Dispersion as Cross-Sectional Correlation

Bruno Solnik and Jacques Roulet

We introduce the concept of cross-sectional dispersion of stock market returns as an alternative to the time-series approach to estimating the global correlation level of equity markets. Our objective is to derive a simple, instantaneous measure of the general level of global market correlation. Our cross-sectional method of estimating global correlation is dynamic and, using cross-sectional data, gives instantaneous information on the trend of global correlation. The traditional time-series method requires a long period of observations, and overlapping data have to be used to study the change in correlation. Both methods yield similar estimates for a "long" period, however, so a combination of the cross-sectional and time-series approaches should be of practical use to global asset managers.

The case for international investing has been built on the rather low correlations between national equity markets. If global correlations are low, spreading investments among countries allows diversification of the total risk of a portfolio and active managers have more opportunities to find assets and markets that will have a large return. So, from both a risk and a return viewpoint, low global correlations are beneficial.

The traditional approach to estimating the extent of markets' comovement is to conduct a time-series estimation of market correlations over a fixed period of time—for example, an estimation window of five years of monthly data.[1] This method is poorly suited to studying changes in correlations over time because a large number of observations are required to estimate just one correlation coefficient. To study time variation in correlation, one has to resort to overlapping observations with a rolling estimation window. This method is not satisfactory because two successive correlation estimates are based on almost the same data set and differ only because one "old" observation is dropped and replaced by a "new" one. A long time is needed for a permanent change in the general level of global correlation to be reflected in the estimation. A temporary change is hardly noticed because it affects only a few observations in the estimation window.

The issue of changes in the level of global correlation has strong practical relevance. Investment managers optimize their asset allocations partly on the basis of the international covariance structure of market returns. The level of global correlation affects the degree of diversification needed in their portfolios, and the size of the excess returns of asset classes is determined by the dispersion. What managers need is an indicator that will instantaneously track the time variation in global correlations. In the absence of such a quantitative indicator, they are left with "emotional impressions" that can be misleading and difficult to integrate into a structured, quantitative asset allocation process.

We propose a new and simple way to measure the instantaneous changes in the general level of global correlation. The method, based on the cross-sectional dispersion of returns, is a simple model linking cross-sectional dispersion to global correlation. It needs refinement, and its practical implementation in asset management is not straightforward, but we suggest here various ways to make the concept of cross-sectional global correlation useful for practitioners.

"Refreshed" estimates of market correlations are particularly important to portfolio managers in periods when market returns are fluctuating widely over time. One reason is that the globalization of investments and the instantaneous flow of information have rendered markets prone to contagion whenever a national or regional crisis devel-

Bruno Solnik is professor of finance at the HEC School of Management, France. Jacques Roulet is head of Quantitative Research & New Product Development at Morgan Stanley Capital International. At the time this article was written, Mr. Roulet was with Lombard Odier & Cie.

ops. Our dispersion measure and the global correlation measure that we derive from it are practical, albeit partial, answers to this need. They allow a manager to instantaneously detect changes in correlations in the international marketplace.

Time-Series Approach

The traditional approach is to measure the time-series correlation for each market pair over a fixed time period. Typically, five years of monthly data (the estimation window) are used because a sufficient number of observations (e.g., 60 months of data) is required to derive statistically significant estimates. (From now on, we use one month as the *observation period* and five years as the *estimation window* to illustrate the discussion.) The distribution of returns is assumed to be multivariate normal with constant parameters. Basically, correlation, volatility, and expected returns are assumed to be constant over the five-year estimation window.

In this method, one computes the mean five-year return for each market and studies the deviations of monthly returns from their mean returns for a market pair. Several points should be stressed:

- Each pairwise correlation coefficient is computed separately. To address the issue of the general level of global correlation, however, one needs to compute a cross-sectional average of the correlations across all markets.
- The time-series method provides an *unconditional* estimate; correlation is assumed to be constant over the 60-month estimation window, as are return distributions. Because correlation needs to be estimated over a time series of data, concluding anything about the *change* in correlation over time is difficult, at least in a statistically meaningful way. For example, a total period of 15 years yields only three independent observations of the 60-month correlation coefficient. Although the "moving correlation" can be computed with overlapping estimation windows (by replacing one monthly observation each month), two successive correlation estimates are strongly, and somewhat spuriously, dependent.

Refinements to the time-series method have been proposed. Pairwise correlations can be computed using weights that give more importance to recent periods. This approach is somewhat *ad hoc* because the set of weights is chosen arbitrarily. An attempt to simultaneously use the information about all markets can be based on the multivariate GARCH (generalized autoregressive conditional heteroscedasticity) approach with time-varying covariances.[2] Unfortunately, this approach has severe limitations, as discussed in Kroner and Ng (1998). A "larger" model in use is the VECH model, in which the covariance between two markets is estimated as a weighted average of past cross-products of the unexpected returns of two markets. The time-varying covariance between two markets is still only a function of the returns on those two markets (a "diagonal" model), and the approach has two practical shortcomings: The number of parameters to be estimated is huge (630 for 20 countries), even in a diagonal model, and it is difficult to ensure that the correlation matrix is indeed semi-definite positive as it should be. To circumvent this nonpositiveness problem, a BEKK model has been proposed. Unfortunately, this model calls for 1,010 parameters to be estimated for 20 countries, and the model is still diagonal. Simpler models can be used, but their treatment of time variation in correlations is even more simplistic.

Cross-Sectional Approach: Dispersion

Another way to look at the general issue of stock market comovement and global risk measurement involves the dispersion of returns to the various markets for any given observation period. This approach is a cross-sectional method rather than a time-series one and works as follows:

Let us define world return as the average return on all markets. If markets are "moving together" ("highly correlated"), all markets will provide a similar return in any given month. Although the world return will vary over time, the dispersion of all national market returns around this world return will be small in any given month. If, in contrast, there is a large dispersion of national market returns around the world return in a given period, markets are not really moving together and managers have ample room for global risk diversification (or profit opportunities for active investors).

Take 1997 as an example. In that year, the return (in U.S. dollars, dividends reinvested) on the Morgan Stanley Capital International (MSCI) World Index was 16.3 percent. The dispersion of returns in U.S. dollars among national markets was large. For example, for Italy, it was 35.8 percent; for Germany, 43.8 percent. Switzerland, at 44.4 percent, topped the list; France, at 13.1 percent, was somewhere in the middle; and the worst performers were Singapore, at –30.2 percent, and Japan, –23.8 percent. This dispersion of returns is clearly large.

The approach illustrated in this article is to calculate, for each period, the cross-sectional dispersion of country index returns around the world return, as measured by the standard deviation of national market returns for that period. All returns were measured in a common currency, and we use the U.S. dollar in our examples.

The observation frequency had to be selected. Daily returns could not be used because of the international time difference. Markets are not open at the same time, and there is often little or no overlap in trading in different regions. This nonsimultaneity problem also pollutes estimations based on weekly returns (one day out of five), but it is less serious for estimations based on monthly returns. Using monthly data, we could obtain one independent measure of dispersion each month, and we could study its change over time.

The inverse relationship between cross-sectional dispersion and time-series correlation will be detailed later; we will first look at the data and illustrate the concept of dispersion.

Empirical Investigation. We used monthly returns in U.S. dollars from January 1971 to September 1998 for the 15 developed stock markets originally included in the MSCI World Index.[3] A list of these markets is given in **Table 1**, together with return and correlation data. For each month t, we estimated the dispersion of returns, as measured by the cross-sectional standard deviation of the 15 market returns, $\sigma_c(t)$. These data are plotted in **Figure 1**. The monthly dispersion looks remarkably stable over time and ranges around an average value of 4.5 percent a month, or 15.6 percent a year.[4] The monthly world return fluctuates widely from one month to the next, but it has little influence on the dispersion. For example, a large negative world return, –20 percent, appears in October 1987, but the global dispersion is equal to 8.5 percent in that month. A similar dispersion was observed in January 1987, when the monthly return was a strong 7 percent, and in July 1986, when the world return was zero. A simple time-series regression between the world return and monthly dispersions confirms the visual impression, with an adjusted R^2 of only 1.0 percent. Dispersion was high in the 1973–75 period and appears to have been lower in the 1994–97 period. It increased again in 1998. The "Fitted Dispersion Line" in Figure 1 is the regression line between dispersion and time. Dispersion slowly declined with time, which suggests that markets have been increasingly moving together, but the R^2 is only 5 percent. The "Fitted Dispersion Line" moved from 5.1 percent a month at the start of 1971 to 3.9 percent a month in September 1998 (but the actual dispersion in 1998 was well above the fitted line), and the slope has a t-statistic of 4.28.

To summarize, the dispersion around the world return has remained quite stable: It (slowly) decreased in the first 27 years of the studied period but jumped up again in 1998. In other words, the potential benefits of international diversification have not disappeared over the past 20 years.

Table 1. Some Statistics on the Stock Markets in the Study

Country	Beta	Standard Error	R^2 with World	Correlation with World
Australia	1.051	0.076	0.354	0.595
Belgium	0.988	0.046	0.587	0.766
Canada	0.867	0.053	0.484	0.696
Denmark	0.788	0.049	0.493	0.702
France	1.201	0.063	0.475	0.690
Germany	1.015	0.053	0.526	0.725
Italy	1.070	0.082	0.326	0.571
Japan	0.878	0.068	0.361	0.601
Netherlands	1.003	0.037	0.683	0.826
Norway	1.240	0.082	0.359	0.599
Spain	0.971	0.066	0.399	0.631
Sweden	1.011	0.061	0.446	0.668
Switzerland	1.018	0.048	0.573	0.757
United Kingdom	1.211	0.068	0.438	0.662
United States	0.687	0.034	0.644	0.802

Figure 1. Dispersion of Monthly Returns for 15 Developed Markets

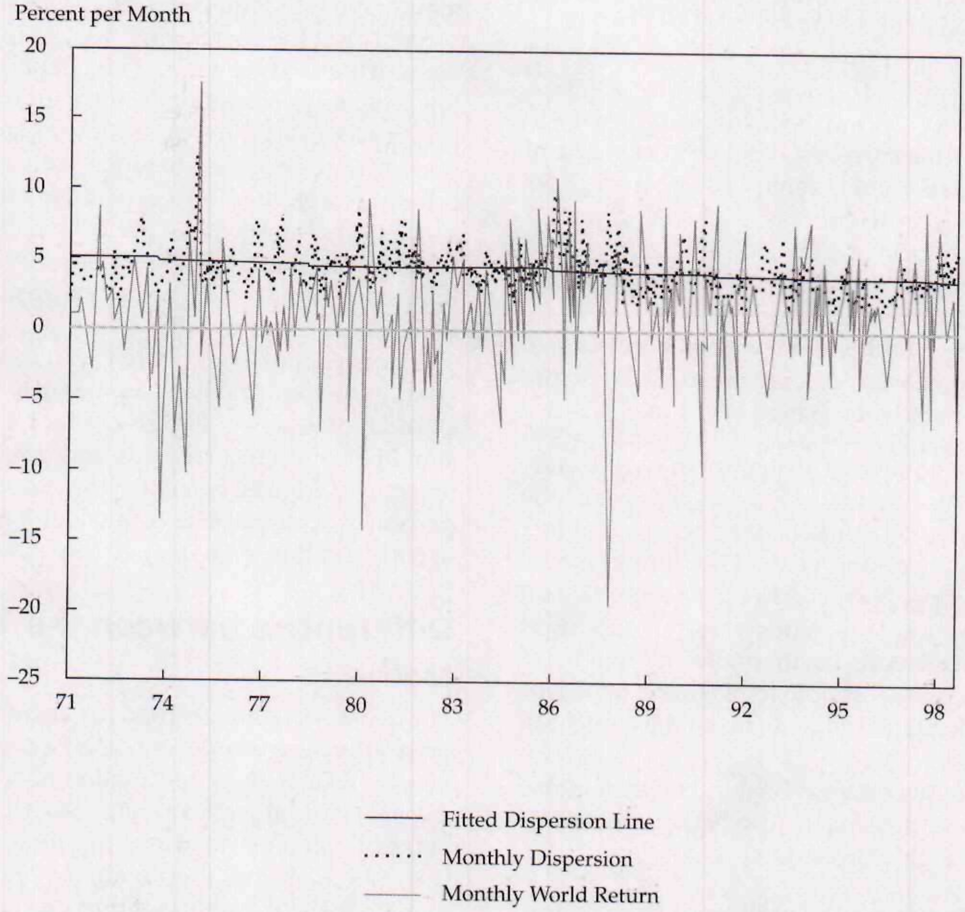

Note: Ending date is September 1998.

We now turn to a formal discussion of the link between time-series correlation and cross-sectional dispersion.

Simple Model of Cross-Sectional Correlation. Assume that the return on any national market i at time t, $R_{i,t}$, is driven by the return on the world market, $R_{w,t}$, plus an independent term $e_{i,t}$. Then,

$$R_{i,t} = R_{w,t} + e_{i,t}. \tag{1}$$

The return on the world market is the average return on all markets, and for simplicity, assume an equally weighted average. (A market-capitalization-weighted average could be used, but it would require the knowledge of the market weights at each point in time.) Equation 1 assumes that all national markets have a beta of 1 relative to the world index. This assumption is made to simplify the mathematical exposition and can be eliminated, as shown later.

Further assume that deviations from the world index are multivariate normal and independent:

$$e_{i,t} \sim N[0, \sigma_e(t)]. \tag{2}$$

Equation 2 implies that in any month, the return of all national markets deviates from the world return by a "tracking error" that is drawn from the same normal distribution, but the standard deviation of this tracking error, $\sigma_e(t)$, can fluctuate over time. As stated earlier, this standard deviation of returns on all markets around the world is called "dispersion." Because the world market is the equally weighted average of all markets, its volatility, which is identified as $\sigma_w(t)$, can change over time.

At each point in time, the correlation between market i and the world market, $\rho_{i,w}(t)$, can be time varying. Given Equations 1 and 2, the instantaneous correlation $\rho_{i,w}(t)$ is

$$\rho_{i,w}(t) = \frac{\sigma_w(t)}{\sigma_i(t)} = \frac{1}{\sqrt{1 + \sigma_e^2(t)/\sigma_w^2(t)}}. \tag{3}$$

To simplify the notation, we now drop the time notation, but the reader should be aware that all variables are assumed to be time varying:

$$\rho_{i,w} = \frac{1}{\sqrt{1+\sigma_e^2/\sigma_w^2}}. \quad (4)$$

Similarly, the correlation between any pair of national markets i and j is simply

$$\rho_{i,j} = \frac{1}{1+\sigma_e^2/\sigma_w^2}; \quad (5)$$

thus,

$$\rho_{i,j} = \rho_{i,w}^2 \quad (5a)$$

Admittedly, this model of return is somewhat simplified; it implies that all markets have the same correlation with the world market because of the assumption that all markets have betas equal to 1 with the same tracking error. This simple model makes it possible, however, to easily compute and study a useful measure, namely, the average level of global correlation. Furthermore, although we have kept the model simple to make our approach easier to understand, refinements to the model can be introduced.

Note the inverse relationship between global correlation and dispersion. The higher the dispersion, the lower the correlation.

Cross-Sectional Estimation of Correlation. Global dispersion can be used to estimate global correlation based on Equation 3. The time variation in world market volatility cannot be directly measured, so we made the assumption that it is constant over time and equal to its long-run average (a monthly standard deviation of 4.5 percent, or 15.6 percent annually).[5] Hence, we estimated the monthly global correlation by using Equation 3, an estimate of the world market volatility, and our estimate of monthly dispersion. We call the resulting estimate "cross-sectional correlation" (to differentiate it from traditional time-series correlation).

This global correlation for January 1971 to September 1998 is plotted in **Figure 2**. The flat unconditional mean correlation with world returns, 69 percent for monthly returns, was obtained simply by computing the traditional time-series correlation coefficient for each market over the whole period and averaging these coefficients across all markets.[6] This traditional correlation is very close to the long-run average correlation computed by our cross-sectional approach; they differ only in the third decimal.

Remember that the monthly correlation estimates in the cross-sectional approach are independent from one month to the next because they are based on the global dispersion of returns for that month; the estimation does not use historical data, certainly not the past 60 months, as the time-series approach usually does. A stable period of high (or low) monthly correlation is not a spurious statistical artifact caused by overlapping data as in the case of time-series estimations with rolling windows. Clearly, monthly correlation was quite low in the 1972–74 period and the 1985–87 period (until October). In the mid-1990s, there was a period of high correlation, and correlation declined at the end of 1997 and into 1998.

Correlation has apparently increased somewhat in the past 25 years, as shown by the "Fitted Correlation Line" in Figure 2, which is the regression line between correlation and time. This fitted line moves from 66 percent at the start of 1971 to 75 percent in September 1998; the slope has a t-statistic of 5.91, and the R^2 is equal to 9.1 percent.

Differences between the Two Methods

The cross-sectional method of estimating global correlation is dynamic and, using short-term data, gives instantaneous information about change in the level of global correlations. This approach makes it possible to study the time variation in correlation and has the advantage of being immediately available on a monthly or even more frequent basis. The traditional time-series method requires a long period of observations, and overlapping data have to be used to study the change.

A drawback of the dispersion-based methodology is that it requires a sufficiently large number of markets to obtain statistical significance, so it gives information only on movements in the general level of global correlation, not on individual correlations between countries or regions.

The two methods could lead to somewhat different conclusions about the level of global correlations, and the reason should be stressed. The cross-sectional method is conditional on the world return, whereas the time-series method is unconditional, in the sense that for each country, it considers the deviations of national returns from their long-term means.[7] One method looks at relative returns; the other looks at absolute returns.

For illustration, consider this simple intuitive example. Suppose in three periods, national markets exhibit the same dispersion of 4 percent around the world index but the world return in the different periods is –10 percent, +10 percent, and zero. In this case, the "long-term" (three-period) mean return is zero. In the first period, the returns

Figure 2. Global Correlation: 15 Developed Markets

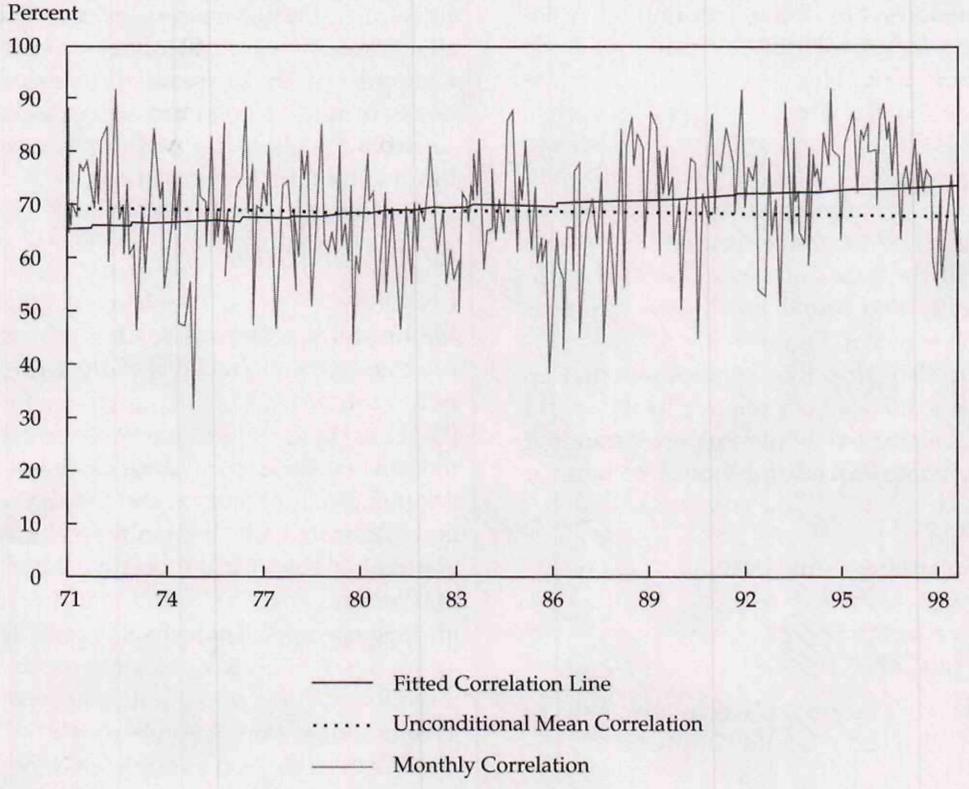

Note: Ending date is September 1998.

on all national markets are negative (centered around –10 percent). In the second period, the returns on all national markets are positive (centered around +10 percent). In the third period, the returns are slightly positive or negative and centered around zero. Traditional time-series correlation estimation would look at deviations of national returns from their long-term means (close to zero) and consider that markets are strongly correlated in Period 1 and Period 2 because all national markets exhibit together a much lower (or higher) return than their mean return. In Period 3, they appear to be uncorrelated. In terms of relative returns, however, the national markets are exhibiting the same dispersion around the world index in all three periods.[8]

We believe our approach is the more appropriate for the current asset management paradigm, in which performance is measured relative to a benchmark rather than in absolute terms. In this reported study, we used the equally weighted world index as a benchmark, but the approach could be extended seamlessly to use a market-cap-weighted world index or any other index as a benchmark.

A "mathematical" comparison of the two approaches is impossible because they are fundamentally different; a link between them depends on the time dynamics postulated for the true correlation and volatility. But a brief empirical comparison of the two indicators of global correlation can be performed.

For each month, we computed independently two measures of global correlation—one derived from our dispersion-based approach and one derived from the traditional time-series approach, taking the average time-series correlation for all markets (calculated over the past 36 months). As might be expected, the dispersion-based correlation was much more volatile than the time-series correlation that used the past 36 months. The correlation over time, however, was equal to 26.7 percent. Furthermore, a change in the dispersion-based correlation helped forecast a change in the time-series correlation.

Extension and Practical Use

The simple model presented here assumed constant world market volatility (in Equation 3). An easy extension would be to use the classic times-series estimate based on historical returns up to t. Another possible extension would be to model time variation in world market volatility by using a simple univariate GARCH(1,1) representation for $\sigma_w(t)$, as practitioners often do. Time variation in such a model depends only on the last innovation (shock) in the world return. So, Equation 3 would still yield a fresh estimate of correlation, based on returns observed at time t.[9] We estimated the cross-sectional correlation with the three different models for world market volatility and found results that are very similar to those given in Figure 1.

We could also drop the assumptions that all markets have a unitary beta relative to the world market. This relaxation would require assuming that betas are stable. Equation 1 would become

$$R_{i,t} = \beta_i R_{w,t} + e_{i,t}. \qquad (1a)$$

Then, the correlations would become

$$\rho_{i,w} = \frac{1}{\sqrt{1 + \sigma_e^2 / \beta_i^2 \sigma_w^2}}. \qquad (4a)$$

Similarly, the correlation between any pair of national markets, i and j, would simply be

$$\rho_{i,j} = \frac{1}{\left(\sqrt{1 + \sigma_e^2 / \beta_i^2 \sigma_w^2}\right)\left(\sqrt{1 + \sigma_e^2 / \beta_j^2 \sigma_w^2}\right)} \qquad (5b)$$

These correlations could be estimated under the assumption that the betas are constant over time but different across countries. Betas estimated over the whole 27-year period are reported in Table 1, with their standard errors and R^2s. The unconditional correlation of each country with the equally weighted world index is given in the last column; the average, as mentioned, is 0.69. The betas are not far from unity, and their low standard errors suggest that they are quite stable over time. The U.S. stock market has the lowest beta (0.687) and a high correlation with the world (0.8).

In practical asset allocation, one could go a step farther and combine the time-series and cross-sectional approaches with a periodic estimation of individual time-series correlations for each country pair and a dynamic updating of the general level of correlation via our dispersion method. This approach would be useful because, for obvious structural, socioeconomic, and sociopolitical reasons, correlations differ among countries. Some regional blocs are closely linked, and some countries have only loose economic ties with each other. For example, the correlation between the German and Dutch stock markets is higher than the correlation between the U.S. and Hong Kong markets. In periods of considerable economic shock, however, when global factors strongly influence all economies, the general level of correlation between all markets increases. In other words, the structural relationships that explain differences in correlations in market pairs are likely to remain quite stable in the short run, but this is not the case for the general level of correlation.

Conclusions

Estimating the cross-sectional dispersion of stock-market returns is an alternative approach to estimating the global level of stock market correlation.

Our objective was to derive an instantaneous measure of the general level of global market correlation. We first defined the concept of dispersion and illustrated its changes from January 1971 to September 1998 for the original MSCI developed market universe. We found that the dispersion, at an average of 4.5 percent a month (15.6 percent annually), was quite stable over the period but had a tendency to decrease with time, which suggests greater integration of equity markets.

We then derived a simple model of global market correlation in a setting in which each country was postulated to have a beta of 1 relative to the world market. We showed that global correlation is inversely proportional to dispersion. On the basis of our model, we showed how to use cross-sectional (or contemporaneous) short-term data to estimate the global market correlation.

We illustrated the usefulness of the model by studying whether developed markets have become increasingly correlated over time. We found that the correlation has had a positive trend, increasing from 66 percent at the beginning of 1971 to 75 percent in September 1998, but that the slope of the regression is quite weak. These findings suggest that equity markets are becoming more integrated but at a slower pace than some practitioners have proposed. This finding is consistent with the findings in Solnik, Boucrelle, and Le Fur (1996), who stressed that the growth of new markets partly offsets the trend toward increasing correlations.

Our cross-sectional method of estimating global correlation is dynamic and, using short-term data, provides instantaneous information on the changes in the level of global correlation. The traditional time-series method requires a long period of observations, and overlapping data have to be used to study the changes in correlations. The methods yield similar estimates, however, for a period of five years or more.

The concept of using dispersion in estimating cross-sectional correlation is important in the current controversy about the relative importance of country allocation and industry allocation in global asset management. In light of all the published empirical evidence based on long-term historical data, structuring asset allocation along industry lines rather than country lines is an act of faith.[10] Such an industry approach makes sense only if the dispersion of returns of global industry indexes is larger than that of country indexes. It would mean that country factors have become less distinctive (country returns have become more homogeneous). Clearly, global asset allocation should be based on factors that maximize global dispersion and thus minimize global correlation.

Notes

1. For example, Solnik, Boucrelle, and Le Fur (1996) looked at the correlation of several markets with the U.S. stock market over the 1958–95 period. They computed the "moving" correlation for all country pairs over the previous 36 months—an estimation window that is quite short for obtaining statistically reliable correlation estimates. They found that the correlation fluctuated over time but showed only a modest increase over the period studied. See also Erb, Harvey, and Viskanta (1994).
2. See Longin and Solnik (1995); Kroner and Ng (1998); Ramchand and Susmel (1998).
3. The MSCI World Index had 22 markets as of the end of 1998. Several countries were later included in the index: for example, Hong Kong and Singapore.
4. To annualize the monthly standard deviation, multiply it by the square root of 12.
5. Other assumptions could easily be incorporated without much influence on the conclusion. The simplest adjustment would be to use the standard time-series estimation based on a rolling window of 60 months. A more sophisticated approach would be to use a GARCH model. A univariate GARCH(1,1) is often used to model time variation in stock market volatility. Most of the variance of the correlation will come from the variance of the dispersion because the correlation in Equation 3 is a function of the dispersion (estimated cross-sectionally, so with a large variance) and the world market volatility (estimated in time series using overlapping data, so with a small variance).
6. The average unconditional monthly correlation of national markets with the world was 0.69 percent, whereas the average pairwise correlation between two national markets was 0.49 (the square of 0.69).
7. By "long term," we mean over the window used to estimate correlation, typically five years of monthly data.
8. To avoid such a problem caused by a couple of extreme observations (small-sample bias), the time-series approach must use a large number of observations so that the empirical distribution of returns will approach the "true" distribution. Unfortunately, the time-series estimation also requires that the distribution remain unchanged over the whole estimation period, but such an assumption is more likely to be violated over long time periods than over short periods. If the three observations used in our simple example were indeed representative of the true distribution of returns (rather than random occurrences in a small sample), the implication would be that the stock markets are extremely volatile, and our cross-sectional method would also find a high correlation because σ_w would be much larger than σ_e in Equation 3.
9. Again note that Equation 3 is consistent with the phenomenon, observed in the literature, that correlation increases when market volatility increases.
10. For a review, see Lombard, Roulet, and Solnik (1999).

References

Aked, M., C. Brightman, and S. Cavaglia. 1999. "On the Increasing Importance of Industry Factors: Implications for Global Portfolio Management." Working paper. Brinson Partners (December).

Erb, C.B., C.R. Harvey, and T.E. Viskanta. 1994. "Forecasting International Correlation." *Financial Analysts Journal*, vol. 50, no. 6 (November/December):32–45.

Kroner, K.F., and V.K. Ng. 1998. "Modeling Asymmetric Co-Movements of Asset Returns." *Review of Financial Studies*, vol. 11, no. 4 (Winter):817–844.

Lombard, T., J. Roulet, and B. Solnik. 1999. "Pricing of Domestic versus Multinational Companies." *Financial Analysts Journal*, vol. 55, no. 2 (March/April):35–49.

Longin, F., and B. Solnik. 1995. "Is the Correlation in International Equity Returns Constant: 1960–1990?" *Journal of International Money and Finance*, vol. 14, no. 1 (February):3–26.

Ramchand, L., and R. Susmel. 1998. "Volatility and Cross Correlation across Major Stock Markets." *Journal of Empirical Finance*, vol. 5, no. 5 (October):397–416.

Solnik, B., C. Boucrelle, and Y. Le Fur. 1996. "International Market Correlation and Volatility." *Financial Analysts Journal*, vol. 52, no. 5 (September/October):17–34.