Have European Stocks Become More Volatile? An Empirical Investigation of Volatilities and Correlations in EMU Equity Markets at the Firm, Industry and Market Level

Colm Kearney and Valerio Potì*

14th January 2004

Abstract

We examine the dynamics of idiosyncratic risk, market risk and return correlations in European equity markets using weekly observations from 3515 stocks listed in the Euro-area stock markets in the period 1974-2004. Similarly to Campbell, Lettau, Malkiel and Xu (2001), we find an increase in idiosyncratic volatility, implying that it now takes more stocks to diversify away idiosyncratic risk. Contrary to their findings, however, market risk is trended upwards and correlations among the stocks are only mildly trended downwards. Market volatility tends to lead the other volatility measures in EMU markets whereas idiosyncratic volatility leads in the US ones. Both the volatility and the correlation measures are pro-cyclical and they increase at times of low market returns, implying a skewed market portfolio return distribution. We suggest a number of implications of these findings for portfolio management, trading and asset pricing.

Key Words: Correlation dynamics, idiosyncratic risk, asset pricing.
JEL Classification: C32, G11, G12, G12, G15.

Contact details:

Colm Kearney, Professor of International Business, School of Business Studies and Institute for International Integration Studies, Trinity College, Dublin 2, Ireland. Tel: 3531-6082688, Fax: 3531- 6799503. Email: colm.kearney@tcd.ie.
Valerio Potì, Finance Lecturer, Dublin City University and Trinity College Dublin, Dublin City University Business School, Glasnevin, Dublin 9, Ireland. Tel: 3531-7005823, Fax: 3531-6799503. Email: valerio.poti@dcu.ie.

*Contact author.

Previous versions of this paper were presented at the Academy of Accountancy and Finance Conference, New Orleans, December 2002, at the European meeting of the Financial Management Association in Dublin, June 2003 and at the Annual Meeting of the European Finance Association in Glasgow, August 2003. Both authors acknowledge financial support from the *Irish Research Council for the Humanities and the Social Sciences.*

Have European Stocks Become More Volatile? An Empirical Investigation of Volatilities and Correlations in EMU Equity Markets at the Firm, Industry and Market Level

1. Introduction

The awareness that financial asset volatility exhibits time varying behaviour, together with the notion that returns are somewhat predictable and that their mean is also time-varying, is central to the view that returns are not identically and independently distributed (*i.i.d.*) as they were once thought to be (see, for a nice discussion, Cochrane (1999)). A large literature has therefore focused on modelling volatility time series. Within the capital asset pricing model (CAPM), originally formulated by Sharpe (1964) and Lintner (1965), only systematic market volatility matters because idiosyncratic volatility can be diversified away. Most empirical studies have consequently focussed on the former (see for example the pioneering work of Schwert (1989) and Bollersev, Chou and Kroner (1992) and Campbell, Lo and MacKinlay (1997) for surveys).

More recently, however, financial researchers have begun to re-examine the nature of risk in equity markets. Evidence, provided by (among others) Barber and Odean (2000), Benartzi and Thaler (2001) and Falkenstein (1996), show that investors often hold undiversified portfolios. Barberis and Thaler (2003) provide an exhaustive review of contributions on this "insufficient diversification" puzzle. For under-diversified investors, as suggested by Malkiel and Xu (2002), the relevant measures of risk may well be total volatility (that comprises *both* market and idiosyncratic volatility). Moreover, Benartzi and Thaler (2001) find that even investors who diversify usually hold only a limited number of assets in order to reduce transaction costs ("naive diversification"). As the level of average idiosyncratic risk influences the Sharpe ratio of naive or less than complete diversification strategies, their success crucially depends on the level of idiosyncratic volatility of the average stock. Therefore, while systematic, market-wide volatility is important to holders of diversified portfolio, such as pension funds, investment houses and other institutional investors, total and idiosyncratic volatility are relevant for under-diversified investors. The study of the level and dynamics of average idiosyncratic risk is especially important in European equity markets as more and more individuals are investing in stocks, or might soon start to do so. The desire to supplement social security benefits and public pension provisions, shrinking because of a rapidly ageing population, contributes towards this shift in investment habits. See, among many others, Guiso, Haliassos and Jappelli (2002) for an extensive review of the empirical evidence on increasing stock market participation in Europe and the importance of its demographic determinants. As suggested by abundant behavioural evidence (see for a survey Barberis and Thaler (2003)), these investors will likely adopt some form of imperfect diversification strategy that will leave them exposed to a certain amount of average idiosyncratic risk. Total asset volatility, including its idiosyncratic component, is also relevant to the pricing and dynamic hedging of options and other contingent-claims contracts. These arguments suggest that *both* systematic and idiosyncratic volatility should be studied. In this vein, Campbell, Lettau, Malkiel and Xu (2001), henceforth CLMX (2001), analyse long-term trends and short-run dynamics in *both* market and idiosyncratic (industry and firm-level) volatility in United States (US, henceforth) stock markets from 1962 to 1997. They use daily data on all stocks traded throughout the period on three markets (AMEX, NASDAQ and the NYSE).

In this paper, we extend CLMX's (2001) study to European equity market. In particular, we analyse the behaviour over the period 1974-2004 of systematic and aggregate firm and industry-level volatility of the 3515 stocks listed on the markets of the current members of the European Monetary Union[1] (EMU henceforth). We also construct comparable US series from CLMX's (2001) data.

[1] In this study we neglect the country level, traditionally prominent in the literature on volatility in European markets (see. for example, Baele (2002)) and we focus instead on the firm, industry and aggregate level of the EMU stock market as a whole. This choice is motivated by the considerable evidence on a substantial degree of equity market integration, which has gathered pace in Europe since the mid-1990s (Hardouvelis, Malliaropulos and Priestley (2000) and Fratzschler (2002)). Moreover, following the introduction of the Euro, equity markets of the

We analytically show that there exists an exact relationship[2] between systematic variance, average idiosyncratic variance and average correlation and we use it to simplify the computation of the average correlation amongst the large number of stocks in our sample. We study both *long-term trends* in the variance and correlation series and *shorter run relationships* that link these series to each other and to market returns. Our aim is to gain insights in their predictability and on the portfolio management, trading and asset pricing implications of their behaviour over time.

Regarding *long-term* trends, our main findings are *first* that European stocks have indeed become more volatile. In particular, because of the larger idiosyncratic volatility of the typical stock, it takes many more stocks to diversify it away. For example, the number of stocks required for residual idiosyncratic portfolio volatility to be reduced to 5.0 percent increased from 35 in 1974 to a maximum of 166 in the second semester 2003. *Second*, average EMU stock correlation appears to be a fast-moving series that quickly mean-revert to a long-run value of around 20 percent. *Third*, idiosyncratic volatility is by far the largest component of the volatility of the typical stock and the potential benefit of diversification strategies is both substantial and fairly stable in the long run.

Regarding *short-run dynamics*, our main findings are *first* that EMU market volatility forecasts both industry and firm-level volatility, whereas US firm level volatility predicts US market and industry-level volatility. This suggests a role reversal of market and firm-level volatility in the EMU and the US markets, with potential asset pricing implications. *Second*, EMU and US average stock correlations move very closely together over time and, interestingly, the former is more explained by lagged values of the latter than by its own lags, thus stressing the importance of the US market as a driving force of global risk

countries that have adopted the new currency have become almost perfectly correlated, as reported by Cappiello, Engle and Sheppard (2003) and by Kearney and Potì (2003).

[2] The analytical derivation of this relationship, while intuitive, is to our knowledge novel. We already introduced this derivation in a previous working paper, but we report and further discuss it here for the reader convenience.

factors. *Third*, regarding returns predictability, EMU market returns are positively related to lagged market variance but negatively related to lagged idiosyncratic variance. This confirms the result recently found by Guo and Savickas (2003) using comparable US data. We interpret this finding as evidence either that market and average idiosyncratic variance jointly proxy for the predictable portion of market risk or that idiosyncratic variance is a conditioning variable that helps explain variation in risk aversion.

Our paper is structured as follows. We begin by introducing, in Section 2, the CLMX's (2001) decomposition of average stock variance and we analytically derive the relation between market variance, average idiosyncratic variance and average correlation. In Section 3, for consistency with CLMX (2001), we extend our analysis to further decompose idiosyncratic variance into an industry and a firm level component. In Section 4, we describe our data set and construct the variance and correlations series for EMU stock markets. We also construct, from CLMX (2001) data, comparable US variance and correlations series. In Section 5, we perform a range of statistical tests to discern more formally the time series behaviour of variances and average correlation and we study their lead-lag relationships and how they help explain each other and the return on the market portfolio. In Section 6, we discuss the implications of our findings for portfolio management, trading strategies and asset pricing. In the final Section, we summarise our main findings and present our conclusions.

2. Variance Components and Average Correlation

Consider the following empirical models of stock and industry returns:

$$r_{i,j,t} = r_{f,t} + \beta_{i,j,m}(r_{m,t} - r_{f,t}) + u_{i,j,t}$$

$$= r_{f,t} + \beta_{i,j}(r_{j,t} - r_{f,t}) + \beta_{i,j}\beta_{j,m}(r_{m,t} - r_{f,t}) + e_{i,j,t} \qquad (1)$$

Where:

$$r_{j,t} = r_{f,t} + \beta_{j,m}(r_{m,t} - r_{f,t}) + \varepsilon_{j,t}$$

Here, $r_{i,j,t}$ is the return on the firm i that belongs to industry j taken from portfolio m, $r_{f,t}$ is the risk free rate, $r_{m,t}$ is the return on the portfolio m, $r_{j,t}$ is the industry j return, $\beta_{i,j,m}$, $\beta_{i,j}$ and $\beta_{j,m}$ are regression coefficients, $u_{i,j,t}$ is an idiosyncratic residual and $e_{i,j,t}$ and $\varepsilon_{j,t}$ are, respectively, firm and industry-level idiosyncratic regression residuals. Also, denote by k the maximum number of assets in each one of the n industries and by $w_{j,t}$ the weight of industry j in the portfolio m and by $w_{i,j,t}$ the weight of asset i in industry j. Then, define the average variance of the idiosyncratic, firm and industry-level regression residuals, the average industry variance and the portfolio return variance as follows:

$$\begin{aligned}
IDIO_t &= \sum_{j=1}^{n} w_{j,t} \sum_{i=1}^{k} w_{i,j,t} Var(u_{i,j,t}) \\
FIRM_t &= \sum_{j=1}^{n} w_{j,t} \sum_{i=1}^{k} w_{i,j,t} Var(e_{i,j,t}) \\
IND_t &= \sum_{j=1}^{n} w_{j,t} Var(\varepsilon_{j,t}) \\
VAR_t^{ind} &= \sum_{j=1}^{n} w_{j,t} Var(r_{j,t}) \\
MKT_t &= Var(r_{m,t})
\end{aligned} \tag{2}$$

Applying CLMX's (2001) analytical framework, we can decompose the average total variance (VAR_t) of the returns on the assets in the portfolio as follows:

$$\begin{aligned}
VAR_t &= \sum_{i=1}^{n} w_{i,j,t} Var(r_{i,j,t}) = VAR_t^{ind} + FIRM_t \\
&= MKT_t + IDIO_t
\end{aligned} \tag{3}$$

Where:

$$VAR_t^{ind} = MKT_t + IND$$

$$IDIO_t = FIRM_t + IND_t$$

Moreover, if the portfolio of assets is the *market* portfolio, $u_{i,j,t}$ is a CAPM idiosyncratic residual. In this case, (3) provides a CAPM-equivalent decomposition[3] of average total variance into market variance and average idiosyncratic variance, with the considerable advantage that it bypasses the need to estimate *betas* for each asset.

To derive the analytical relationship between systematic variance, average idiosyncratic variance and average correlation, we first rewrite the MKT_t and VAR_t terms in (2) by converting them to matrix notation:

$$MKT_t = w_t' H_t w_t \tag{4}$$

$$VAR_t = w_t' D_t D_t i \tag{5}$$

Where,

$$H_t \equiv D_t C_t D_t$$
$$[H_t]_{i,j} = h_{i,j,t}$$
$$[C_t]_{i,j} = c_{i,j,t} \in [-1, 1] \;\; \forall \;\; i \neq j, \;\; \text{and} \;\; [C_t]_{i,j} = c_{i,j,t} = 1 \qquad \forall \;\; i = j$$

In (4) and (5), w_t is an *nx1* vector of weights, i is an *nx1* unit vector, C_t is an *nxn* correlation matrix, H_t is an *nxn* variance-covariance matrix, D_t is an *nxn* diagonal matrix, with the elements $d_{i,j,t}$ on its main diagonal being the standard deviations of returns. In particular:

$$[D_t]_{ij} = d_{i,j,t} = \sqrt{h_{i,j,t}} \quad \forall \;\; i = j, \;\; \text{and} \;\; [D_t]_{ij} = d_{i,j,t} = 0 \;\; \forall \;\; i \neq j$$

[3] As discussed by CLMX (2001), this is an approximate decomposition. In particular, $IDIO_t$ is only approximately equal to the average variance of the CAPM idiosyncratic residuals. CLMX

We can re-write (4) as:

$$
\begin{aligned}
MKT_t &= w_t' D_t C_t D_t w_t \\
&= \sum_{i=1}^{n} \sum_{j=1}^{n} w_{i,t} w_{j,t} c_{i,j,t} d_{i,j,t} d_{i,j,t} \\
&= w_t' D_t D_t i w_t' C_t w_t \\
&= (w_t' D_t D_t i) CORR_t \qquad (6)
\end{aligned}
$$

Here:

$$CORR_t = \sum_{i=1}^{n} \sum_{j=1}^{n} w_{i,t} w_{j,t} c_{i,j,t} = w_t' C_t w_t$$

Using (5), we can rewrite (6) as follows:

$$MKT_t = VAR_t CORR_t \qquad (7)$$

This is the first appealing analytical result. It says that the task of modelling the variance of a group of assets can be broken up in the simpler tasks of modelling the average variance and correlation processes and, possibly, their interaction. This provides a justification for 2-step estimation procedures of the conditional variance-covariance matrices of large systems such as the DCC-GARCH model used by Cappiello, Engle and Sheppard (2003). Solving (7) for $CORR_t$, we have a suggestive expression for the average correlation:

$$CORR_t = \frac{MKT_t}{VAR_t} \qquad (8)$$

(2001), however, show that their difference is negligible if the cross-sectional variance of the beta coefficients is not too volatile.

Equation (8) provides an intuitively appealing result. Average correlation is the ratio between systematic and average idiosyncratic variance. This result can be used to substantially simplify the construction of average correlation time series amongst a large number of assets. Also, from (3) and (8), we obtain a suggestive expression for average idiosyncratic variance:

$$IDIO_t = (1 - CORR_t)VAR_t \tag{9}$$

Using (8) and (9), we can rewrite the variance decomposition equation in (3) as:

$$VAR_t = \underbrace{CORR_t VAR_t}_{MKT_t} + \underbrace{(1 - CORR_t)VAR_t}_{IDIO_t} \tag{10}$$

Interestingly, (10) tells us that we can interpret average correlation as the parameter that, for any given level of average total variance, splits the latter into systematic variance and average idiosyncratic variance. The variance of the average asset is, in other words, the weighted average (or a convex combination) of systematic and average idiosyncratic variance, with weights given by average correlation and by its complement to one.

4. Data and Variable Construction

We use weekly returns and semi-annual capitalization data from Datastream for the period December 1974 to March 2004. By using weekly returns we overcome the problem of asynchronous trading across the EMU stocks markets. Firm level data comprise total returns and market capitalisation for the 3515 stocks listed on the stock markets of the EMU member countries[4]. Industry level data comprise Datastream Level 4 fixed history industry indices for the EMU equity market[5]. Market level data comprise total returns on the Datastream fixed

[4] We include all the countries that had adopted the Euro as of March 2004, as reported in Panel A of the Data Appendix.

[5] Datastream Level 4 Industry Indices classify EMU stocks into 35 industries (Panel B in the Data Appendix), thus providing enough cross-sectional variation to be able to discriminate their behaviour from sources of variation common to all the stocks (e.g. the market).

history[6] index for the overall EMU stock market[7]. We use unconditional estimators of variances based on sums (or averages) of return innovations squares and cross-products[8]. We therefore define variance over a period T of length p as the average of the squared deviations of returns (or their components) r_t, t = 1,...., p, from their mean $\bar{r}_T$. In all computations we apply the convention that each year comprises 52 weeks and each semester comprises 26 weeks. Therefore, to compute our semi-annual variance of weekly returns, we set p equal to 26. Formally:

$$Var(r_t)_T = \sum_{t=1}^{p} (r_t - \bar{r}_T)^2 \tag{11}$$

EMU Series

Using (11), we first construct variance series computed over non-overlapping semi-annual periods for the individual stocks $Var(r_{i,j,t})_T$, j = 1,.....k, for the individual industries $Var(r_{j,t})_T$, i =1,....n, and for the market portfolio $Var(r_{m,t})_T$. We then compute the average total variance time series and, using (3), we derive the average idiosyncratic variance time series as the difference between VAR_T and MKT_T[9]. Formally, we first compute MKT_T and VAR_T as follows:

$$MKT_T = Var(r_{m,t})_T \tag{12}$$

---

[6] The choice of using fixed history indices is necessary to ensure consistency with our average variance computation methodology and with the procedure followed by CLMX (2001).

[7] We constructed a value-weighted index of all the stocks included in our dataset for the shorter period 1st semester 1997 – 1st semester 2004 and found that its correlation with the Datastream EMU market index was almost perfect (96.8 percent) over this period and over various sub-periods. We felt that, since we could use the excellent proxy represented by the Datastream EMU market index (that represents at least 75% of the capitalization of the EMU equity market), it was not necessary to construct the value-weighted index of our stocks for the entire 1974-2004 sample period, a computationally very intensive task that would have likely lead to errors.

[8] Many researchers have used this approach (see, for example, Schwert (1989) and CLMX (2001)) because of its simplicity. The implicit assumption is that the variance of a process is an observable variable. It follows that, as pointed out by Merton (1980), it can be estimated to any desired degree of accuracy by sampling squared deviations of the process realisations from their means at sufficiently high frequency.

[9] The Datastream Global Index computation methodology updates the stock weights every 3 months until 1995 and, from then on, every January. In the value-weighted calculations we update weights semi-annually, using mid period capitalizations.

$$VAR_T = \sum_{i=1}^{n}\sum_{j=1}^{k} w_{i,T} w_{i,j,T} Var(r_{i,j,t})_T \qquad (13)$$

Using (3), we then compute, from (12) and (13), the average idiosyncratic variance $IDIO_T$:

$$IDIO_T = VAR_T - MKT_T \qquad (14)$$

Turning to the decomposition of average idiosyncratic variance into an industry and a firm level component, we are not able to compute VAR_t^{ind} in (2) from our stock level data. This would entail classifying our 3515 stocks into industries, creating a value-weighted index for each index and computing the variance of each one of these industry indices over time. Because of the format of our database, this is currently impossible[10]. Instead, we use Datastream Level 4 fixed history industry indices for the EMU equity market[11]. Applying (3), we construct IND_T by subtracting MKT_T from $VAR^{ind}{}_T$ and we derive $FIRM_T$ by subtracting IND_T from $IDIO_T$. Applying (8) to the constructed market and average stock and industry variance series, we compute the equally-weighted and value-weighted average correlation among the stocks and the industries.

We obtain 61 non overlapping semi-annual variance and correlation data points (T = 1, 2,....., 61) computed from weekly returns data. The variance series are annualized by multiplication by a factor of 2 (to minimise rounding errors and to display the results in a more legible numerical format). While we construct both equally-weighted and value-weighted series, we focus and, therefore, report mainly on the latter. The constructed series and their definitions are summarised

[10] The main problem is that our source, Datastream Advance, downloads data into MS Excel spreadsheets. The latter suffer the limitations of being able to contain only 256 columns. Therefore, we are forced to download the data by rows (one stock per row) because their number is much larger (65,536) and to display the return and variance time series by columns. This means that we can download only roughly a year of weekly data at a time, making it impossible to automate the procedure to the extent necessary to reliably perform the industry classification and the ensuing variance and weighting computations. It would help considerably if Datastream, or another provider, made the data available for download in text files (e.g. files with extension .txt, which can contain a much larger number of rows and columns).

[11] They represent at least 75% of the capitalization of the relevant industry.

in Panel C of the Data Appendix. In Figure 1, we plot the decomposition of annualised value-weighted average total stock variance[12] into market and average idiosyncratic-level variance (Panel A) and the ratio of firm to industry variance (Panel C). In Figure 2 we report the equally-weighted (Panel A) and value-weighted (Panel B) average stock and industry correlations.

US Series

To be able to accurately examine to what extent our constructed EMU variance and correlation series share similar features to those displayed by the series used by CLMX (2001) for US markets, we construct comparable variance and correlation series from CLMX (2001) data[13]. To do so, we simply aggregate at a semi-annual frequency the monthly CLMX (2001) market, industry and firm-level variance series constructed from weekly returns and we then multiply the results by a factor of two to annualise[14]. The average idiosyncratic variance series is computed, according to (3), as the sum of the average industry and firm-level variance series. Using (8), we also derive the average stock correlation series, not available in the CLMX (2001) study, from the ratio of the market variance to the average total variance series. We obtain 48 non overlapping semi-annual value-weighted variance and correlation data points (T = 1, 2,....., 48) computed from weekly returns data. The US variance series are plotted in Figure 1 (Panel B). The US average correlation series is reported in Figure 2 (Panel B).

Graphical Analysis

EMU total, idiosyncratic and market variance series start off relatively low and tend to rise towards the end of the period. This tendency, however, is more pronounced for idiosyncratic variance and its firm-level component. US total and

[12] The equally-weighted average total variance series (not reported but available upon request) is much higher, thus suggesting that the greater the capitalization, the smaller, on average, stock volatility. However, since the equally-weighted market variance is smaller than the value-weighted one, small-capitalization stocks are on average less correlated than large-capitalization ones.

[13] We thank CLMX (2001) for kindly making their constructed variance series available.

idiosyncratic variance series also tend to increase over the portion of CLMX's (2001) sample period that overlaps with ours (1974-1997). EMU average stock correlation closely mirrors the level and behaviour of the US one. Therefore, from (9), the relative portion of the total variance of the average EMU and US stock represented by idiosyncratic variance is similar throughout most of the sample period. In particular, since average stock correlation is usually well below 50% (with the noticeable exception of the oil 1974 crisis and the 1987 stock market crash), idiosyncratic variance is the largest component of both EMU and US average total variance and, therefore, the potential benefit to diversification strategies is substantial in both cases. Average EMU stock correlation is dramatically lower in the equally-weighted case than in the value-weighted one. This implies that most of the variance of the average EMU stock can in principle be diversified away.

As reported in Panel C of Figure 1, the EMU ratio of firm level to industry volatility increased much less than the US one. A closer inspection of the data reveals that the reason why this ratio is initially so low for EMU markets is the little cross-sectional dispersion within industries due to the limited number of listed stocks. Essentially, unlike in the much more mature US markets, EMU industry indices used to comprise only a few stocks (usually with relatively similar operations and, therefore, homogeneous fundamental sources of returns variability). This also explains why, as reported in Panel B of Figure 2, the average correlation amongst EMU industries is remarkably similar to the average correlation amongst EMU stocks until the mid 80s. As shown in Figure 3, this started to radically change from the mid 80s onwards. The EMU ratio of firm to industry variance first rapidly catches up with the US one and in the 90s they move closely together. At the same time, average EMU stock correlations considerably decreased relative to average EMU industry correlations.

5. Time Series Behaviour

[14] Notice, however, that MKT_T constructed from CLMX (2001) data is the variance of the market portfolio excess-return.

We begin our formal study of the time series behaviour of our constructed variance and correlation series by reporting descriptive correlations and autocorrelations. We will then conduct *Wald* tests on the presence of a time-trend. Finally, we will examine the short run interactions amongst the series, with special regard to whether and to what extent they help forecast each other and market returns and to whether their relationship with the latter helps explain the systematic skewness of asset returns.

Descriptive Correlations

In Table 1 we report descriptive correlations of the EMU and US market variance, average idiosyncratic variance and average correlation series with their own and each other's lags. All the series, with the exception of average idiosyncratic risk and its industry and firm-level components, display very little persistence (but this slightly increases in the 1974-2004 sample) and the US series are somewhat more persistent than the EMU ones[15]. It is worth noticing that the low persistence is an unusual result. It is explained by the circumstance that our variance and correlation series are constructed using weekly instead of daily returns and that the semi-annual sampling period is relatively long compared to the usual monthly and daily frequencies. The low persistence of the variance and correlation series implies that these are unlikely to contain a unit root. This is also the case for the more persistent average idiosyncratic variance and its industry and firm-level components[16]. We therefore treat the constructed variance and correlation series as stationary and work with them in levels without differencing[17].

Wald-Type Tests

---

[15] Neither the results for the 1974-2004 sample nor those for the average industry and firm level variance series are reported to save space but they are available upon request.
[16] While they are more auto-correlated, they appear far from containing a unit root.
[17] To double check on whether the series are stationary, however, we also conduct Dickey-Fuller and augmented Dickey-Fuller tests and we analyse the spectral density function of the series. These results are available upon request.

We perform Wald-type tests on the presence of a deterministic time-trend in our constructed EMU and US value-weighted variance and correlation time series. Since the residuals from a static model that includes only a constant and a deterministic time-trend are auto-correlated, we estimate a dynamic model that includes also a lag of the dependent variable[18]. We then test the restriction that the deterministic time trend coefficient in the dynamic model is zero.

The results are reported in Table 2. Panel A reports the results for the EMU series. Average idiosyncratic variance (largely because of its firm-level component) and EMU market variance contain a statistically significant deterministic trend. The coefficient estimated for the former (0.10 percent) is greater then the coefficient estimated for the latter (0.056 percent) and they explain a substantial portion of the increase of these series over time. For example, after 10 years the projected increase in MKT_T is 0.56 percent whereas the increase in $IDIO_T$ is 1.0 percent (corresponding, respectively, to almost 7.5 percent and 10 percent in volatility terms). The results for the US series are reported in Panel B. US average correlation contains a marginally significant negative time-trend. All other trend coefficient estimates are insignificant. The signs of the estimated deterministic time trend coefficients, however, agree with the CLMX's (2001) study.

The low own lag regression coefficient estimate, for both the EMU and US average stock correlation series, means that they are fast moving variables that quickly revert back to their long run average (the half-life of a shock to the EMU and US series is, respectively, 2.19 and 3.41 months). The long run average of the EMU and US series are, respectively, stationary (20.5 percent) and possibly decreasing over time due to a marginally significant deterministic trend (at a rate of roughly 4 percent every 10 years from an initial 24.58 percent).

---

[18] We include among the regressors only one lag of the regressand because, from Table 1, higher order auto-correlations do not appear to be important. To check that the estimated residuals from this model are serially independent we use the Durbin's *h* statistic because, in the presence of a lagged value of the dependent variable among the regressors, the *DW* test is biased towards acceptance of the null of no autocorrelation. We use the generalised version of Durbin's *h*-test, developed by Godfrey and Breusch, based on a general Lagrange Multiplier test. Even though

VAR Analysis

So far we have analyzed the long-run trends (that unfolded over a 30 year period) of our constructed variance and correlation series. However, a casual examination of their descriptive correlations, reported in Table 1, suggests that there is also a rich set of dynamic interactions of a shorter-term nature (at a zero to two semi-annual lags horizon). We will now more formally study these relationships. In particular, we will examine to what extent our constructed variance components help predict each other and the market return within a VAR system. This exercise is particularly important because, since these variables are little persistent at our chosen data frequency, they would be difficult to forecast to any interesting level of accuracy using only their own lags. It also offers preliminary indications on whether the variance series help predict market returns and which variables, in a trivariate relation between MKT_T, $IDIO_T$ and the market return, can be considered (weakly) exogenous.

We set up a VAR system of our MKT_T, IND_T and $FIRM_T$ to study Granger-causality relationships among them. Both the AIC and the SBC (not reported to save space) suggest the inclusion of only one lag of each variable in the VAR system[19]. In the analysis that follows, therefore, we include only one lag of each variable in the VAR system. We first perform block-exogeneity tests to determine whether lags of one variable Granger-cause any of the other variables in the system. If all the lags of one variable can be excluded from the equations of the other two variables, we can model these two variables using a simple 2-variable VAR. To perform the test, we estimate unrestricted 2-variable VAR systems with one lag of each of the three variables that enter as a predetermined variable the equations for the other two and restricted versions of the same systems without the lags of the predetermined variable. To test these restrictions we use the Likelihood Ratio test. The constructed statistics, modified with the

this procedure can detect higher order serial correlation, we test only test the null of no first-order residual autocorrelation.

inclusion of Sims's (1980) multiplier correction to improve the small sample properties of the test, is distributed as a Chi-Squared with degrees of freedom equal to the number of lags (1 in each equation, for a total of 2 restrictions) excluded from each equation in the restricted system.

Table 3 reports the results. For EMU series, from Panel A, the only non block exogenous variable is MKT_T[20]. Moreover, this variable is the only one that directly Granger-causes another one[21], namely IND_T. Table 3 also reports, in Panel B, the variance decomposition of the EMU market variance and average industry and firm level variance series. The variance decomposition imposes the restriction that there are no contemporaneous effects of IND_T on MKT_T and of $FIRM_T$ on both IND_T and MKT_T. This ordering is suggested by the results of the block-exogeneity tests and it is consistent with an industry factor structure, augmented by a market factor, of both asset returns and asset returns second moments. This innovation accounting suggests that a large portion of the variance of IND_T and $FIRM_T$, over 30% and 45% respectively one period ahead, is explained by variation in MKT_T, whereas very little of the latter in explained by variation in IND_T and $FIRM_T$. Regarding causality links between the US series, the relevant results are also reported in Table 3[22]. We find that there is a clear role reversal between market and firm-level variance. The only non block exogenous variable is average firm level variance. This series helps predict the market and industry variance series whereas the latter have very little predictive power[23]. The overall picture that emerges is that market variance plays a much more important role in the EMU whereas the latter is more relevant, in terms of predictive ability, in the US. This suggests that their role in predicting stock market returns and asset pricing, which might be relevant as suggested by

---

[19] Moreover, for EMU series, a Likelihood Ratio test does not reject the restriction that the lag length is one instead of two (the Chi-squared statistics is 7.83 with significance level 0.550).
[20] The significance level of MKT_T is only slightly higher than the 5 percent level.
[21] This is not reported in the table to save space.
[22] These results are different from those reported by CLMX (2001) for the US equity markets as they find that their variance series predict each other to a greater extent. It should be kept in mind that CLMX's (2001) variance series are constructed at a monthly frequency from daily returns, whereas we use series and underlying returns with lower frequency, semi-annual and weekly respectively. This explains to a large degree the weaker Granger-causality links between our variance series and their lower persistence.
[23] Details on these results are not reported in the table but they are available upon request.

(among many others) Goyal and Santa Clara (2003) and Guo and Savickas (2003), might be different depending on whether the stocks are drawn from a EMU or a US sample.

After studying the multivariate predictive relations between the market and the idiosyncratic variance components, we focus on the multivariate predictive relations between the variance components and the market return. In particular, we set up a VAR system of our MKT_T, $IDIO_T$ and R_{mT} series to study Granger-causality relationships among them. Again, both the AIC and the SBC (not reported to save space) suggest to include only one lag of the variables in the VAR system. In Figure 5 we report the impulse response functions and their confidence intervals to visualize the impact of shocks to MKT_T and $IDIO_T$ on MKT_T, $IDIO_T$ and the market return. Confidence intervals are constructed using a Monte Carlo integration method[24]. Since a reduced form VAR is under-identified, we impose a set of identifying restrictions in a manner that is consistent with asset pricing theory and with the results of our block-exogeneity and Granger-causality tests. In particular, we use a Cholesky decomposition of the VAR variance-covariance matrix that rules out contemporaneous effects of $IDIO_T$ on MKT_T and of the market return on both variance series. With this "ordering" of the variables ($MKT_T \rightarrow IDIO_T \rightarrow R_{mT}$), MKT_T explains most of its own variance and a large portion of the variance of $IDIO_T$ and R_{mT} whereas $IDIO_T$ and R_{mT} explain only most of their own variance (this is not reported to save space). The impulse response functions highlight the large contemporaneous positive effect that a shock to MKT_T has on $IDIO_T$ and the even larger but negative effect on R_{mT}. However, while the effect on $IDIO_T$ quickly fades away, the lagged effect on the market return is substantial and of opposite sign (positive). Therefore, higher market variance causes first a contemporaneous negative return and then, the next period, a positive return. This can be interpreted as suggesting that higher market variance causes higher market expected return and, hence, an initial drop in price followed by a higher return. This analysis confirms previous evidence provided by Harvey (1989) and

[24] See, for example, Enders (2004) for details on how to construct confidence bands for impulse response functions using a Monte Carlo integration procedure.

by Turner, Startz and Nelson (1989) on a positive relation between market risk and return. Under the assumption that lagged market variance proxies for expected market variance, our results are also consistent with those reported by French, Schwert and Stambaugh (1987) and, more recently, Scruggs (1998). They found that, while the expected component of the stock market excess return is positively related to the predictable stock market volatility, actual volatility and actual returns are negatively correlated. A positive shock to $IDIO_T$, instead, has no initial impact on MKT_T. This implies, by (8), a contemporaneous drop in average correlation, accompanied by an initial positive effect on R_{mT}. The shape of the confidence bands suggests that we can be 95% confident about the sign of these effects. In the following period, there is a positive affect on MKT_T and a negative one on R_{mT} (again, a rise in MKT_T is accompanied by a contemporaneous negative return). Therefore, a positive shock to $IDIO_T$ has a negative impact on average correlation and future returns. Guo and Savickas (2003) find evidence that this puzzling inverse relationship between aggregate stock market returns and lagged average idiosyncratic variance is also present in US stock data. To sum up, we find that the market return is positively and negatively associated with the lag of, respectively, market and average idiosyncratic variance. These findings suggest the possibility of a positive predictive relationship between aggregate stock market returns and our variance series.

Predictive Regressions

We now formally test for predictive relations between market returns and lags of our variance series. In Panel A of Table 4 we report predictive regressions of EMU and US average correlation series using a constant term and lagged values of both dependent variables as regressors. Interpreting these regressions as Granger-causality tests, we find that EMU average correlation is Granger-caused by US average correlation but this is not true the other way around. In particular, US average correlation explains alone 12 percent of its EMU counterpart and 12 percent of its own variation. EMU average correlation instead explains a mere 2 percent of its own variation and virtually none of the variation of its US counterpart. This confirms the common perception that the US stock market is an important risk factor for EMU markets. This finding has important implications both for pricing of EMU assets and from a risk management point of view. It is also relevant, from a systemic risk and financial stability perspective, for EMU financial regulators.

In Panel B of Table 4 we report predictive regressions of the stock market return, using a constant and our lagged variance series as regressors. We only report estimates of the regressions for EMU data (results for the US are already available in Guo and Savickas (2003)). We find that both market variance and average idiosyncratic variance help predict market returns but the relationship with lagged market variance is positive whereas the relationship with lagged idiosyncratic variance is negative. This result is consistent with our VAR analysis and, in particular, with the impulse response functions reported in Figure 5. The relationship between market returns and average idiosyncratic variance, however, is statistically significant at a conventional level only in the 1974-2004 period but not in the 1974-1997 one (interestingly, both in our EMU sample and in the US stock sample used by Guo and Savickas (2003)). The significance levels of the reported t-statistics are confirmed by a bootstrap experiment. Within a rational expectation framework, it is unlikely that lagged market variance be a proxy for expected market variance because of its low persistence[25]. This role, however, is not excluded if we drop the rational expectation assumption. Alternatively, within Merton's (1973) ICAPM framework, the circumstance that market variance and average idiosyncratic variance jointly predict future market-wide returns suggests that they might jointly proxy either for the predictable component of systematic risk or for changes to the future investment opportunity set. More generally, they might represent possible candidates as conditioning variables in asset pricing models with time-varying conditional risk premia.

[25] If market expectations about risk are not rational and investors' beliefs are consistent in the sense implied by Sargent's (1993) discussion of the rational expectation equilibrium framework,

Second Moments vs. Aggregate Returns

We finally ask what stylised features about the distribution of asset returns in EMU equity markets emerge from the analysis of the behaviour of the variance and correlation series. To this end, we regress our constructed variance and correlation series on contemporaneous and lagged aggregate market returns. The findings are reported in Table 5. At this (relatively low) frequency, there is no evidence that past negative returns lead to higher volatility for the average stock or to stronger average correlation among the stocks as implied, at lower frequencies (daily or weekly) by models of asymmetric conditional volatility and correlations such as the DCC-GARCH model employed by Cappiello, Engels and Sheppard (2003). There is, however, clear evidence that higher volatilities are associated on average with low values of the market returns and that, in these circumstances, average correlation also tends to be high (negative contemporaneous relationship of, respectively, average variance and correlation with market returns). On the basis of (7), this explains why market variance tends to be high when market returns are low (negative contemporaneous relationship of market variance with market returns). It therefore explains why the distribution of market returns is skewed to the left[26]. In turn, this implies that, in order to effectively model the asymmetry in the multivariate distribution of asset returns, the asymmetric behaviour of both asset volatilities and asset correlations should be taken into account.

As noticed by Cappiello, Engle and Sheppard (2003), little theoretical framework is available to explain the asymmetric relationship between correlations and negative and positive returns innovations. Tentatively, we could resort to the two popular explanations, the leverage and volatility feed-back effect, used to account for the relationship between volatilities and returns in a univariate setting and generalize them to the multivariate relationship between

investors should recognize that lagged market variance is a poor predictor of future variance and they should therefore look for a better proxy.

26 The skewness of the distribution of total weekly returns on the EMU Datastream equity index over the period 1974-1997 and 1974-2004 reported by Kearney and Potì (2003) is, respectively, -0.47 and -0.64 (significantly different from zero at, respectively, the 7% and the 13% level).

the latter and correlations. Suppose for simplicity that there are only two firms. Consider first that these firms own each other's liabilities (liabilities for the issuer, assets for the underwriter, the holder or the lender). Following negative news for one firm (firm A), the value of its assets decreases, both its equity and debt value decrease but the former decreases more because it is more junior. Its equity returns become thus more volatile because the firm's financial leverage increases. However, the other firm (firm B) owns firm's A liabilities. These are assets that are now worth less. The joint (system) leverage thus increases more than the leverage of firm A and the volatility of their aggregate return must increase more than the volatility of returns on the equity of firm A. This implies, from (8), that their correlation increases. If joint negative pieces of news depress the value of both firms' assets, this mechanism is mutually reinforced. We label this mechanism the *multivariate leverage effect*. Alternatively, negative market returns lead, by the volatility feed-back effect, to higher expected and realised market volatility. If the return on the market is the relevant risk factor, higher market variance implies, in a CAPM-like world, an increase to the variance of the average stock proportional to the explanatory power of the market model. Since the latter is the square of average stock correlation which, in turn, must be less than one, the variance of the average stock must increase less than proportionally relative to the variance of the market return. This implies, from (8), that the average stock correlation increases. We label this mechanism the *correlation feed-back effect*. A further mechanism that can explain the rise of correlations at times of low market returns, that we dub the *behavioural correlation effect*, would use frame dependence and loss aversion to explain the rise of expected returns at times of low market returns, as in Barberis, Huang and Santos (2001). Correlations would then rise because of the correlation feed-back effect, the correlation leverage effect or a combination of both.

We leave a more careful formalization of the *multivariate leverage effect*, *correlation feed-back effect* and *behavioural correlation effect* for future research. However, it is worth pointing out that, while high correlations conditional on low level of returns do not imply that the unconditional distribution of returns is not normal or, more generally, asymmetric (Forbes and

Rigobon (2002), Ang and Chen (2000)), these effects imply an asymmetric (skewed to the left) unconditional distribution of asset returns.

6. Implications for Portfolio Management and Asset Pricing

There are three main set of implications of our findings. The *first* relate to the consequences for portfolio management advice of the rise in average idiosyncratic volatility and of the level and dynamics of average correlation. The *second* set of implications relate to dynamic trading strategies based on the level of average correlation. The *third* set of implications relate to asset pricing, the predictability of asset returns and the volatility puzzle.

Portfolio Management Implications

A conventional rule of thumb, based on Bloomfield, Leftwich and Long (1977), suggests that a randomly chosen portfolio of 20 stocks produces most of the reduction in idiosyncratic risk that can be achieved through diversification. However, as suggested by CLMX (2001), the higher the average idiosyncratic variance, the larger the number of stocks needed to achieve a relatively complete diversification, given a random portfolio selection strategy. In Panel A of Figure 4 we report the residual portfolio idiosyncratic volatility as a function of the number of stocks included in equally-weighed portfolios formed by drawing randomly from our stock sample for various levels of average idiosyncratic risk at different points in time. To reduce idiosyncratic volatility to 5.0 percent it took 166 stocks were needed in 2003, 43 stocks in 1989 but just 35 stocks in 1974. It is worth noticing that most of the increase has taken place in the second half of the sample period. CLMX (2001)'s findings are similar. They report that a residual portfolio idiosyncratic volatility as low as 5 percent required 50 US stocks in the period 1986-1997 whereas it would have taken only roughly 20 stocks in the period 1974-1985[27].

[27] It would also have taken about the same number of stocks in the earlier 1962-1973 period.

On the other hand, the lower the correlation among stock returns, the higher the fraction of average total variance represented by idiosyncratic variance and the higher the potential benefit from diversification. The low level of average correlation, especially in the equally-weighted case, suggests that diversification can be an important source of improvement in the portfolio risk-return ratio (even though the full benefit in terms of variance reduction will not be available to the average investor since the value-weighted series is substantially higher than the equally-weighted one[28]). In EMU markets the potential diversification benefit is fairly stable over time as average stock correlation, while very noisy, is a fast moving variable that reverts back very quickly to a stationary long-run value. The potential diversification benefit using US stocks, however, might be slowly increasing over time, as US average stock correlation is mildly trended downward.

The low level of average correlation (especially in the equally weighted case) implies a low explanatory power of the market model. This means that, if the CAPM does not hold, the market model is able to accurately price only portfolios with a very large number of assets. The decrease of intra-industry correlations also leads to a weaker explanatory power for multi-factor models that, in the spirit of the APT, add industry indices alongside the market index to capture the sources of common variation in asset returns.

Trading Strategies Implications

Turning to trading implications of our findings and to the applicability of the methodology proposed for estimating the average correlation of a large set of assets, suppose that a derivatives trader observes Figure 2 and concludes that, since equally weighted average correlation is close to an all time low and value-weighted correlation is also below its historical average (but it seems to be picking up), it would be desirable to take a position that gained from an average correlation increase. Is it possible to construct such a position using commonly

[28] This also means that, trivially, the market portfolio is not the minimum variance one.

traded financial instruments? Of course, it is. Implied average correlation can be traded by trading options on a basket against a basket of options. To buy implied average correlation, e.g. on the Eurostoxx50 index, the investor should buy an "at the money option" (ATM) on the index and sell ATM options on the single stocks proportionally to the weights of the latter in the index. This position can be replicated by delta-hedging its mirror image (this way the investor can trade implied vs. realized average correlation). This is useful to hedge large portfolios of derivatives priced on the basis of given levels of implied average correlations. Risk premia for this trading activity are likely to be high, at least at the beginning, as the market for correlation risk is relatively untapped (trading it would make the market more complete). Estimating the risk premium that would accrue to a market maker that initiated such a trading activity would be a challenging research possibility. The mechanism of dynamic replication of exposures to the level of correlations is also useful for understanding how market expectations feed into actual correlations and it might suggest a less than perfect association between fundamentals correlations and returns correlations if market expectations about fundamentals are systematically biased.

Asset Pricing Implications

There are two main set of implications of our findings on market and average idiosyncratic volatility. *First*, the circumstance that they can be jointly used to predict stock market returns has conditional asset pricing implications. To see this, first consider why the relation between market return and lagged market and average idiosyncratic variance is, respectively, positive and negative. Changes in *IDIO* imply, by (8) and (9), changes of *CORR* of opposite sign if we keep *MKT* constant. For example, a positive shock to *IDIO*, keeping *MKT* constant, implies an increase of average total variance and a decrease of average correlation. Therefore, once we control for changes in market variance, average idiosyncratic variance is negatively correlated with average correlation. A negative sign of the regression coefficient on lagged *IDIO*, as reported in Table 4 (Panel B), therefore implies a positive association between market returns and lagged average correlation. This is the mechanism highlighted by the impulse response function to shocks to *IDIO* in Figure 5. Moreover, since *MKT* is, from (8), positively related to *CORR*, a positive association between the market return and lagged *CORR* implies a positive association between the former and lagged *MKT*. This can also be seen by considering that *CORR* is, by (8), a positive function of *MKT* and, by (9), a negative function of *IDIO*. Therefore their inclusion on the right hand side of the market return predictive regression is equivalent to regressing market return on the portion of average correlation that is their linear combination. We will denote this component of average correlation by $CORR^*$. Notice that, since the fitted market return $\hat{R}_m$ from the predictive regression is much more persistent[29] than both the actual return and average correlation, $CORR^*$ must be the persistent component of average correlation.

The interesting question then becomes: why do low predictable average stock correlations forecast low market returns? There are a number of plausible explanations that can be tentatively advanced, either from a rational asset pricing or from a behavioural perspective. To economize on space we will focus on the former. Low predictable average stock correlations imply by (7) low predictable stock market risk and therefore, within a Merton's (1973) ICAPM perspective, low expected and future market returns. Since $CORR_t^*$ is both persistent (slow moving) and it forecasts aggregate market returns, it also displays the two properties required for it to be a good candidate as a conditioning variable in models where risk premia are time-varying and the parameters in the stochastic discount factor conditionally depend on investors' expectations of future excess returns, such as the conditional version of the CAPM used by Lettau and Ludvigson's (2001). As nicely discussed in Cochrane (1999), the small predictability driven by a slow moving explanatory variable builds up over time, adding up to the substantial predictability of aggregate returns observed at long multi-annual horizons. Including average idiosyncratic variance amongst the set of conditioning variables used in a conditional CAPM empirical specification could be particularly useful because it is more readily available and it can be estimated with less delay than other successful candidates to this role, such as

[29] Its one and two-lag serial correlations are, respectively, 0.3729278 and 0.3540218.

Lettau and Ludvigson's (2001) consumption-wealth ratio. In a continuous time setting, the same empirical properties of average idiosyncratic volatility make it a suitable candidate as a conditioning variable that proxies for the state of the investment opportunity set in an empirical asset pricing specification based on Merton's (1980) Inter-temporal Capital Asset Pricing Model (ICAPM).

There are a number of other variables, such as Lettau and Ludvigson's (2001) consumption-wealth ratio, the dividend yield, the price-earnings ratio and the stochastically detrended short term interest rate, that to various degrees and over differing sample periods display this predictive ability. As suggested by Pesaran and Timmerman (1995, 2000), the researcher might use common model selection criteria like the AIC and the SBC to determine which variables should enter the predictive model in the period of interest. Arguing that there is learning in the marketplace and that predictive performance improves when one switches models over time based on formal model selection criteria, Pesaran and Timmerman (1995, 2000) and Dell'Aquila and Ronchetti (2004) suggest that this fact could be successfully exploited in investment strategies.

As shown in Table 6, average stock correlation is pro-cyclical. In particular, it is more pro-cyclical than market variance, whereas the market return displays an anti-cyclical property (trivially, by the forward looking nature of asset prices). The fitted market return $\hat{R}_m$ from the predictive regression is even more pro-cyclical, implying that the predictable portion of average stock correlation captured by $CORR_t^*$ is pro-cyclical too. This relation between (actual and predictable) average correlation and the state of the economy is in contrast with the possibility that average total and idiosyncratic variance influence asset prices because they proxy for the intensity of uninsurable idiosyncratic income shocks, as suggested by Goyal and Santa Clara (2003). Since idiosyncratic income shocks are concentrated in periods of economic recession, they play a key role in generating the mean equity premium, the low risk-free rate, and the predictability of returns in Constantinides and Duffie's (1996) model. This is nicely discussed in Constantinides (2002). The lower average correlation and the higher average

The *second* set of implications of our findings on average idiosyncratic risk relate to the volatility puzzle. In particular, if the relevant measure of risk was average total variance or some average of the former and of market variance, the high level of the average total variance relative to market variance (implied by the sizeable amount of average idiosyncratic variance) would render the volatility puzzle considerably milder. To see this, consider Hansen and Jagannathan (1991) volatility bound on the volatility of the stochastic discount factor, $\frac{E(R_m)}{\sigma(R_m)} \leq \sigma_m$. Here, m is the stochastic discount factor (essentially, the growth rate of consumption marginal utility), σ_m is its unconditional volatility and the market unconditional Sharpe ratio represents the lower bound on the volatility of the stochastic discount factor. Suppose that, with incomplete markets or because of market frictions and liquidity problems, investors are not able to fully diversify as recommended by financial theory. In this circumstance, the relevant unconditional Sharpe ratio would likely be an average of the *heterogeneous* Sharpe ratios of the individual investors and, as such, it would be the ratio of the market return to some convex combination of market volatility and the larger average total volatility. In the extreme case of investors being totally unable to diversify (each one holds only one asset), the minimum volatility bound on the stochastic discount factor would be reduced by a factor almost as large as one less the square root of the average stock correlation[30]. For example, with an average stock correlation equal to its long run average value (20.5 percent), the minimum volatility bound would be almost 55 percent lower. With an average stock correlation equal to 50 percent, the minimum volatility bound would be almost 29 percent lower. Notice that this discussion does not address the issue of whether models with incomplete markets and idiosyncratic sources of risks can solve the equity premium puzzle, as in Lettau (2001). Rather, more modestly, its aim is to point out that, if we accept that investors often hold imperfectly diversified portfolios, the puzzle itself may be milder than it is under the assumption that everybody truly holds the market portfolio. Therefore, while we accept Lettau's argument (2001) that models with uninsurable idiosyncratic income shocks cannot generate a volatility of the stochastic discount factor large enough to match the market portfolio Sharpe ratio of around 0.5, we argue instead that the latter might not be the relevant Sharpe ratio. We suggest that the relevant Sharpe ratio is instead an average of individual investors' portfolios (*heterogeneous*) Sharpe ratios. The higher the degree of diversification of the typical investor, the more the denominator of the ratio should be close to market volatility and the ratio itself should resemble the market Sharpe ratio. On the other hand, the lower the degree of investors' portfolios diversification, the more the denominator of the ratio should be close to average total volatility and the ratio itself lower than the market Sharpe ratio.

7. Summary, Conclusions and Future Research

In this paper we applied the variance decomposition proposed by CLMX (2001) and we derived the relation between market volatility, idiosyncratic volatility and average correlation. This derivation is to our knowledge novel. We applied this analytical framework to construct market variance, idiosyncratic variance (and its industry and firm components) and correlation series. We also constructed, for comparative purposes, analogous (albeit slightly shorter) US variance and

[30] This is an approximation because the square root is a concave function. By Jensen's inequality, the square root of the average total variance is therefore higher than the average total volatility (average of square roots of individual total variances). The true reduction in the bound is therefore slightly lower.

correlation series from CLMX (2001) data. Like in most empirical papers, our approach has been mainly descriptive. We applied econometric methods to infer the salient features of our constructed variance and correlation time series and we discussed their implications for portfolio management, trading strategies and financial asset pricing theory.

Regarding *long term* trends, our main findings are that, *first*, the variance of both the average European stock and of the EMU market portfolio has increased over time and that a large portion of this increase is explained by a long-run deterministic trend. European stocks, therefore, have indeed become more volatile. One consequence of the rise of average idiosyncratic risk is that it takes increasingly more stocks to capture the benefit of diversification. *Second*, value-weighted average stock correlation tends to mean revert quickly to a (roughly) 20 percent long-run mean after a shock and, as a consequence, idiosyncratic volatility accounts for the main portion of the variance of the typical stock. The potential benefits to diversification strategies are, therefore, substantial.

Regarding *short run* dynamics, EMU variance series are best forecast by market variance, whereas US variance series are best forecast by average idiosyncratic variance. Market and average idiosyncratic variance, as already documented by Goyal and Santa Clara (2003) and by Guo and Savickas (2003) using US data, predict market-wide returns. We suggested a number of implications that these findings have for asset pricing. Further investigating these implications, with special regards to the cross-section of average asset returns, might be a fruitful area for future research.

A useful extension of this work would be to study trends in average variance and correlation series constructed from daily instead of weekly returns. Care should be taken, in this case, to overcome the problems associated with asynchronous trading in the various European markets. Also, it would be useful to construct the average industry-level variance series directly from stock returns, instead of using Datastream industry indices.

Bibliography

Ang, A. and J. Chen (2000), "Asymmetric Correlation of Equity Portfolios", Working Paper, Columbia University.

Baele, L. (2002), "Volatility Spillover Effects in European Equity Markets: Evidence from a Regime-Switching Model", Working Paper.

Barber, B.M., and T. Odean (2000), "Trading Is Hazardous To Your Wealth: the Common Stock Investment Performance of Individual Investors", Journal of Finance 55, 773-806.

Barberis, N., M. Huang and T. Santos (2001), "Prospect Theory and Asset Prices", Quarterly Journal of Economics 116:1−53.

Barberis, N. and R.H. Thaler (2003), "A Survey of Behavioural Finance", in "Handboof of Economics of Finance", Eds. Constantinides, M.H. and R. Stulz, Elsevier Science B.V.

Benartzi, S. and R.H. Thaler (2001), "Naïve Diversification Strategies in Defined Contribution Saving Plan", American Economic Review 91, 79-98.

Bloomfield, T., R. Leftwich and J.B. Long Jr. (1977), "Portfolio Strategies and Performance", Journal of Financial Economics 5, 2001-2018.

Bollersev, T., R. Chou and K. Kroner (1992), "ARCH Modelling in Finance: a Review of the Theory and Empirical Evidence", Journal of Econometrics 52, 5-59.

Campbell, J.Y, M. Lettau, B.G. Malkiel and Y. XU (2001), "Have Individual Stocks Become More Volatile? An Empirical Exploration Of Idiosyncratic Risk", Journal of Finance 51, 1-43.

Campbell, J.Y, A.W. Lo and A.C. MacKinlay (1997), "The Econometrics of Financial Markets", Princeton University Press, Princeton, New Jersey.

Cappiello, L., R.F. Engle, and K. Sheppard (2003), "Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns", ECB Working Paper No. 204.

Cochrane, J.H. (1999), "New Facts in Finance", Economic Perspectives Federal Reserve Bank of Chicago 23, 36-58

Cochrane, J.H. (2001), "Asset Pricing", Princeton University Press, Princeton, New Jersey.

Dell'Aquila, R. and E. Ronchetti (2004), "Stock and Bond Retrun Predictability: the Discrimination Power of Model Selection Criteria", Universitè de Geneve Working Paper no 2004.05.

Enders, W. (2004), "Applied Econometric Time Series", Wiley, NJ.

Falkenstein E.G. (1996), "Preferences for Stock Characteristic as revealed by Mutual Fund Portfolio Holdings", Journal of Finance 51, 111-135.

Forbes, K. J. and R. Rigobon (2002), "No Contagion, Only Interdependence: Measuring Stock Market Co-movement", Journal of Finance 57, 2223-2261.

Fratzschler, M. (2002), "Financial Market Integration in Europe: On the Effects of EMU on Stock Markets", International Journal of Finance and Economics 7, 165-193.

French, K.R., Schwert G.W. and R. Stambaugh (1987), "Expected Stock Returns and Volatility," Journal of Financial Economics.

Goyal, A. and P. Santa-Clara (2003), "Idiosyncratic Risk Matters!", Journal of Finance 58, 9755-1008.

Guo, H. and R. Savickas (2003), "Idiosyncratic Volatility, Stock Market Volatility, and Expected Stock Returns", Federal Reserve Bank of St. Louis Working Paper nr 2003-028A.

Luigi Guiso, Haliassos, M. and T. Jappelli (2002), "Household Stockholding in Europe: Where Do We Stand and Where Do We Go?", CSEF Working Paper no 88, Salerno University.

Hansen, L.P. and R. Jagannathan (1991), "Implications of Securities Market Data for Models of Dynamic Economies", Journal of Political Economy 99, 225-262.

Hardouvelis, G., D. Malliaropulos and R. Priestley (2000), "EMU and European Stock Market Integration", CEPR Discussion Paper.

Harvey, C.R. (1989), "Time-varying Conditional Covariances in Tests of Asset Pricing Models", Journal of Financial Economics 24, 289-317.

Kearney, C. and V. Potì (2003), "Idiosyncratic Risk, Market Risk and Correlation Dynamics in European Equity Markets", EFA 2003 Annual Conference Paper nr. 744, http://ssrn.com/abstract=424521.

Lettau, M. (2001), "Idiosyncratic Risk and Volatility Bounds, or Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle?", Working Paper (forthcoming in the Review of Economics and Statistics).

Lettau, M. and S. Ludvigson (2001), "Consumption, Aggregate Wealth and Expected Stock Returns", Journal of Finance 56, 815-849.

Lintner, J. (1965), "Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets", Review of Economics and Statistics, 47, 13-37.

Malkiel, B.G and Y. Xu (2002), "Idiosyncratic Risk and Security Returns", Working Paper.

Merton, R.C. (1980), "On Estimating the Expected Return on the Market: an Exploratory Investigation", Journal of Financial Economics 8, 323-361.

Newy, W.K. and D.K. West (1987), "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix", Econometrica 55, 703-8.

Pesaran, M.H. and A. Timmerman (1995), "Predictability of Stock Returns: Robustness amd Economic Significance", Journal of Finance 50, 1201-1228.

Pesaran, M.H. and A. Timmerman (2000), "A Recursive Modelling Approach to Predicting UK Stock Returns", The Economic Journal, 150-191.

Schwert, G.W. (1989), "Why Does Stock Market Volatility Change Over Time?", Journal of Finance 44, 1115-1153.

Sargent, T. (1993), "Bounded Rationality in Macroeconomics", Oxford University Press.

Scruggs, T.J. (1998), "Resolving the Puzzling Intertemporal Relation between the Market Risk premium and Conditional Market Variance: A Two Factor Approach", Journal of Finance 53, 575-603.

Sharpe, W.F. (1964), "Capital Asset Prices: a Theory of Market Equilibrium Under Conditions of Risk", Journal of Finance 19, 425-442.

Sims, C.A. (1980), "Macroeconomics and Reality" , Econometrica 48, 1-49.

Tuner, C.M., R. Startz and C.R. Nelson (1989), "A Markov Model of Heteroskedasticity, Risk, and Learning in the Stock Market", Journal of Financial Economics 25, 3-22.

Figure 1:
Value-Weighted Variance Time Series

Panel A

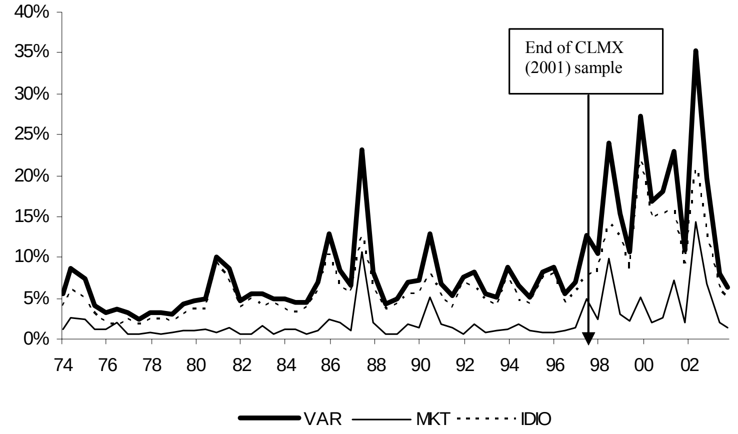

Panel B

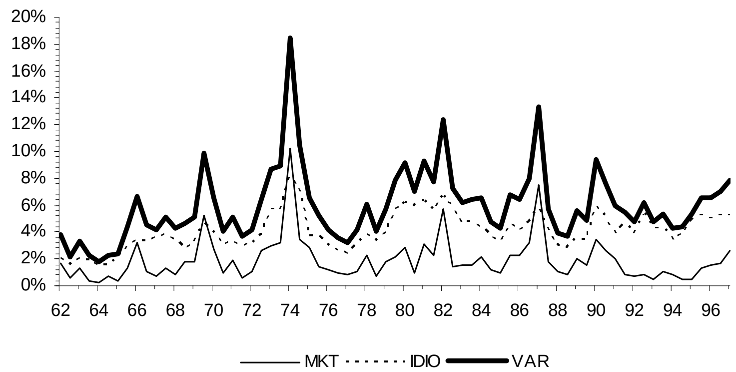

Panel C

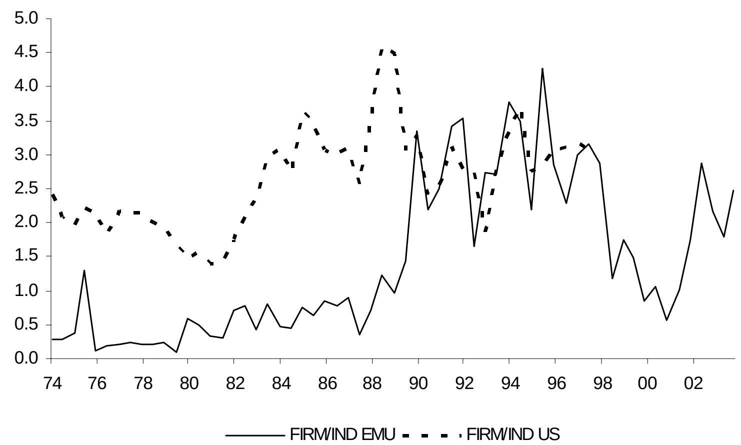

Note. Panel A and B plot the decomposition of average total stock variance into systematic and average idiosyncratic variance for, respectively, a value-weighted portfolio of 3515 stocks listed on the EMU stock markets from 1974 to 2004 and for the US stocks listed on the NYSE, NASDAQ and AMEX over the period 1962-1997. Panel C plots the ratio of $FIRM_T$ to IND_T for the EMU and US series.

Figure 2:

Average Correlation Time Series

Panel A

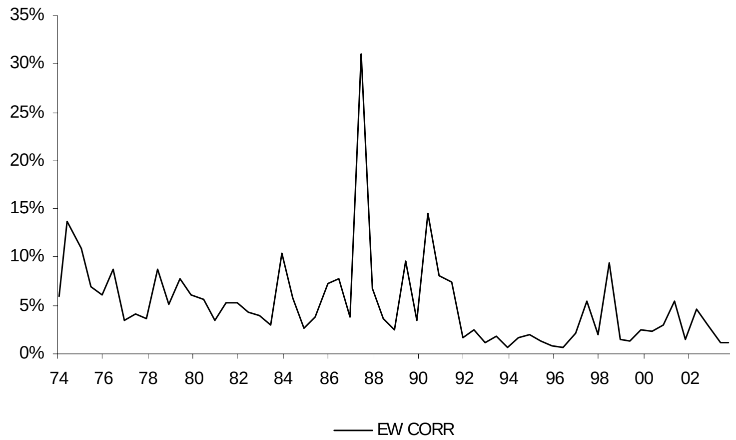

Panel B

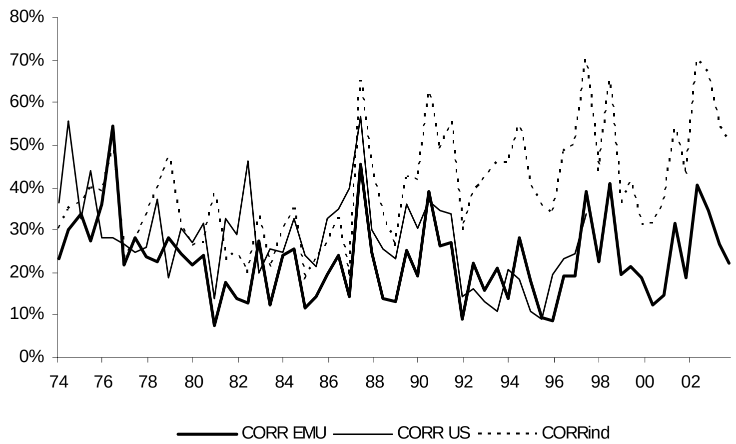

Note. Panel A of this Figure plots the equally weighted average correlation (*EW CORR*) amongst the firms included in a portfolio of 3515 stocks listed on the *EMU* stock markets over the period 1974-2004; Panel B plots the value-weighted average correlation amongst these firms (*CORR EMU*), amongst the 35 Datastream Level 4 industry indices for the *EMU* stock markets ($CORR^{ind}$) over the period 1974-2004 and amongst the US stocks in the CLMX (2001) sample for the period 1974-1997 (*CORR US*).

Figure 3
Diversification in EMU Industry Indices

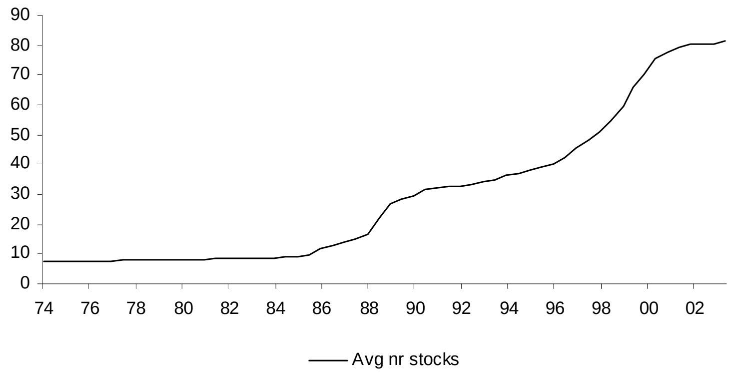

Note. This Figure plots the number of listed stocks in the average EMU industry over the period 1974-2004.

Figure 4:

Residual Portfolio Idiosyncratic Volatility

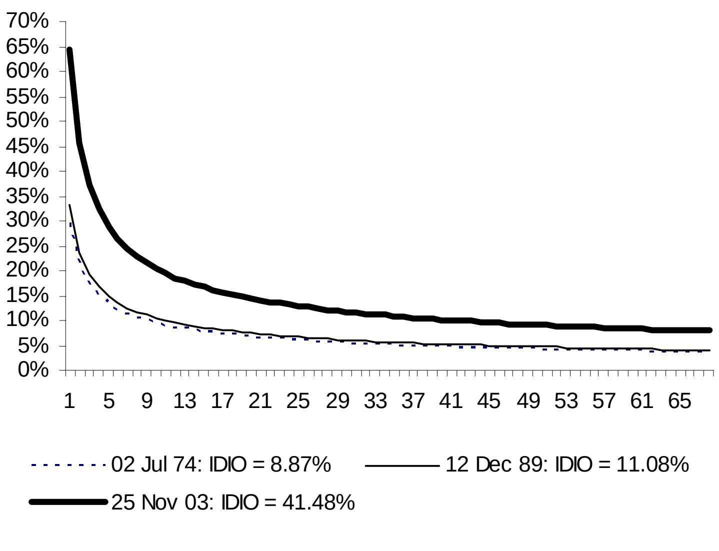

Notes. This figure reports the residual undiversified idiosyncratic volatility as a function of the number of stocks included in equally-weighted portfolios formed by randomly drawing from our stock sample at different points in time with varying levels of average idiosyncratic variance.

Figure 5:
Impulse Response Functions

Panel A
(Impulse Response Functions)

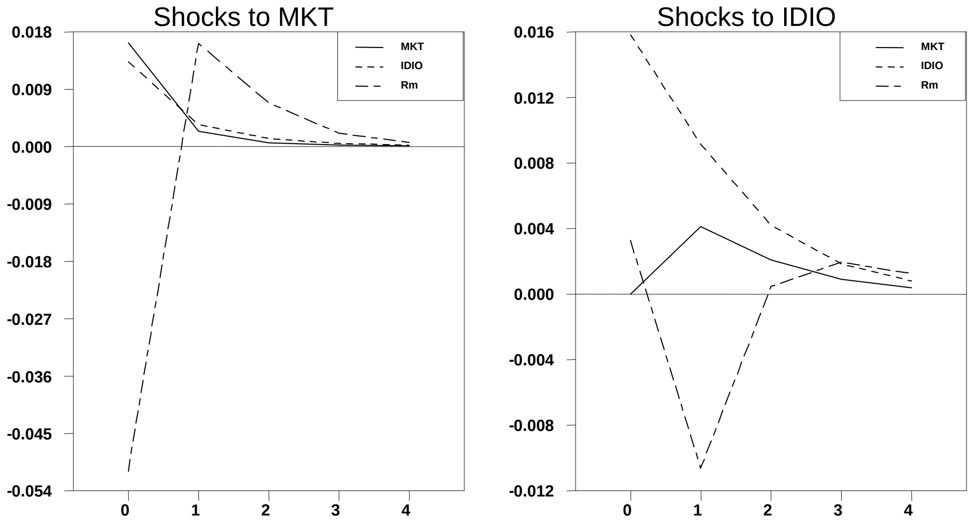

Panel B
(Confidence Intervals)

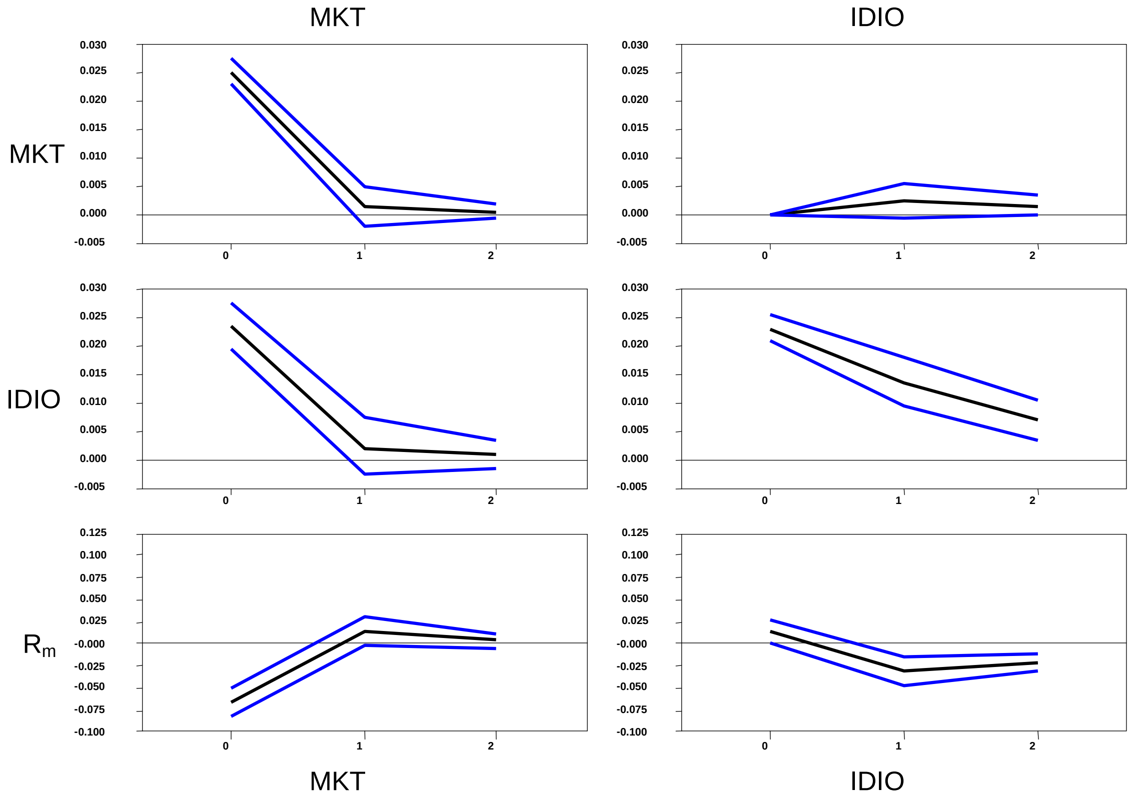

Notes. This figure plots the impulse response functions and their 95% confidence bands (constructed using a Montecarlo integration procedure) of the MKT, IDIO and R_m series to shocks to MKT and IDIO (value-weighted, EMU data). The symbols retain the usual meaning as in the text. The sample period is 1974-2004.

Table 1

Descriptive Correlations

Panel A

	Lags	MKT^{EMU}	$IDIO^{EMU}$	MKT^{US}	$IDIO^{US}$
	0	1.00	0.62	**0.62**	0.30
MKT^{EMU}	1	*0.04*	0.17	**0.24**	0.11
	2	-0.03	-0.02	0.03	-0.07
	0	0.62	1.00	0.52	**0.50**
$IDIO^{EMU}$	1	0.01	*0.30*	0.17	**0.21**
	2	-0.15	-0.06	-0.01	0.03
	0	**0.62**	0.52	1.00	0.68
MKT^{US}	1	**0.04**	0.24	*0.23*	0.32
	2	-0.00	0.28	0.12	0.14
	0	0.30	**0.50**	0.68	1.00
$IDIO^{US}$	1	-0.02	**0.28**	0.38	*0.64*
	2	-0.24	0.12	-0.06	0.31
	0	-0.47	-0.29	-0.36	-0.10
r_m^{EMU}	1	0.14	-0.01	0.27	0.18
	2	-0.12	-0.04	0.03	0.14
	0	-0.27	-0.25	-0.24	-0.05
r_m^{US}	1	0.10	-0.05	0.35	0.10
	2	-0.05	0.04	-0.08	-0.01

Panel B

	Lags	MKT^{EMU}	$IDIO^{EMU}$	$CORR^{EMU}$	MKT^{US}	$IDIO^{US}$	$CORR^{US}$
	0	0.67	-0.01	1.00	0.29	0.01	**0.45**
$CORR^{EMU}$	1	0.11	0.02	*0.20*	0.22	0.02	**0.38**
	2	-0.01	-0.15	0.19	0.02	-0.11	0.26
	0	0.57	0.31	**0.45**	0.80	0.35	1.00
$CORR^{US}$	1	0.16	0.23	**0.09**	0.20	0.19	*0.41*
	2	0.11	0.23	0.03	0.26	0.11	0.41

Notes. Correlations of the variables reported in the first column with lags of the variables reported at the top of the other columns over the period 1974 – 1997 (correlations for the period 1974-2004 for EMU stocks, available upon request, are very similar). First order autocorrelation coefficients are highlighted in italics.

Table 2
Specification and Wald Tests

	Static Model	**Dynamic Model**				
	DW-stat.	α [*t*-stat.]	δ [*t*-stat.]	β [*t*-stat.]	*h-stat.* [sign.]	*Wald-stat.* [sign.]
		Panel A (EMU 1974-2004)				
$IDIO_T$	1.26	1.16%	.10%	35.86%	.52	9.55
		[1.29]	[3.09]	[2.80]	[.47]	[.003]
$FIRM_T$	1.67	-.40%	.11%	12.79%	2.30	24.81
		[-.88]	[4.98]	[.93]	[.13]	[.000]
IND_T	.99	1.1%	.01%	50%	1.47	.79
		[1.88]	[.89]	[4.34]	[.22]	[.375]
MKT_T	1.86	.35%	.056%	5.91%	.35	7.99
		[.53]	[2.83]	[.44]	[.55]	[.006]
VAR_T	1.54	1.67%	.17%	21.32%	.33	11.66
		[1.15]	[3.41]	[1.60]	[.56]	[.001]
$CORR_T$	1.65	20.50%	-.04%	16.98%	2.37	.37
		[4.98]	[-.60]	[1.30]	[.12]	[.54]
		Panel B (US 1974-1997)				
$IDIO_T$	.72	1.61%	.001%	63.53%	2.38	.02
		[2.55]	[.16]	[5.48]	[.12]	[.870]
$FIRM_T$	.69	1.0%	.006%	63.99%	3.27	.92
		[2.53]	[.96]	[5.65]	[.07]	[.341]
IND_T	.67	.5%	-.004%	66.39%	1.24	.80
		[2.38]	[-.89]	[5.86]	[.26]	[.374]
MKT_T	1.54	2.28%	-.028%	22.42%	3.24	2.089
		[3.26]	[-1.44]	[1.51]	[.071]	[.155]
VAR_T	1.14	4.26%	-.019%	42.69%	2.88	.487
		[3.20]	[-.69]	[3.11]	[.089]	[.488]
$CORR_T$	1.39	24.58%	-.197%	29.50%	3.20	3.093
		[4.04]	[-1.75]	[2.00]	[.073]	[.085]

Notes. This tables reports estimates of the parameters of the model of the variance and correlation series with a deterministic time trend. All the variables are defined as in the text. All the series are semi-annual (annualised). *DW* denotes the Durbin-Watson statistics of the static model. All the other columns report estimated coefficient and *t*-statistics for the dynamic model. The rightmost columns report the Durbin's *h*-statistic of the null that the dynamic model residuals are not first-order autocorrelated and the Wald statistic (in both cases with the associated significance levels) of the restriction that δ is equal to zero. All the Wald-Test statistics, standard errors and significance levels have been computed using a Newy-West adjusted variance–covariance matrix with Parzen weights to correct for heteroschedasticity and autocorrelation.

Static Model:

$$y_t = \alpha + \delta t + u_t \qquad u_t \sim i.i.d.\ N(0, \sigma^2)$$

Dynamic Model:

$$y_t = \alpha + \beta\, y_{t-1} + \delta t + u_t \qquad u_t \sim i.i.d.\ N(0, \sigma^2)$$

Table 3
VAR Analysis

Panel A
(Block-exogeneity Tests)

Variable	$ln\lvert\Sigma_R\rvert$	$ln\lvert\Sigma_U\rvert$	Chi-Squ.(2)	Sig.
		EMU		
MKT_T	-16.336	-16.236	5.476	0.064
IND_T	-16.214	-16.204	0.580	0.748
$FIRM_T$	-15.856	-15.840	0.851	0.653
		US		
MKT_T	-22.415	-22.411	0.213	0.898
IND_T	-18.981	-18.911	3.177	0.204
$FIRM_T$	-19.904	-19.511	17.656	0.000

Panel B
(Decomposition of Variance)

Series	**Step**	**MKT_T**	**IND_T**	**$FIRM_T$**
MKT_T	1	100	0	0
	2	98.5	0.5	1
	3	98.4	0.6	1
IND_T	1	32	68	0
	2	25	74	1
	3	23	76	1
$FIRM_T$	1	45.5	1.1	53.4
	2	45	2	53
	3	45	2	53

Notes. Panel A of this table reports the log-determinants of the unrestricted ($ln|\Sigma_U|$) and restricted ($ln|\Sigma_R|$) 2-variable VAR systems where the variable specified in the left-most column is restricted to be block-exogenous. The sample period is 1974-2004 for the EMU and 1974-1997 for the US. The Chi-Squared statistic is computed as $(T - c)(ln|\Sigma_R| - ln|\Sigma_U|)$, where $T = 61$ and $T = 48$ for, respectively, the EMU and US sample and c is Sims' (1980) multiplier correction. Panel B reports the percentage of the variance of the series (computed using EMU data over the period 1974-2004) reported in the first column explained by the series reported at the top of each row. The variance decomposition imposes the restriction that IND has no contemporaneous effect on MKT and FIRM has no contemporaneous effect on MKT and on IND. All the variables are linearly de-trended.

Table 4
Average Correlation and Market Return Predictive Regressions

Dependent Variable	**Explanatory Variables**			
	Panel A (EMU and US Average Correlations Predictive Regressions)			
	1994-1997			
$CORR_T^{EMU}$	$CORR_{T-1}^{US}$		.33	.34
			[2.78]	[3.88]
	$CORR_{T-1}^{EMU}$	0.20	0.03	
		[1.41]	[0.21]	
	Adj. R^2	0.02	0.10	0.12
$CORR_T^{US}$	$CORR_{T-1}^{US}$	0.40	0.45	
		[2.93]	[2.61]	
	$CORR_{T-1}^{EMU}$		-0.12	0.10
			[-0.76]	[0.88]
	Adj. R^2	0.14	0.14	-0.01
	Panel B (EMU Market Return Predictive Regressions)			
	1974-1997			
$r_{m,T}$	MKT_{T-1}	0.96		1.59
		[1.97]		[2.02]
	$IDIO_{T-1}$		-0.03	-0.78
			[-0.06]	[-1.15]
	Adj. R^2	0.02	0.0	0.03
	1974-2004			
$r_{m,T}$	MKT_{T-1}	0.42		1.95
		[0.62]		[2.67]
	$IDIO_{T-1}$		-0.69	-1.55
			[-2.28]	[-3.83]
	Adj. R^2	0.01	0.04	0.12

Notes. Panel A and Panel B report, respectively, predictive regressions of EMU and US average correlations and of the EMU market return. In brackets are t-statistics adjusted for heteroskedasticity and auto-correlation and regressions always include a constant.

Table 5
Second Moments vs. Market Return

Dependent Variable	Explanatory Variables		
MKT_T^{EMU}	$r_{m,T-1}^{EMU}$	-0.01	
		[-0.41]	
	$r_{m,T}^{EMU}$		-0.10
			[-2.88]
	Adj. R^2	-0.01	0.29
VAR_T^{EMU}	$r_{m,T-1}^{EMU}$	-0.03	
		[-0.56]	
	$r_{m,T}^{EMU}$		-0.21
			[-2.92]
	Adj. R^2	-0.01	0.23
$IDIO_T^{EMU}$	$r_{m,T-1}^{EMU}$	-0.02	
		[-0.69]	
	$r_{m,T}^{EMU}$		-0.11
			[-2.69]
	Adj. R^2	-0.01	0.15
$CORR_T^{EMU}$	$r_{m,T-1}^{EMU}$	-0.02	
		[-0.25]	
	$r_{m,T}^{EMU}$		-0.25
			[-2.49]
	Adj. R^2	-0.01	0.09

Notes. This table reports regressions of the variance and correlation series on contemporaneous and lagged market returns over the period 1974-2004 (parameters estimates and significance levels are similar for the shorter 1974-1997 period). The reported t-statistics (in squared brackets) are adjusted for heteroskedasticity and auto-correlation and regressions always include a constant. All variables are de-trended and all regressions include a constant.

Table 6
EMU GDP Growth Correlations

Leads of GDP Growth	CORR	MKT	R_m	$\hat{R}_m$
2	0.13	0.26	-0.37	0.16
1	-0.15	-0.10	-0.03	-0.15
0	0.19	0.14	-0.17	0.20
-1	0.02	-0.07	0.22	0.01
-2	0.16	-0.00	0.10	0.16

Notes. This table reports the correlations of leads of GDP growth with the variables reported at the top of each column (average stock correlation, market variance and the actual and predicted market portfolio return) using EMU data over the period 1980-2004.

Data Appendix
Data and Variables Definitions

Panel A

EMU Countries (as of March 2004)			
1	Austria	7	Ireland
2	Belgium	8	Italy
3	Denmark	9	Luxembourg
4	France	10	Netherlands
5	Germany	11	Norway
6	Greece	12	Spain

Panel B

Industries – Datastream Level 4			
1	MINING	19	RETAIL, GENERAL
2	OIL & GAS	20	LEISURE + HOTELS
3	CHEMICALS	21	MEDIA, ENTERTAIN
4	CONS.&BLDG MAT.	22	SUPPORT SERVICES
5	FORESTRY&PAPER	23	TRANSPORT
6	STEEL&OTH.METALS	24	FOOD&DRUG RETLRS
7	AEROSPCE,DEFENCE	25	TELECOM SERVICES
8	DIVERSIFIED INDS	26	ELECTRICITY
9	ELECTR. EQUIP.	27	OTHER UTILITIES
10	ENG.&MACHINERY	28	INF.TECHN.HARDW.
11	AUTO & PARTS	29	SFTWR&COMP.SERV.
12	H'HLD GDS&TEXTLS	30	BANKS
13	BEVERAGES	31	INSURANCE
14	FOOD PRDR./PRCR.	32	LIFE ASSURANCE
15	HEALTH	33	INVESTMENT COS.
16	PER.CARE&HSHLD	34	REAL ESTATE
17	PHARM. & BIOTECH	35	SPC&OTH. FINANCE
18	TOBACCO		

Panel C

Variables		
1	$\mathbf{r_i}$	Weekly return on industry i
2	$\mathbf{r_{i,j}}$	Weekly return on stock j from industry i
2	$\mathbf{r_m}$	Weekly return on the stock market portfolio
3	**VAR**	Average total variance of stock returns
4	**MKT**	Annualised semi-annual variance of r_m
5	**IDIO**	VAR – MKT
6	$\mathbf{VAR^{ind}}$	Average total variance of industry returns
7	**IND**	VAR^{ind} – MKT
6	**FIRM**	VAR – VAR^{ind}

Notes. Panel A of this table reports the industries included in our sample based on the Datastream Level 4 classification. Panel B reports the countries. Panel C summarizes the main variables. The market portfolio is the Datastream EMU index for the EMU and the CRSP index for the US. All returns are total returns (they include accrued dividends). All indices are "fixed history" (they are not recalculated following modifications to the index composition).