods, Cyclical Services, Financials, General Industries, Information Technology, Non-cyclical Consumer Goods, Non-cyclical Consumer Services, Resources, and Utilities. Using this industry classification, we compute ten industry indexes by building market value-weighted indexes, e.g. we aggregate the industry indexes across the eleven EMU countries.

[3] This is, of course, not surprising since both diversification strategies essentially eliminate the EMU factor risk.

[4] The starting date of our sample is motivated by data availability for all EMU participants and freedom of capital movements. The movement toward full liberalization of capital movements in Europe starts only in the late 1970s and early 1980s. Complete freedom of capital movements was attained only in the late 1980s (Bakker, 1996).

We also employ the 110 industry indexes, ten sectors across each of the eleven countries, without aggregating them across countries. This portfolio strategy at a lower level of index aggregation represents a simultaneous cross-country and cross-industry diversification. We call this portfolio strategy "country–industry" diversification. But we limit these country–industry indexes to the indexes that have observations over the full period (subperiods). Finally, we form an EMU market portfolio constructed from the Datastream country or our industry indexes.

2.2. Methodology

We measure the mean-variance efficiency of a diversification strategy using the Basak et al. (2002) test. Consider the existence of a benchmark asset or portfolio with return r and assume well-defined first two moments, $E(r_t)=\beta$ and $\mathrm{Var}(r_t)=v$, for $t\in[0, T]$. The matrix R contains the returns of a set of primitive assets or portfolios with expected return $E(R_t)=\mu$ and covariance matrix $\mathrm{Cov}(R_t)=\sum$. Mimicking portfolios are constructed from the primitive assets.

Let w denote the vector of portfolio weights of the primitive assets in the mimicking portfolio and $\overline{1}$ represents a vector of ones. A mimicking portfolio with expected return equal to the benchmark can be obtained from the following unconstrained optimization problem:

$$\begin{aligned} &\min_{w} w'\sum w - v \\ &\text{s.t. } w'\overline{1} = 1 \\ &w'\mu = \beta. \end{aligned} \tag{1}$$

The return of the mimicking portfolio r_t^{β} equals, if it exists, by construction the return on the benchmark asset $E(r_t^{\beta})=\beta$ and its variance is $\mathrm{Var}(r_t^{\beta})$. The measure of efficiency of the mimicking portfolio with respect to the benchmark is defined as $\lambda=\mathrm{Var}(r_t^{\beta})-v$. Whenever λ is positive, the mimicking portfolio has a higher variance than the benchmark, and therefore it is mean-variance inefficient. Conversely, a negative value implies that the mimicking portfolio is mean-variance efficient.

BJS (2002) show that the test statistic under the null hypothesis $\lambda_0=0$ takes the form

$$\xi = \frac{T^{1/2}(\lambda - \lambda_0)}{\lambda_\sigma} = \frac{T^{1/2}\lambda}{\lambda_\sigma} \tag{2}$$

where λ is the sample measure of efficiency, λ_σ denotes the standard deviation of the measure of efficiency, and T is the sample size. The test statistic is standard normally distributed for large T.

A major strength of this approach is that one can easily incorporate common constraints faced by investors and portfolio managers and that represent important obstacles in the overall allocation decision. In particular, we consider the following optimization problem with short-selling constraints[5]:

$$\begin{aligned} &\min_{w} w'\sum w - v \\ &\text{s.t. } w'\overline{1} = 1 \\ &w'\mu = \beta \\ &w \geq 0. \end{aligned} \tag{3}$$

Now the measure of efficiency for both problems is the value of the Lagrangian L:

$$\lambda = L = w'\sum w - v + \delta_1(w'\overline{1} - 1) + \delta_2(w'\mu - \beta) - \delta_3' w \tag{4}$$

where δ_1, δ_2, and δ_3 represent Lagrange multipliers. The multiplier δ_3 is active only for the problem in (3).

[5] The BJS (2002) test can also be performed with lower and upper bounds on portfolio weights, i.e., $l \leq w \leq u$.

We perform the following steps to test the efficiency of geographic and industry diversification. First, we solve for the unconstrained as well as the constrained tangency and minimum variance portfolios for country and industry diversification in order to construct benchmark portfolios. Second, we solve for each of the benchmarks the problems in (1) and (3) by way of utilizing the rival strategy as mimicking portfolio. For example, when the country tangency portfolio constitutes the benchmark, then the industry indexes form the mimicking portfolio, and conversely when the industry tangency constitutes the benchmark. Third, we perform the BJS (2002) test to infer the mean-variance efficiency of diversification strategies.

Investors, of course, are not just limited to geographic or industry diversification. Therefore, we would like to examine the efficiency of geographic and industry diversification against other investment strategies. First, we consider an active diversification strategy at a lower level of index aggregation: country–industry indexes. Again these country–industry indexes are respectively the ten industry indexes within each of the countries. Second, we address the efficiency of a passive strategy such as the market portfolio relative to geographic and industry diversification. Our market index is a value-weighted index constructed from the Datastream country or our industry indexes.

To sum up: aside from the four country or industry motivated benchmarks (country tangency portfolio, country minimum variance portfolio, industry tangency portfolio, and industry minimum variance portfolio) we also consider the country–industry tangency portfolio, the country–industry minimum variance portfolio, and the EMU market portfolio as benchmark portfolios. Mimicking portfolios are always constructed either from industry or country indexes.

Portfolio efficiency with constraints is also analyzed in Kandel et al. (1995), Wang (1998), and Li et al. (2003). As for spanning tests (Huberman and Kandel, 1987) our approach focuses on mean-variance efficiency, but the BJS (2002) test provides a major advantage over the traditional spanning test as we can test directly one strategy against the other[6]. Notice that our comparison of country and industry indexes with country–industry indexes (ten indexes against say 110 indexes) based on the BJS (2002) test can be interpreted as an informal spanning test.

Finally, to assess the small-sample properties of the BJS (2002) test, we introduce an estimation-based block-bootstrap BJS test. The bootstrap data is then employed to evaluate the importance of estimation error in portfolio weights as well as to address the out-of-sample performance of geographic and industry diversification. The block-bootstrap simulations are described in detail in the appendix.

3. Empirical results

3.1. Descriptive statistics

Table 1, Panel A, reports basic descriptive statistics for the eleven country index returns and their averages in this study. In Panel B we report the ten industry indexes and the average of the ten industries. We notice that the average annualized country index return, 6.75%, is slightly higher than the industry average, 6.14%. The average standard deviation of countries, 20.68%, is also higher than the average standard deviation of industry indexes, 19.57%.

Table 1 also shows average correlations for our country Corr(C) and industry indexes Corr(I). We exclude the correlation of each index with itself from Corr(C) and Corr(I). Average

[6] The spanning test sets, by construction, the benchmark as a combination of geographic and industry diversification. This is, however, an implausible diversification strategy since each single stock shows, then, up once in a country portfolio and a second time in an industry portfolio. See also the Gerard et al. (2002) mean-variance spanning test which helps to circumvent this problem.

Table 1
Descriptive statistics of country and industry indexes

	Mean	Stdv.	Corr(C)	Corr(I)
Panel A: country indexes				
Austria	1.71%	18.42%	0.47	0.38
Belgium	5.58%	16.81%	0.58	0.48
Finland	11.12%	30.58%	0.42	0.51
France	6.31%	17.88%	0.63	0.54
Germany	4.17%	19.23%	0.65	0.54
Greece	14.02%	30.59%	0.38	0.63
Ireland	8.68%	18.64%	0.51	0.65
Italy	4.61%	21.24%	0.54	0.69
Netherlands	6.74%	17.02%	0.63	0.74
Portugal	3.12%	18.13%	0.52	0.76
Spain	8.16%	18.97%	0.61	0.76
Average	6.75%	20.68%	0.54	0.61
Panel B: industry indexes				
Basic industries	4.34%	17.45%	0.66	0.71
Cyclical consumer goods	3.22%	20.31%	0.63	0.68
Cyclical services	4.69%	17.49%	0.68	0.72
Financials	4.75%	18.53%	0.70	0.73
General industries	4.13%	19.27%	0.67	0.73
Information technology	7.43%	31.73%	0.54	0.55
Non-cyclical consumer goods	8.01%	15.12%	0.58	0.63
Non-cyclical services	7.49%	21.69%	0.59	0.61
Resources	8.93%	19.27%	0.44	0.49
Utilities	8.44%	14.87%	0.57	0.59
Average	6.14%	19.57%	0.61	0.64

This table reports the following basic descriptive statistics for the index returns in this study: annualized mean (Mean), annualized standard deviation (Stdv.), and average correlations (Corr(C) and Corr(I)). Panel A contains the eleven EMU country indexes and the average of the eleven countries. In Panel B we report the ten industry indexes and the average of the ten industries. The continuously compounded weekly returns are calculated in US dollar for the period January 12, 1990 until July 29, 2005, with a sample size of 812 observations. Corr(C) and Corr(I) denote the average Pearson correlation of each index with the eleven country indexes and the ten industry indexes respectively. The correlation of each index with itself is excluded from Corr(C) and Corr(I). Industry indexes are computed by building market value-weighted indexes based on Datastream level three sector classification. Source: Datastream and own calculations.

correlations within countries, 0.54, tend to be lower than average correlations between countries and industries, 0.61. Furthermore, we note that average correlations for industries are higher, 0.64, than the average correlation of an industry with countries, 0.61. Also country indexes, 0.54, are on average much less correlated than industry indexes, 0.64.

Both pieces of evidence, the higher average return and the lower average correlation of country indexes, suggest that country diversification should outperform industry diversification.

3.2. Geographic versus industry diversification under unconstrained maximization

Now we apply the BJS (2002) test to solve the problem in (1) where geographic and industry diversification are used in turn as benchmarks and tested against the opposite strategy.

The annualized means and standard deviations of the benchmark portfolios (tangency portfolios and minimum variance portfolios) for each diversification strategy (country indexes, industry indexes, and country–industry indexes) are shown in Table 2. Unconstrained benchmark portfolio strategies

Table 2
Mean and standard deviation of benchmark portfolios: tangency and minimum variance portfolios

Sample			Pre-convergence		Convergence		Euro	
	1990–2005		1990–1994		1995–1998		1999–2005	
	Mean	Stdv.	Mean	Stdv.	Mean	Stdv.	Mean	Stdv.
Panel A: benchmark portfolios without constraints								
Tangency portfolios								
Countries	21.33%	28.17%	44.53%	41.42%	74.17%	31.36%	49.08%	41.46%
Industries	14.11%	17.18%	64.14%	39.43%	59.89%	22.81%	33.74%	39.33%
C–I indexes	50.38%	24.37%	117.36%	24.66%	85.10%	11.87%	202.68%	48.54%
Minimum variance portfolios								
Countries	5.16%	13.86%	3.89%	12.24%	10.91%	12.03%	4.88%	13.08%
Industries	8.45%	13.30%	5.96%	12.02%	14.86%	11.36%	4.03%	13.59%
C–I indexes	9.94%	10.83%	7.20%	6.11%	20.78%	5.87%	7.55%	9.37%
Panel B: benchmark portfolios with short-selling constraints								
Tangency portfolios								
Countries	10.45%	17.89%	8.67%	15.99%	23.77%	13.71%	6.14%	14.42%
Industries	8.42%	13.70%	6.10%	13.81%	24.47%	14.02%	8.58%	19.12%
C–I indexes	13.36%	13.90%	22.88%	14.38%	35.88%	11.21%	18.88%	13.64%
Minimum variance portfolios								
Countries	4.89%	14.19%	3.65%	12.52%	15.05%	12.65%	3.43%	13.50%
Industries	8.01%	13.59%	5.05%	12.81%	17.68%	12.22%	2.60%	14.19%
C–I indexes	8.37%	11.81%	4.68%	7.82%	15.08%	7.74%	6.60%	11.13%
Panel C: market index								
EMU market index	5.99%	16.40%	0.14%	14.94%	18.58%	14.85%	2.80%	18.21%

The table reports annualized mean (Mean) and annualized standard deviation (Stdv.) of tangency and minimum variance portfolios for country, industry and country–industry (C–I) indexes. Panel A contains unconstrained portfolios. Panel B contains portfolios with short-selling constraints. Panel C reports the descriptive statistics for a market value-weighted index of the sample countries. Country–industry indexes (CI indexes) are respectively the ten Datastream level three sector indexes within each of the eleven EMU countries. C–I indexes represent a simultaneous cross-country and cross-industry diversification strategy. The subperiods pre-convergence, convergence, and euro include 260, 208, and 344 observations respectively. Source: Datastream and own calculations.

are reported in Panel A while benchmark portfolios with no short sales constraints are shown in Panel B. Table 2, Panel C, contains also the corresponding data on the EMU market index.

Again a naive view on the annualized mean returns of the tangency portfolios suggests that countries should outperform industries. This is because tangency returns over the full sample period and in the *convergence* and *euro* subperiods for unconstrained portfolios and over the full sample period and the *pre-convergence* subperiod for no short sales portfolios are higher than the corresponding industry tangency return. The picture for minimum variance portfolios is just the opposite: only in the *euro* subperiod, for both constrained and unconstrained portfolios, the country return is higher than the industry return. Here we can already see that it is important to test for differences in the portfolios since the level of returns contain insufficient information to unambiguously rank the diversification strategies. On top of that, even if it appears that country tangency portfolio returns are higher it does not necessarily mean that the country efficient frontier lies outside of the industry frontier.

Fig. 1 highlights exactly the aforementioned issues. The figure shows the unconstrained country, the industry, and the country–industry unconstrained efficient frontier for the full sample. As expected, the industry frontier lies for the largest part of the spanned returns inside of the

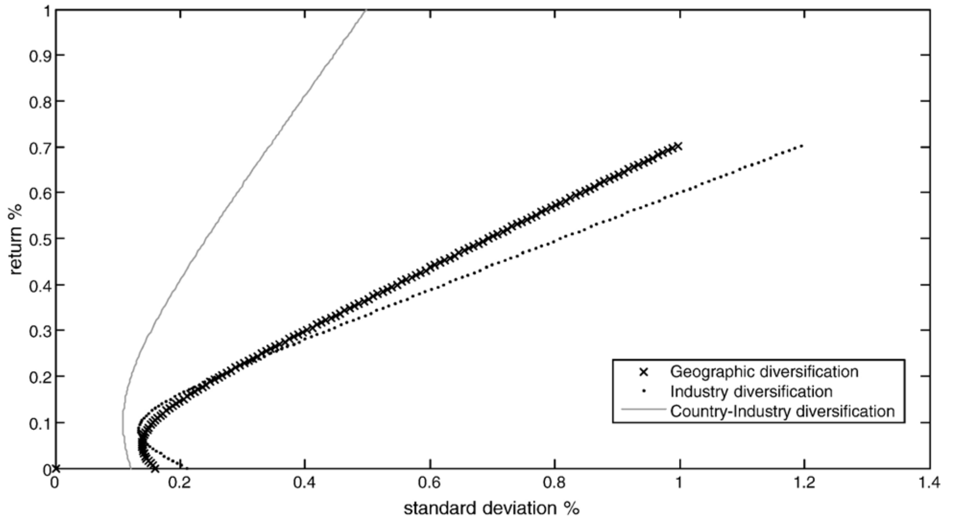

Fig. 1. Unconstrained mean variance frontiers. This figure displays the unconstrained mean variance frontiers of geographic, industry and cross-country and cross-industry (country–industry) diversification. Industry indexes are computed by building market value-weighted indexes based on Datastream level three sector classification. Returns and standard deviations are annualized. The sample period is 1990–2005. Source: Datastream and own calculations.

country frontier. However, the country and industry frontiers cross two times such that the relevant part of the industry frontier, between the minimum variance portfolio and the tangency portfolio, lies outside of the country frontier. This justifies the usefulness of employing a test rather than to rely on unobservable country and industry factors or just on correlations.

Tables 3 and 4 present the BJS (2002) test results. Because the sign of the test statistic is identical to the sign of the measure of efficiency, λ, we do not report λ. Again a positive sign implies that the benchmark strategy is mean-variance efficient, and, a negative sign implies that it is mean-variance inefficient[7].

In Table 3, Panel A, the benchmark strategy is geographic diversification, while industry indexes represent the mimicking assets. The test statistic for the tangency portfolio is slightly negative, but insignificantly so, suggesting that the industry frontier lies outside of the country frontier. At the minimum variance portfolio we find that the country return is so low, compared to industry, that the corresponding point on the industry frontier lies below the industry minimum variance portfolio[8].

In Panel B, industry diversification represents the benchmark, and countries play the role of the mimicking portfolio. Both test statistics for the tangency portfolio and the minimum variance portfolio on the industry frontier are positive, although only the test at the minimum variance portfolio shows statistical significance[9] at the 10% confidence level. Therefore, we cannot reject the null that diversification strategies, country and industry, exhibit identical variance around their

[7] We employ the p-values from the Normal distribution to indicate significant test statistics since the block-bootstrap experiments never reverse our conclusions.

[8] In the tables we mark tests with a (b) when the mimicking portfolio is below the minimum variance portfolio of the mimicking strategy and do not report test statistics since we can find another portfolio with identical risk but higher return.

[9] Note that a negative sign in Panel A does not necessarily imply a positive sign in Panel B, and vice-versa. This is because the tangency portfolios or the minimum variance portfolios show different mean returns, so the frontiers can, in principle, cross.

Table 3
Efficiency test for geographic and industry diversification

Sample		Pre-convergence	Convergence	Euro
	1990–2005	1990–1994	1995–1998	1999–2005
Panel A: benchmark: countries, mimicking portfolio: industries				
Tangency portfolios	−0.00	−2.21**	−0.87	0.81
Minimum variance portfolios	(b)	(b)	(b)	0.94
Panel B: benchmark: industries, mimicking portfolio: countries				
Tangency portfolios	0.77	1.09	0.87	−1.58
Minimum variance portfolios	1.89*	0.46	0.97	(b)
Panel C: benchmark: C–I indexes, mimicking portfolio: countries				
Tangency portfolios	2.23**	2.15**	3.82***	2.68***
Minimum variance portfolios	2.26**	4.70***	4.26***	4.61***
Panel D: benchmark: C–I indexes, mimicking portfolio: industries				
Tangency portfolios	1.80*	3.13***	1.59***	2.00**
Minimum variance portfolios	2.42**	7.54***	6.03***	3.42***
Panel E: benchmark: market index, mimicking portfolio: countries				
EMU market index	− 5.91***	−2.95***	−3.29***	(b)
Panel F: benchmark: market index, mimicking portfolio: industries				
EMU market index	(b)	−3.43***	−2.78***	(b)

The table presents the test statistics of the BJS (2002) efficiency test. The standard normally distributed test measures the difference between the variance of a benchmark and a mimicking portfolio with identical returns. The sign of the test should be interpreted as follows: a positive value implies that the mimicking portfolio is mean-variance inefficient. A negative value implies that the mimicking portfolio is mean-variance efficient. In Panel A country indexes form the tangency portfolio and the minimum variance portfolio and industry indexes represent the mimicking portfolios. In Panel B industry indexes form the tangency portfolio and the minimum variance portfolio and country indexes represent the mimicking portfolios. In Panel C country–industry indexes (C–I indexes) form the tangency portfolio and the minimum variance portfolio and country indexes represent the mimicking portfolios. In Panel D country–industry indexes (C–I indexes) form the tangency portfolio and the minimum variance portfolio and industry indexes represent the mimicking portfolios. In Panel E the EMU market index represents the benchmark and country indexes form the mimicking portfolios. In Panel F the EMU market index represents the benchmark and industry indexes form the mimicking portfolios. Source: Datastream and own calculations. Superscripts **, **, *** denote statistical significance at 10%, 5%, and 1% levels respectively. (a) indicates that the test is not feasible because the mimicking portfolio cannot attain the benchmark return. (b) indicates that the mimicking portfolio is an inefficient portfolio.

respective tangency portfolios but do so at the minimum variance portfolio. Also notice that the relevant part of the industry frontier, between minimum variance portfolio and tangency portfolio, lies entirely outside of the country frontier.

Gerard et al. (2002) also find that geographic diversification and industry diversification are always redundant strategies relative to each other using a spanning methodology. In a follow up paper Eiling et al. (2005) come to the same conclusion. Notice that results in both papers are based on Sharpe ratios and do not address minimum variance portfolios.

The results for the *pre-convergence* and the *convergence* subperiod reinforce our overall findings. Test statistics, in Table 3, Panels A and B, for geographic diversification are negative, at the tangency portfolio, or so low, at the minimum variance portfolio, that the corresponding point on the industry frontier is inefficient. In the *pre-convergence* subperiod the negative test statistics for the country tangency portfolio is even significant at the 5% confidence level. Again, the test

Table 4
Efficiency tests for geographic and industry diversification with short-selling constraints

Sample		Pre-convergence	Convergence	Euro
	1990–2005	1990–1994	1995–1998	1999–2005
Panel A: benchmark: countries, mimicking portfolio: industries				
Tangency portfolios	(a)	(a)	−0.04	0.50
Minimum variance portfolios	(b)	(b)	(b)	1.07
Panel B: benchmark: industries, mimicking portfolio: countries				
Tangency portfolios	1.85*	−0.96	0.08	(a)
Minimum variance portfolios	2.12**	−0.50	0.78	(b)
Panel C: benchmark: C–I indexes, mimicking portfolio: countries				
Tangency portfolios	0.54	(a)	(a)	(a)
Minimum variance portfolios	1.79*	4.96***	3.93***	(a)
Panel D: benchmark: C–I indexes, mimicking portfolio: industries				
Tangency portfolios	(a)	(a)	(a)	(a)
Minimum variance portfolios	0.84	5.66***	5.17***	1.68*
Panel E: benchmark: market index, mimicking portfolio: countries				
EMU market index	−5.31***	−2.74***	−3.60***	−5.18***
Panel F: benchmark: market index, mimicking portfolio: industries				
EMU market index	(b)	−2.25**	−2.86***	(b)

The table presents the test statistics of the BJS (2002) efficiency test with short-selling constraints. The standard normally distributed test measures the difference between the variance of a benchmark and a mimicking portfolio with identical returns. The sign of the test should be interpreted as follows: a positive value implies that the mimicking portfolio is mean-variance inefficient. A negative value implies that the mimicking portfolio is mean-variance efficient. In Panel A country indexes form the tangency portfolio and the minimum variance portfolio and industry indexes represent the mimicking portfolios. In Panel B industry indexes form the tangency portfolio and the minimum variance portfolio and country indexes represent the mimicking portfolios. In Panel C country–industry indexes (C–I indexes) form the tangency portfolio and the minimum variance portfolio and country indexes represent the mimicking portfolios. In Panel D country–industry indexes (C–I indexes) form the tangency portfolio and the minimum variance portfolio and industry indexes represent the mimicking portfolios. In Panel E the EMU market index represents the benchmark and country indexes form the mimicking portfolios. In Panel F the EMU market index represents the benchmark and industry indexes form the mimicking portfolios. Source: Datastream and own calculations. Superscripts *, **, *** denote statistical significance at 10%, 5%, and 1% levels respectively. (a) indicates that the test is not feasible because the mimicking portfolio cannot attain the benchmark return. (b) indicates that the mimicking portfolio is an inefficient portfolio.

statistics with industry diversification as the benchmark, both in the *pre-convergence* and the *convergence* subperiod, provide supportive evidence for the fact that the industry frontier lies outside the country frontier in the mean-variance space. However, none of the positive test statistics comes even close to be significant.

Interestingly, this picture is reversed for the *euro* subperiod. That is, when country indexes form the benchmark both tests at the tangency and the minimum variance portfolio show positive but insignificant statistics. The opposite is true when industry portfolios represent the benchmark. Again we do not obtain statistically significant results although the industry minimum variance portfolio exhibits so low returns that the corresponding portfolio on the country frontier is mean-variance inefficient.

Next we analyze country–industry diversification. We find, in Table 3, Panels C and D, for both country and industry diversification as mimicking strategy that the test statistics are always positive and statistically significant at least at the 10% confidence level. In other words, a diversification strategy based on country–industry indexes clearly outperforms geographic as well as industry diversification.

The same result, obtained by comparing Sharpe ratios, leads Adjaouté and Danthine (2002) to speculate on whether cost structure or a behavioral explanation is more appropriate to explain the fact that portfolio managers have favored the geographic diversification model.

Lastly, we address the efficiency of a passive strategy such as the EMU market portfolio relative to geographic and industry diversification. Our EMU market index is a value-weighted index constructed from the Datastream country or industry indexes. The mean and standard deviation of the EMU index are reported in Table 2. The test results in Table 3, Panels E and F, suggest that either strategy outperforms the EMU index. That is, all reported test statistics are both negative and significant at the 1% confidence level. Again some of the mimicking portfolios, for example industries over the full sample period, are mean-variance inefficient, e.g. lie below the minimum variance portfolio.

3.3. Geographic versus industry diversification with short-selling constraints

Most portfolio managers face short-selling constraints, but standard mean-variance analysis does not take into account these constraints on the performance of strategies. To ascertain the impact of short-selling constraints, we test constrained country and industry diversification against the opposite strategy and against other benchmarks.

As above we begin comparing the diversification strategies over the full sample by plotting the frontiers, see Fig. 2. Short-selling constraints shrink the industry frontier dramatically. Nevertheless, the tiny industry frontier lies again partially outside of the country frontier. This small part of the industry frontier includes both the tangency and the minimum variance portfolio. That is, the industry efficient frontier spans a smaller range of returns, but "peaks" outside of the country frontier.

The test results for the problem in (3) are presented in Table 4. As already suggested by Fig. 2, over the full sample the return of the country tangency portfolio is not attainable by industry allocation[10], see Panel A. We do not report test statistic for the minimum variance portfolio since the portfolio of the mimicking assets is mean-variance inefficient. Panel B shows the results when industry diversification serves as the benchmark. Test statistics for both the tangency and the minimum variance portfolio are positive and significant at the 10% level and 5% level, respectively. Notice that this is not a contradictory result. The country indexes generate a significantly broader frontier than the industry indexes. That is, the industry tangency and minimum variance portfolios show almost identical mean returns, 8.42% and 8.01%, while the country mean return for the tangency portfolio is more than 100% higher than the return of the country minimum variance portfolio, that is, 10.45% versus 4.89%. In other words introducing short-selling constraints shrinks the efficient locus of industry diversification, while country diversification is affected to a much lesser extent. Nevertheless, a small part of the industry frontier lies outside of the country frontier and significantly so. The last point suggests that the BJS (2002) test is not lacking power to reject the null hypothesis.

Gerard et al. (2002) find that under short sales restrictions geographic diversification dominates industry diversification. Again their findings are consistent with our results.

The results for the subperiods are mixed and change over time. When country diversification is the benchmark, Panel A, we find for the *pre-convergence* subperiod the same results as in the full sample: mimicking portfolios cannot attain the tangency portfolio and are mean-variance inefficient

[10] In the tables we mark tests with an (a) when the mimicking portfolio cannot attain the benchmark. Note that (a) is a stronger result than rejecting the null hypothesis.

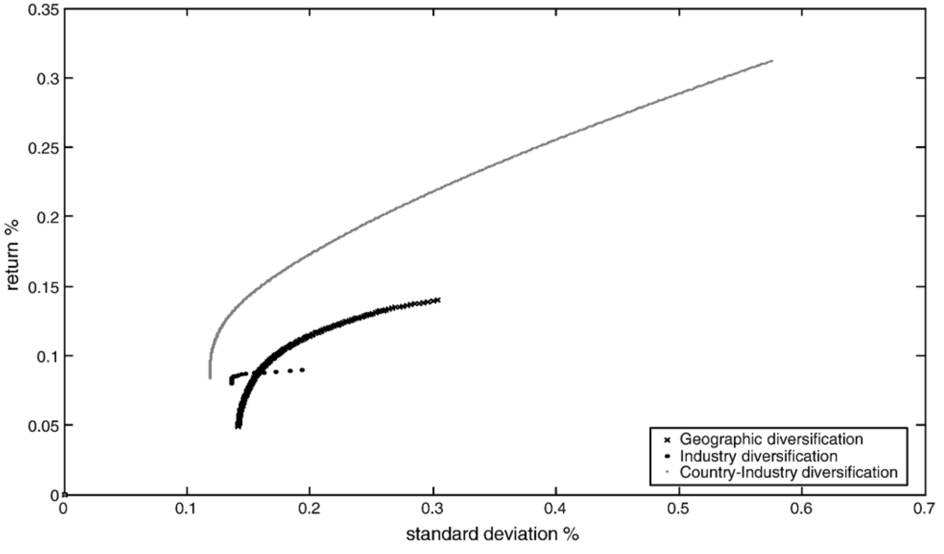

Fig. 2. No short sales mean variance frontiers. This figure displays the no short sales mean variance frontiers of geographic, industry and cross-country and cross-industry (country–industry) diversification. Industry indexes are computed by building market value-weighted indexes based on Datastream level three sector classification. Returns and standard deviations are annualized. The sample period is 1990–2005. Source: Datastream and own calculations.

at the country minimum variance portfolio. In the *convergence* subperiod the test statistics for the tangency portfolio is negative but insignificant and again at the minimum variance portfolio the mimicking portfolio is mean-variance inefficient. As for unconstrained portfolios the *euro* subperiod marks a change in the location of the frontiers and this is reflected in the positive sign of both test statistics. However, we are unable to reject the null hypothesis of equal variance for the tests. Panel B shows exactly the opposite pattern: in the *pre-convergence* subperiod both test statistics are positive and switch sign to negative in the *convergence* subperiod. Again we are unable to reject the null hypothesis of equal variance for each of the tests. In the *euro* subperiod country indexes cannot attain the industry tangency portfolio and are inefficient at the industry minimum variance portfolio. Also, the country efficient frontier is more affected by the short-selling constraint and, hence, spans a smaller range of returns than the industry frontier. Nevertheless, the country frontier lies entirely outside of the industry frontier.

Panels C and D of Table 4 report test results when country–industry diversification represents the benchmark. Neither country nor industry portfolios attain, with the exception of country portfolios in the full sample, the tangency portfolio of the country–industry diversification, confirming the potential advantages of the country–industry structure even with short sales restrictions. For the minimum variance portfolio, the test statistics are always positive and statistically significant, at least at the 10% confidence level, in five of six subperiods and for the full sample when countries form the mimicking portfolios. The most striking result here is that in the euro subperiod country portfolios cannot even attain the minimum variance portfolio of country–industry diversification. We note in passing that the many significant results with the country indexes suggest that the subperiods do not contain too few observations to establish sensible test results. As expected, the passive EMU benchmark, in Panels E and F, is inefficient relative to geographic and industry diversification. That is, the test statistics are both negative and highly significant, at the 1% confidence level, or mimicking portfolios are inefficient.

In brief, our results with short-selling constraints are over time similar to those of our unconstrained portfolios, that is we observe in the *euro* subperiod that the ranking between country and industry is reversed. However, we also observe that industry portfolios are more affected by the no short sales constraint in the full sample and the *pre-convergence* and *convergence* subperiods while country portfolios are more affected by the constraint in the euro subperiod[11].

Over the 1990–2005 period and also in the first two subperiods, industry diversification proves to be at least statistically equivalent to geographic diversification. However, when short-selling constraints are introduced, geographic diversification dominates industry diversification. Both these findings are surprising. The first result is puzzling because of the low correlation structure of countries relative to industries. The second puzzle builds on the first. If these two approaches to international diversification are so similar, then, it is not clear why one diversification strategy should suffer so much more from short-selling constraints than the other.

The same picture arises in the *euro* subperiod. Now, however, with reversed roles. The country frontier lies outside of the industry frontier under the unconstrained optimization. With short-selling constrains the country frontier shrinks dramatically compared to the industry frontier. Therefore, what is the intuition behind these results?

It is well-known that the Markowitz algorithm (1991) leads to extreme portfolio positions (Green and Hollifield, 1992) and that the outcome is almost never balanced. That is, optimal portfolio allocation does not imply that the resulting portfolio will be well diversified[12].

Green and Hollifield (1992) argue that factor risk is the source of unbalanced portfolio weights. If stock or index returns are driven by one dominant factor we can form portfolios with zero factor risk. Such a portfolio will take a large negative position in one stock, or an index, to finance an even larger positive position in other stocks, or indexes. Further, Green and Hollifield (1992) provide empirical evidence that if factor risk stemming from the single dominant factor in stock returns is the reason for unbalanced portfolios, then stock returns must be highly correlated and show a great diversity of betas.

Therefore, if industry portfolios with a higher correlation structure than country portfolios also exhibit higher diversity of betas than country portfolios, then, industry portfolios are better suited to eliminate factor risk. Thus industry portfolios must perform better than country portfolios. To ascertain the validity of this hypothesis, we use the simplest one-factor model possible, the capital asset pricing model, with the value-weighted EMU index as a proxy for the market portfolio. Indeed, in untabulated results we find that industry indexes have a much greater diversity of betas than country indexes. According to these arguments, one should consequently expect the unconstrained industry frontier to lie outside the country frontier.

However, when short positions are not permitted factor risk cannot be eliminated. It is only in these circumstances that the lower correlation structure of countries plays a role. This implies that factor risk stemming from the single dominant factor in stock returns has a much stronger influence on portfolio formation than correlations among particular stocks or indexes.

[11] Avoiding investments in a large fraction of available assets as suggested by portfolio weights obtained with short-selling constraints, e.g. many zero weights, may not be feasible. As pointed out above the BJS (2002) test can deal with lower (l) and upper bounds (u) on portfolio weights, i.e. $l \leq w \leq u$. Hence, so we recompute the portfolio optimization problems adding various lower and upper bounds on the portfolio weights. Untabulated results are easy to summarize: portfolios have very similar return and risk and, hence, we are unable to reject the hypothesis of equality in variance. However, the country efficient frontier dominates, in general, the industry frontier.

[12] A typical reaction is to enforce a balanced portfolio strategy; see Black and Litterman (1992), Jagannathan and Ma (2003) and Michaud (2001). However, country and industry indexes are composed of the same assets. A priori there is no reason why industry indexes would be more subject to measurement error than country indexes.

The *euro* subperiod introduces a major shift in the factor loadings of the country and industry portfolios. Country indexes still show a slightly smaller diversity of betas compared to industry, but industry beta diversity decreases and country beta diversity increases. As discussed above these movements in factor loadings are reflected in our tests for the *euro* subperiod and, thus, are consistent with the Green and Hollifield (1992) argument.

4. Robustness checks

4.1. Bootstrapping portfolio weights

The extreme sensitivity of portfolio weights to changes in the means (see Best and Grauer, 1991) is a major obstacle in mean-variance analysis. To evaluate the importance of estimation error a sequence of bootstrap experiments is carried out. The analysis, aims to illustrate the distribution of the portfolio weights with and without constraints. The distribution of portfolio weights is important from a pure practical point of view; in addition, it is of interest for our purpose to ascertain whether the diversification strategies present different statistical properties for the weights.

Britten-Jones (1999) examines the distribution of unconstrained portfolio weights under independently multivariate Normal distribution. We extend the analysis to portfolio weights with constraints, using a numerical study, without relying on the multivariate Normal distribution. Our block-bootstrap analysis is based on a rolling sampling window method which is helpful in accounting for patterns of long-term dependence in the data.

Table 5 reports the sample mean, the bootstrap mean, the bootstrap standard deviation, and the first quartile of the bootstrap efficient portfolio weights for minimum variance portfolios. Panel A contains the unconstrained and the short-selling constrained country portfolio weights. First, we observe that most of the unconstrained country portfolio weights from the data sample are similar to the mean of the bootstrap data. As in Britten-Jones (1999) we find that efficient portfolio weights have large sampling error. In our case, however, bootstrap confidence intervals are so wide that they are completely uninformative. For example, the sample portfolio weight for the Netherlands is 0.2713, the bootstrap mean portfolio weight is 0.1955, and the bootstrap standard deviation of the portfolio weight is 0.2735.

Because some of the bootstrap portfolio weight distributions are clearly non-normal, standard confidence bounds can be misleading. This fact is evidenced by the large values for the first quartile of the bootstrap distributions. When constructing confidence intervals without relying on the two standard error rule, some confidence bounds are tighter but still uninformative.

Panel A of Table 5 also reports results for constrained country portfolios. Short-selling restricted portfolio weights from the block-bootstrap show mean values very similar to the data. Notice that only five first quartiles, for Austria, Belgium, France, Ireland and Portugal, present values different from zero. Most of the short-selling restricted portfolio weight confidence bounds are significantly tighter than in the unrestricted case but remain large and uninformative.

Panel B of Table 5 contains our results for industry portfolios. In short, the patterns described above for country index portfolios also apply for industry portfolio weights. Therefore, we find it difficult to discriminate among the distributions of country and industry portfolio weights.

It is important to point out that while results in Britten-Jones (1999) and Li et al. (2003) find huge variances in unconstrained tangency portfolios, to our understanding, none of these papers goes so far to conclude that confidence bounds are so huge that no useful information can be deduced from them.

Table 5
Bootstrap distribution of minimum variance portfolio weights

	Sample mean	Bootstrap mean	Bootstrap stdv.	Bootstrap 1st quartile	Sample mean	Bootstrap mean	Bootstrap stdv.	Bootstrap 1st quartile
	Unconstrained portfolio weights				No short sales portfolio weights			
Panel A: country index portfolios								
Austria	0.2823	0.3703	0.1830	0.2497	0.2273	0.3498	0.1724	0.2278
Belgium	0.1960	0.1601	0.0840	0.0963	0.1828	0.1248	0.0856	0.0692
Finland	0.0092	−0.0143	0.0526	−0.0501	0.0000	0.0049	0.0087	0.0000
France	0.1529	0.2008	0.1196	0.0742	0.0345	0.0724	0.0566	0.0229
Germany	−0.3231	−0.2154	0.1441	−0.3593	0.0000	0.0003	0.0023	0.0000
Greece	−0.0114	0.0021	0.0384	−0.0344	0.0000	0.0143	0.0237	0.0000
Ireland	0.1828	0.1500	0.1068	0.0765	0.1642	0.1200	0.0686	0.0705
Italy	0.0522	0.0869	0.0284	0.0638	0.0189	0.0403	0.0351	0.0000
Netherlands	0.2713	0.1955	0.2735	−0.0055	0.1628	0.1609	0.2551	0.0000
Portugal	0.2306	0.1171	0.0445	0.0848	0.2095	0.1007	0.0633	0.0519
Spain	−0.0429	−0.0530	0.1127	−0.1785	0.0000	0.0116	0.0230	0.0000
Panel B: Industry index portfolios								
Basic industries	0.0861	0.1395	0.1789	0.0176	0.0000	0.0528	0.0743	0.0000
Cyclical consumer goods	−0.0725	−0.0662	0.1103	−0.1811	0.0000	0.0000	0.0000	0.0000
Cyclical services	0.2551	0.2905	0.1755	0.1356	0.0912	0.1632	0.1842	0.0000
Financials	−0.2315	−0.0511	0.3058	−0.3804	0.0000	0.0405	0.0716	0.0000
General industries	−0.0494	−0.0637	0.1148	−0.1655	0.0000	0.0001	0.0007	0.0000
Information technology	−0.0361	−0.0586	0.0662	−0.0932	0.0000	0.0003	0.0012	0.0000
Non-cyclical consumer goods	0.4085	0.3471	0.1304	0.2353	0.3252	0.2139	0.1184	0.1229
Non-cyclical services	0.0711	−0.0612	0.1416	−0.1650	0.0205	0.0306	0.0458	0.0000
Resources	0.1669	0.1609	0.0948	0.0941	0.1426	0.1345	0.0980	0.0717
Utilities	0.4017	0.3628	0.0884	0.3081	0.4205	0.3641	0.1031	0.2994

The table presents the sample mean, bootstrap mean, bootstrap standard deviation (stdv.), and bootstrap 1st quartile of minimum variance portfolio weights. Panel A contains unconstrained and no short sales country portfolio weights and Panel B contains unconstrained and no short sales industry portfolio weights. Bootstrap replications involve a block wise sampling across the country or industry return time series. Sample windows are constructed as follows: the size of the sampling window is 256 observations. The first data window ranges from the first observation to the 256th data point. The second window ranges from the second observation to the 257th data point. We proceed until the last window with 256 observations is reached. The bootstrap distribution contains 556 observations. Source: Datastream and own calculations.

On the other hand, standard deviations of tangency portfolio weights in Eun and Resnick (1988) are, at least, as large as the mean of the corresponding portfolio weight and in some cases an order of magnitude larger[13]. Furthermore, Best and Grauer (1991) show that small changes in one of the expected returns of a portfolio lead to extreme changes in many portfolio weights. They report that several portfolio weights change by several hundred multiples of the change in the expected return.

In earlier versions of the paper we implemented a simpler bootstrap procedure that preserved the covariance structure of index returns but failed to account for long-term dependence in the data. Unlike the results above, we found confidence bounds for minimum variance portfolios to be somewhat smaller while tangency portfolios showed larger variations.

4.2. Extended data sample

In this subsection we discuss several additional robustness checks. Our two main concerns are the sensitivity of the test to the number of assets (indexes) and to the sample size. All our untabulated results discussed in this subsection are available upon request.

First, as we test eleven countries versus ten industries, the larger number of countries might give an advantage to geographic diversification. We eliminate Portugal and, thus we can also extend the sample back until April 1988. We do not find that our results depend on the inclusion of Portugal.

Our second concern addresses the small sample properties of our test. Increasing the number of observations may help us to use the standard Normal distribution for the BJS (2002) test with more confidence. Increasing the sample size, however, implies reducing the number of countries in the analysis since the Datastream country indexes for Finland, Greece, and Spain were introduced only shortly before the country index for Portugal. Nevertheless, we collect data back until 1974 for the remaining seven countries and our ten industry indexes. Again the data is weekly US dollar-denominated returns from January 7, 1974 to July 25, 2005 (representing 1647 observations). To control for any biases arising from the smaller number of countries compared to industries we also collect data for Denmark, Switzerland, and the UK. As for the EMU motivated sample, we create three subperiods: *pre-liberalization* ranging from 1974 to 1983, *liberalization* ranging from 1984 to 1993, and *convergence* ranging from 1994 to 2005.

We find that our results are robust to both a longer time series as well as a different mix of countries. The BJS (2002) tests are consistent over our primary 1990–2005 sample with the longer period 1974–2005. We also find that the reduced number of country indexes does not alter our results above. However, we observe that country indexes do relatively better if we include Denmark, Switzerland, and the UK, and this is in particular the case for portfolios with no short sales constraints. Further, there are no apparent patterns in the subperiods for the longer sample that would support the argument that our results in the subperiods for the 1990–2005 sample are due to economic conditions such as the business cycle. Importantly, we cannot detect a change in the factor loadings for the country and industry indexes in the last subperiod for the extended sample size simply because the subperiods are too long.

[13] Our confidence bounds for tangency portfolio weights are even larger than for minimum variance portfolios. Austria for example shows a mean tangency portfolio weight in the bootstrap experiments for unconstrained strategies of 0.21 and a standard deviation of 19.04.

Table 6
Out-of-sample performance of minimum variance portfolios

	Country	Industry	p-value	Country	Industry	p-value
	Unconstrained portfolios			No short sales portfolios		
Mean return	9.59%	8.99%	45.68%	9.60%	9.48%	49.17%
Standard deviation	0.23%	0.23%		0.25%	0.25%	

The table presents annualized country mean and standard deviation, annualized industry mean and standard deviation, and the p-value of a t-test for difference in means for unconstrained minimum variance portfolios and minimum variance portfolios with short-selling constraints, respectively. The t-tests assume equal variances unless the null hypothesis of equal variance (F-test) is rejected at a 10% significance level. The reported p-values assume one-tail distributions. The p-values for two-tail distributions are twice as large as the p-values of a one-tail distribution. Each set of portfolio weights is based on the sample windows from a block-bootstrap return time series. Sample windows are constructed as follows: the size of the sampling window is 256 observations. The first data window ranges from the first observation to the 256th data point. The second window ranges from the second observation to the 257th data point. We proceed until the last window with 256 observations is reached. The bootstrap distribution contains 556 observations. Source: Datastream and own calculations.

4.3. Out-of-sample analysis

Until now we essentially assume that the average returns and the estimated covariance structure are good proxies for the true underlying data generating process. Even if this assumption is indeed satisfied, the out-of-sample performance of mean-variance portfolios may be very different. Of course, since we cannot know the true underlying data generating process the task of constructing efficient portfolios is even harder. Therefore, estimation is key to successful investments and, thus, investors should be interested in the out-of-sample properties of geographic and industry diversification.

Studies find that mean-variance efficient portfolios have a poor out-of-sample performance (see Jobson and Korkie, 1980, 1981; Best and Grauer, 1991). In particular, the sample means of returns represents ingredients to the optimization problem. We are, of course, more interested in out-of-sample differences between geographic diversification and industry diversification than in their absolute performance[14].

We continue to focus on the minimum variance portfolio and analyze the performance of geographic and industry diversification with and without constraints. The out-of-sample experiments take the following form: at the end of each week, we estimate the covariance structure of the country and industry indexes, respectively.

We examine here the out-of-sample performance of the minimum variance portfolios based on estimates of the covariance structure from the block-bootstrap experiments only. In earlier versions of the paper we implemented an exhaustive analysis of out-of-sample portfolios. These results as well as our results for the tangency portfolios are available upon request. Overall, we find that the examples presented below are representative.

For our purpose, the resulting country and industry out-of-sample time series are tested for differences in their means. We assume equal variances for the t-tests unless the null hypothesis of

[14] It is important to remark that we do not examine the out-of-sample performance of geographic diversification and industry diversification based on the predictive power of the usual suspects like the T-bill rate (Breen et al., 1989) or past index returns (Ferson and Harvey (1993)). However, the reader should keep in mind, that when the mean returns are estimated with rolling windows, predictability may affect the results indirectly if one of the diversification strategies exhibits stronger persistence.

equal variance (F-test) is rejected at a 10% significance level. Table 6 reports summary statistics (mean, variance, and the p-value for the t-test) for the out-of-sample experiments. The mean and variance of the time series are annualized.

Three major lessons can be drawn from Table 6. First, it does not appear to be the case that nonnegativity constraints affect the return and risk profile of the analyzed portfolio strategies in a dramatic way. For instance, both the means of returns as well as their standard deviations remain approximately at the same level independently of whether the portfolio strategies involve constraints or not.

Second, untabulated results suggest that the estimation method used for the covariance structure of the country and industry indexes may have a substantial influence on the outcomes. In particular, for unconstrained strategies, rolling windows and increasing estimation windows outperform a portfolio strategy with fixed weights, e.g. known covariance structure. Apparently, an investor could attain a twice as high return with unconstrained country diversification if he consistently updates his estimate instead of focusing on a constant covariance structure.

Third, country diversification almost always outperforms industry diversification in absolute return terms. However, we can never reject the null hypothesis that the means of the two experimental out-of-sample time series are identical. The smallest p-value we observe is as high as 30%.

Overall, our results match well with conclusions in Eun and Resnick (1988). Controlling estimation risk is almost always beneficial. The minimum variance portfolio performs nearly as well as the tangency portfolio when estimation risk is taken into account. But, in general, statistical discrimination among strategies is difficult.

Finally, notice that while we find evidence for the existence of factor risk driving extreme weights in unconstrained optimization (see Green and Hollifield, 1992) we also find evidence of estimation risk. Our results support the view that constraints are helpful because estimation risk has a much larger impact on the performance of portfolio strategies than factor risk.

5. Conclusions

We have compared country and industry diversification based on standard mean-variance theory using the Basak et al. (2002) test. Our approach allows us direct comparison of the two diversification strategies from a pure performance perspective without relying on unobservable country or industry factors. Furthermore, the test is flexible enough to incorporate short sale constraints.

According to the mean-variance tests, geographic diversification and industry diversification are statistically equivalent strategies. Geographic and industry portfolios outperform the EMU index but underperform a diversification strategy based on country–industry indexes. When restricting short sales positions, geographic diversification outperforms industry diversification. The industry efficient frontier lies inside the country frontier, spans a smaller range of returns, but “peaks” outside of the country frontier.

The difference in the way geographic diversification and industry diversification are affected by short-selling constraints is intriguing. We investigate the plausibility of an explanation provided by Green and Hollifield (1992): If stocks, or indexes, are highly correlated and exhibit a wide diversity of betas, one can form portfolios with essentially zero factor risk. Such a portfolio, however, will take a large negative position in one index to finance an even larger positive position in another index. As we compare the features of both country and industry data, we find that industry index data follow that pattern perfectly. Industries not only exhibit higher correlations than countries, but also show large negative positions in optimal portfolios and a much greater diversity of betas with respect to the EMU market index.

A striking aspect in the subperiods is that industry diversification is at least as successful as geographic diversification in the full sample and the first two subperiods. Yet in the *euro* subperiod country diversification clearly outperforms industry diversification. Our evidence seems to be at odds with recent research that advocates the raise in industry factors over the last decade and, hence, also recommends industry diversification.

Overall, country diversification performs better when we impose realistic constraints, which supports the traditional wisdom on geographic allocation. However, the statistical evidence for a difference between the strategies is weak.

Acknowledgement

Financial support from the National Center of Competence in Research, "Financial Valuation and Risk Management" (Ehling) and the Fundacao para a Ciencia e Tecnologia (Ramos), is gratefully acknowledged. The National Center of Competence in Research is managed by the Swiss National Science Foundation on behalf of the Swiss federal authorities. We would like to thank Gopal Basak, Jean-Pierre Danthine, Vihang Errunza, Christian Gourieroux, Albert Holly, Raymond Kan, Michael Rockinger, Olivier Scaillet, Tim Simin, René Stulz, Ernst-Ludwig von Thadden, John Wald, Christian Wolff (the editor), an anonymous referee, and seminar participants at the FAME Workshops, the 5th Conference of the Swiss Society for Financial Market Research (2002), the Euro Conference at NYU (2002), the University of Geneva seminar, the ECB workshop on Capital Markets and Financial Integration in Europe (2002), the European Investment Review Conference (2002), the NFA Annual Meeting (2002), the Macro Brown Bag Lunch at the Wharton School (2003), MFA Annual Meeting (2003), CEMAF/ISCTE Lisbon (2003), EFMA Annual Meeting (2003), Global Asset Management and Performance, McGill University (2005), and Journal of Empirical Finance Conference on International Finance (2005) for helpful discussions and comments.

Appendix A. Bootstrapping portfolio weights

We implement a block-bootstrapping procedure in which blocks of country and industry data are used to compute the BJS test statistics. The block-bootstrap data is constructed by the sampling window method of Hall and Jing (1996). For convergence of this method to the asymptotic distribution of the (normalized) sample mean see Hall et al. (1998). We use this method because the sampling window method is, in particular, helpful in accounting for patterns of autocorrelation or other long-term dependence. In contrast, a simple replication of data with replacement (across the time series) may preserve the covariance structure, the skewness as well as the fat tails of the return distributions but disregards long-term dependence.

The bootstrap experiments are carried out as follows: first recall that the data are indexed by $t \in [0, T]$. The sampling window method uses the first $\hat{T}$ observations in T to compute the first bootstrap BJS test statistics, ξ_1. The size of the sampling window, $\hat{T}$, is assumed to be strictly smaller than T, that is, $t=1, \ldots, \hat{T}$ and $\hat{T}<T$. The subscript denotes that the BJS test statistics is from the first data window of size $\hat{T}$. Next we use the second window ranging from $t=2$ to the $(\hat{T}+1)$th observation to compute ξ_2. The last window is $t=T-\hat{T}+1, \ldots, T$.

Hall et al. (1998) argue that the procedure should be implemented with $\hat{T}=c\times T^{1/2}$ where $c=1, \ldots, 9$. We chose $c=9$ to maximize the likelihood that mean returns are positive. We also conducted block-bootstrap experiments with $c=1$, 3, 4, 5, 6 and do not find that our results depend on the choice of c.

BJS also conduct a sampling window simulation to show that the normality assumption is adequate. Based on the squared correlation between the quantiles of their data and the quantiles of the Standard Normal distribution they conclude that the normality assumption is imperfect but sufficiently realistic. We confirm their results by showing that the bootstrap p-values are sufficiently close to the p-values from the Standard Normal distribution and, further, never affect our conclusions.

References

Adjaouté, K., Danthine, J.P., 2002. European financial integration and equity returns: a theory-based assessment. Second European Central Banking Conference Paper.

Arshanapalli, B., Doukas, J., Lang, L., 1997. Common volatility in the industrial structure of global capital markets. Journal of International Money and Finance 2, 189–209.

Baca, S.P., Garbe, B., Weiss, L., 2000. The rise of sector effects in major equity markets. Financial Analysts Journal 56, 34–40.

Bakker, A.F.P., 1996. The Liberalization of Capital Movements in Europe. Kluwer Academic Publishers.

Basak, G., Jagannathan, R., Sun, G., 2002. A direct test for the mean-variance efficiency of a portfolio. Journal of Economic Dynamics and Control 26, 1195–1215.

Beckers, S., Grinold, R., Rudd, A., Stefek, D., 1992. The relative importance of common factors across the European equity markets. Journal of Banking and Finance 16, 75–99.

Best, M.J., Grauer, R.R., 1991. On the sensitivity of mean-variance-efficient portfolios to changes in assets means: some analytical and computational results. Review of Financial Studies 4, 315–342.

Black, F., Litterman, R., 1992. Global portfolio optimization. Financial Analysts Journal 48, 28–43.

Breen, L., Glosten, L.R., Jagannathan, R., 1989. Economic significance of predictable variations in stock index returns. Journal of Finance 44, 1177–1189.

Britten-Jones, M., 1999. The sampling error in estimates of mean-variance efficient portfolio weights. Journal of Finance 54, 655–671.

Brooks, R., Del Negro, M., 2004a. The rise in comovement across national stock markets: market integration or IT bubble? Journal of Empirical Finance 11, 659–680.

Brooks, R., Del Negro, M., 2004b. A latent factor model with global, country, and industry shocks for international stock returns. Federal Reserve Bank of Atlanta Working Paper 2002-23b.

Carrieri, F., Errunza, V., Sarkissian, S., 2004. Industry risk and market integration. Management Science 50, 207–221.

Carrieri, F., Errunza, V., Sarkissian, S., 2005. The dynamics of geographic versus sectoral diversification: is there a link to the real economy? McGill University Working Paper.

Cavaglia, S., Brightman, C., Aked, M., 2000. The increasing importance of industry factors. Financial Analysts Journal 56, 41–56.

Drummen, M., Zimmermann, H., 1992. The structure of European stock returns. Financial Analysts Journal 48, 15–26.

Eiling, E., Gerard, B., de Roon, F., 2005. International diversification in the Euro-zone: the increasing riskiness of industry portfolios. Working Paper.

Errunza, V., Padmanablan, P., 1988. Further evidence on the benefits of portfolio investments in emerging markets. Financial Analysts Journal 44, 76–78.

Eun, C.S., Resnick, B.G., 1988. Exchange rate uncertainty, forward contracts, and international portfolio selection. Journal of Finance 43, 197–215.

Ferson, W., Harvey, C., 1993. The risk and predictability of international equity returns. Review of Financial Studies 6, 527–566.

Galati, G., Tsatsaronis, K., 2001. The impact of the Euro on Europe's financial markets. BIS Working Paper 100.

Gerard, B., Hillion, P., de Roon, F., 2002. International portfolio diversification: industrial structure, country and currency effects revisited. Working Paper.

Green, R.C., Hollifield, B., 1992. When will mean-variance efficient portfolios be well diversified? Journal of Finance 47, 1785–1809.

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

Grinold, R., Rudd, A., Stefek, D., 1989. Global factors: fact or fiction? Journal of Portfolio Management 79–88 (Fall).

Hall, P., Jing, B.Y., 1996. On sample reuse methods for dependent data. Journal of Royal Statistical Society. Series B. 58, 727–737.

Hall, P., Jing, B.Y., Lahiri, S.N., 1998. On the sampling window method for long-range dependent data. Statistica Sinica 8, 1189–1204.

Heckman, L., Narayanan, S., Patel, S., 1998. Country and industry importance in European returns. Salomon Smith Barney Working Paper.

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

Huberman, G., Kandel, S., 1987. Mean-variance spanning. Journal of Finance 42, 873–888.

Isakov, D., Sonney, F., 2004. Are practitioners right? On the relative importance of international portfolios in international stock returns. Swiss Journal of Economics and Statistics 140, 355–379.

Jagannathan, R., Ma, T., 2003. Risk reduction in large portfolios: why imposing the wrong constraints helps. Journal of Finance 57, 1651–1683.

Jobson, J.D., Korkie, B., 1980. Estimation for Markowitz efficient portfolios. Journal of the American Statistical Association 75, 544–554.

Jobson, J.D., Korkie, B., 1981. Putting Markowitz theory to work. Journal of Portfolio Management 7, 70–74.

Kandel, S., McCulloch, R., Stambaugh, R.F., 1995. Bayesian inference and portfolio efficiency. Review of Financial Studies 8, 1–53.

Li, K., Sarkar, A., Wang, Z., 2003. Diversification benefits of emerging markets subject to portfolio constraints. Journal of Empirical Finance 10, 57–80.

Markowitz, H., 1991. Portfolio Selection: Efficient Diversification of Investments. Blackwell.

Michaud, Richard, O., 2001. Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation. Oxford University Press.

Roll, R., 1992. Industrial structure and the comparative behaviour of international stock market indexes. Journal of Finance 47, 3–42.

Rouwenhorst, K.G., 1999. European equity markets and the EMU. Financial Analysts Journal 55, 57–64.

Wang, Z., 1998. Efficiency loss and constraints on portfolio holdings. Journal of Financial Economics 48, 359–375.