Mico Loretan
(+1 202) 452 2219
mico.loretan@frb.gov

William B English
(+1 202) 736 5645
william.b.english@frb.gov

III. Special feature: Evaluating changes in correlations during periods of high market volatility[*]

In computing measures of the market risk of a portfolio, such as Value at Risk, portfolio managers typically rely on estimates of correlations between returns on the financial instruments in the portfolio and on the volatility of those returns. This task is relatively simple if the correlations and volatilities do not change over time, and if there are sufficient data to allow them to be estimated fairly precisely. The task is vastly more difficult if the correlations change abruptly as a result of structural breaks in the mechanisms that determine asset returns – perhaps owing to the impact of contagion on the links between markets, changes in the sources of shocks, or new market structures or practices.[12] However, changes in correlation patterns may be no more than the natural and predictable effects of fluctuations in asset return volatility. In such cases, the problem facing risk managers should be less difficult, as the empirical challenge then consists of modelling the time-varying nature of asset return volatilities.

In periods of heightened market volatility, correlations between returns on financial assets tend to increase relative to correlations estimated during periods of normal volatility. For example, the average correlation between yield spreads for selected fixed income securities rose to 0.37 following the Russian crisis in August 1998 from 0.11 in the first half of 1998 (Committee on the Global Financial System (1999), Table A18). The increased correlation of returns during periods of high volatility is often explained as resulting from changes in the underlying relationships that determine returns.[13] Yet, probability theory shows that correlations between asset returns depend on market volatility even if the underlying relationships between returns have not changed; variations in correlations measured over different periods of time may merely be the consequence of variations in realised volatility.

This article explores the link between volatility and correlation, which has until recently largely been overlooked in the economics and finance literature.[14] The next subsection provides two numerical examples that demonstrate the dependence of correlations on volatilities, and also states a theorem that links variances and correlations. An empirical application is presented next, focusing on the behaviour

[*] This article is based on a longer BIS conference research paper (Loretan and English (2000)). The authors are members of the research staff of the Board of Governors of the Federal Reserve System. The analysis and conclusions in this article are those of the authors and do not indicate concurrence by the Board of Governors, by the Federal Reserve System or by the BIS. We thank Jim Clouse, Mike Gibson, Michael Gordy, Brian Madigan, Henri Pagès, Matt Pritsker, Vince Reinhart and participants at the 1999 Central Bank Economists' Autumn Meeting at the Bank for International Settlements for helpful comments and discussions.

[12] Recent discussions of possible routes for contagion include Drazen (1998), Eichengreen et al (1996) and Gerlach and Smets (1995). Kodres and Pritsker (1999) present a structural model of the contagion-like transmission of shocks.

[13] For example, the increase in market volatility of US and other government securities in 1994 was accompanied by an increase in sampling correlations. In its 1995 annual report, Bankers Trust (1995) stated that movements in interest rates in 1994 were "unusual in the degree to which interest rates across international markets moved together" (p 23). The bank went on to note that "this phenomenon of increased correlation among interest rates reduced the risk management benefits derived from diversification across interest-sensitive instruments" (p 23). The bank responded to this situation by withdrawing from substantial market positions (p 24).

[14] For previous economic and finance studies of the link between volatility and correlations, see Ronn (1995), Boyer et al (1999) and Forbes and Rigobon (1999).

of equity returns in the United Kingdom and Germany during the past decade. We find that quarters in which the volatility of equity returns was high also tended to be quarters with above average correlations, in a manner that is consistent with a constant unconditional data generating process for equity returns. The final subsection discusses the implications of the link between volatility and correlation for risk management and for financial supervision.

The link between volatility and correlation

According to probability theory, when the movements of random variables are more volatile, sampling correlations between those variables *should* be elevated even if the underlying process generating the variables remains unchanged. Boyer et al (1999) provide a formal proof of this link (see the box on the next page).

To demonstrate the intuition that underlies this theoretical result, consider a pair of random variables, x and y, and suppose that the possible outcomes for these variables are distributed jointly normally with means equal to zero, variances equal to one and contemporaneous correlation equal to 0.5. A sample of 1,000 independent draws of such pairs is shown in the graph below. The thick and thin ellipses denote the areas that contain 50% and 95% of the total mass of the distribution respectively. Now suppose we split this sample into two subsamples based on the outcome of the variable x. One subsample would be "low volatility" and would include all x and y pairs for which the absolute value of x is less than 1.96.[15] The other subsample would be "high volatility" and would include all pairs for which the absolute value of x is greater than or equal to 1.96. Intuitively, the effect of excluding observations with large values of x should be to reduce the sample correlation between x and y. By contrast, the correlation for the high volatility subsample should be enhanced because one portion of that subsample picks up the large positive values of both variables while the other portion picks up the large negative values. As is noted in the graph below, the difference between the correlations in the two subsamples is large: the correlation for the high volatility sample is 0.81, while that for the low volatility sample is 0.45. Note that the correlation in the latter subsample is close to the population value of 0.5; this result may not be surprising since the low volatility subsample includes 95% of the data.

Bivariate normal random numbers, ρ = 0.5

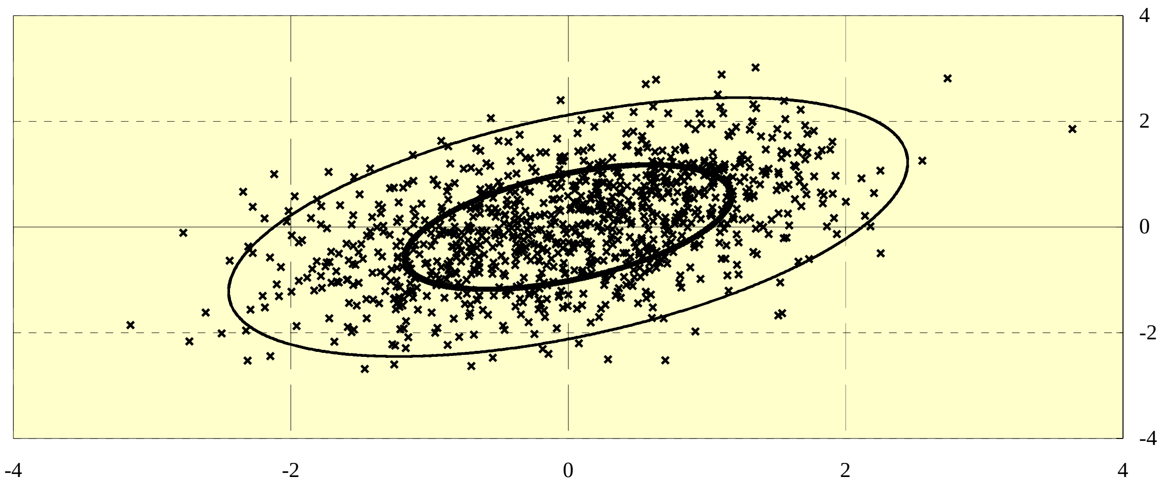

$\rho\,((x,y) \mid |x|<1.96)=0.45$, $\rho\,((x,y) \mid |x| >= 1.96) = 0.81$

[15] The distributions of x and y are standard normal by assumption. Hence, the absolute value of x is less than 1.96 with a probability of 95%.

A formal result

The intuitive link between volatility and correlation can be derived formally. Boyer et al (1999) provide the following theorem.

Theorem. Consider a pair of i.i.d. bivariate normal random variables x and y with standard deviations σ_x and σ_y, respectively, and covariance σ_{xy}. Let ρ $(=\sigma_{xy}/(\sigma_x\sigma_y))$ denote the unconditional correlation between x and y. The correlation between x and y conditional on an event $x \in A$, for any $A \subset \mathbb{R}$ with $0 < \text{Prob}(A) < 1$, is given by

$$\rho_A = \rho\left[\rho^2 + (1-\rho^2)\frac{\sigma_x^2}{\text{Var}(x \mid x \in A)}\right]^{-1/2} \tag{1}$$

Proof. [1] Let u and v be two independent standard normal random variables. Now construct two bivariate normal random variables x and y with means μ_x and μ_y, respectively, standard deviations σ_x and σ_y, respectively, and correlation coefficient ρ:

$$x = \mu_x + \sigma_x u \tag{2}$$

$$y = \mu_y + \rho\sigma_y u + \sqrt{1-\rho^2}\,\sigma_y v \tag{3}$$

Consider an event $x \in A$, for any $A \subset \mathbb{R}$ with $0 < \text{Prob}(A) < 1$. By definition, the conditional correlation coefficient between x and y, ρ_A, is given by

$$\rho_A = \frac{\text{Cov}(x, y \mid x \in A)}{\sqrt{\text{Var}(x \mid x \in A)}\sqrt{\text{Var}(y \mid x \in A)}} \tag{4}$$

By substituting for u in (3) using equation (2), then substituting the resulting expression for y into (4), and using the fact that x and v are independent by construction, one can rewrite this as

$$\rho_A = \frac{(\rho\sigma_y/\sigma_x)\,\text{Var}(x \mid x \in A)}{\sqrt{\text{Var}(x \mid x \in A)}\sqrt{(\rho^2\sigma_y^2/\sigma_x^2)\,\text{Var}(x \mid x \in A) + (1-\rho^2)\sigma_y^2}} \tag{5}$$

which can, in turn, be simplified to yield the expression in (1).

Thus, the conditional correlation between x and y is larger (smaller) than ρ in absolute value if the conditional variance of x given $x \in A$ is larger (smaller) than the unconditional variance of x.

[1] This proof is based on the property of bivariate normal random variables that each component can be expressed as the weighted average of the other and of an independent variable that is also normally distributed. See, for example, Goldberger (1991, p 75).

The theoretical link between volatility and correlation holds in a time series context as well. Consider subdividing a long time series of two variables, *x* and *y*, which are observed daily, into quarterly subsamples. For each subsample, calculate the variance of *x* and the correlation between *x* and *y*. Finally, order the subsamples by the variance of *x*. The table on the next page shows the results of such an exercise under the assumption that *x* and *y* are independent and normally distributed, with unit variances, and a constant correlation coefficient equal to 0.5 (as in the graph on the previous page). The first column of the table shows ranges for the ratio of the quarterly sampling variance in *x* to its population value (which is 1). The other three columns show the distribution of quarterly correlation values for the samples in those ranges. For quarters with in-sample variance of *x* close to its population value (0.9 to 1.1), the median sampling correlation is 0.50. However, the distribution of sampling

In-sample correlations when conditioning on volatility

Range of variances of *x* relative to its population variance	Conditional correlation of *x* and *y*		
	Bottom 5%	**Median**	**Top 5%**
0.3–0.5	0.17	0.36	0.53
0.5–0.7	0.24	0.43	0.58
0.7–0.9	0.29	0.47	0.61
0.9–1.1	0.34	0.50	0.64
1.1–1.3	0.38	0.54	0.67
1.3–1.5	0.41	0.57	0.69
1.5–1.7	0.45	0.59	0.71
1.7–1.9	0.48	0.61	0.72

Note: The random variables x and y are i.i.d. bivariate normal with a population correlation coefficient ρ of 0.5. Reported values for Corr(x,y) are based on 2.5 million random draws of "quarters" (consisting of 60 "daily" data pairs). There were too few observations with a variance of x less than 0.3 or greater than 1.9 times its population value for values of Corr(x,y) to be reported with confidence.

correlations is fairly wide, with a 90% confidence interval running from 0.34 to 0.64. In contrast, for quarters with in-sample variance of x between 1.7 and 1.9 times its population value, the median correlation is 0.61, with the 90% confidence interval running from 0.48 to 0.72. In other words, in this time series example, periods of increased sampling volatility are also periods of relatively high measured correlations, even when the population correlation remains constant.

An empirical application

In order to assess the real-world applicability of this theoretical link between volatility and correlation, we need to consider whether it can explain the historical relationship between pairs of asset returns. Are contemporaneous changes in sampling variances and sampling correlations empirically consistent with an unchanged underlying distribution of asset returns, and, in particular, with a constant population correlation?

We consider stock prices as measured by the FTSE and Dax stock price indices.[16] These data series represent large and liquid markets and reflect market conditions at roughly the same time, and so we do not have to be concerned about the implications of non-synchronous data collection.[17] Our data are daily observations from the beginning of 1991 to the middle of 1999. The returns are calculated as daily percentage changes in the respective price indices.

The graph on the next page shows time series plots of the within-quarter variances (left-hand panel) and correlations (right-hand panel) of the daily stock market returns. It is clear that autumn 1998 was a period of high volatility and, just as the theoretical results would suggest, one of elevated correlation.

To evaluate the importance of the theoretical link between volatility and correlation more generally, we show in the graph on page 34 a scatterplot of the quarterly in-sample correlations against the in-

[16] In Loretan and English (2000) we present results for returns on government bonds and foreign exchange as well. See also Forbes and Rigobon (1999) for a detailed examination of the link between volatility and correlation in equity prices. The FTSE and Dax data are from Bloomberg, and reflect closing quotes.

[17] For a discussion of the problems associated with non-synchronous data collection, see RiskMetrics (1996, pp 184-196).

sample volatility of the return on the Dax (the lines in the graph are discussed below).[18] The graph clearly shows a generally increasing relationship between the sample variances and sample correlations; the observations for the final two quarters of 1998 comprise two of the three observations at the top right. Although the upward slope in the graph on the next page is consistent with theoretical expectations, the data also show a considerable dispersion in the sample correlation for a given level of sample volatility. In order to provide a more compelling test of whether the population correlation is constant, we need to determine whether the empirical relationship lies mostly within a confidence band around the expected average relationship between volatility and correlation, where the expected relationship and the confidence band are based on the assumption of a constant distribution of the asset returns. One way to construct the theoretical expectations and confidence band is to use a bootstrap, which is based on repeatedly drawing observations from the actual data. Specifically, we select a random sample of a quarter's worth of observations (60 pairs of returns) from the observed data series and calculate the sample variances of the two returns and the sample correlation between the return series. We then repeat the process a large number of times (2 million random samples in total), thereby producing a very large number of correlation-variance pairs. We then use these random observations to calculate the median value of the correlation as a function of the volatility as well as 90% confidence intervals around that median.[19] The resulting lines are plotted in the graph on the next page.

Within-quarter variances and correlations, Dax and FTSE indices

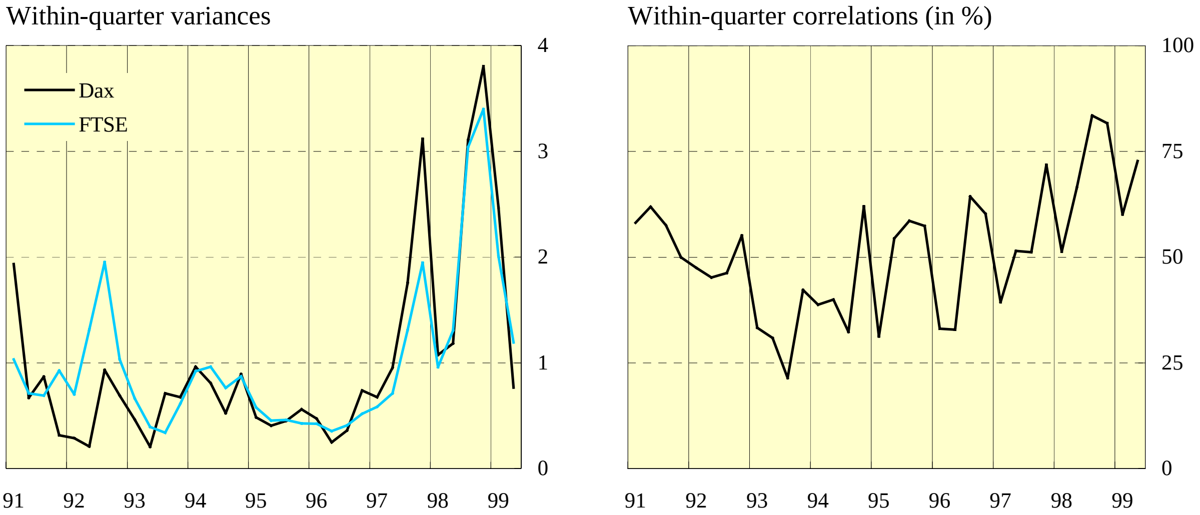

[18] The correlations could be plotted relative to the volatility of either return; use of the FTSE index yields similar results. Note that the within-quarter variance of the return on the Dax has been expressed relative to its full-sample value.

[19] Note that the median correlation and the confidence intervals are based on the actual distributions of the data series rather than on an assumed distribution, such as the bivariate normal. Our earlier study, Loretan and English (2000), shows both the bootstrap results and those based on a bivariate normal distribution. The median values are similar, but the confidence contours are wider under the bootstrap; this appears to be due to the fact that the actual returns have more outlier observations than would be implied by a normal distribution. The bootstrap procedure preserves the unconditionally heavy-tailed nature of the distributions as well as the contemporaneous correlation structure of the data. However, it does not take account of serial dependence features such as GARCH, which, as discussed in Loretan and English (2000), appear to be present in the data.

Quarterly variances vs quarterly correlation, bootstrapped confidence bands

In percentage points[1]

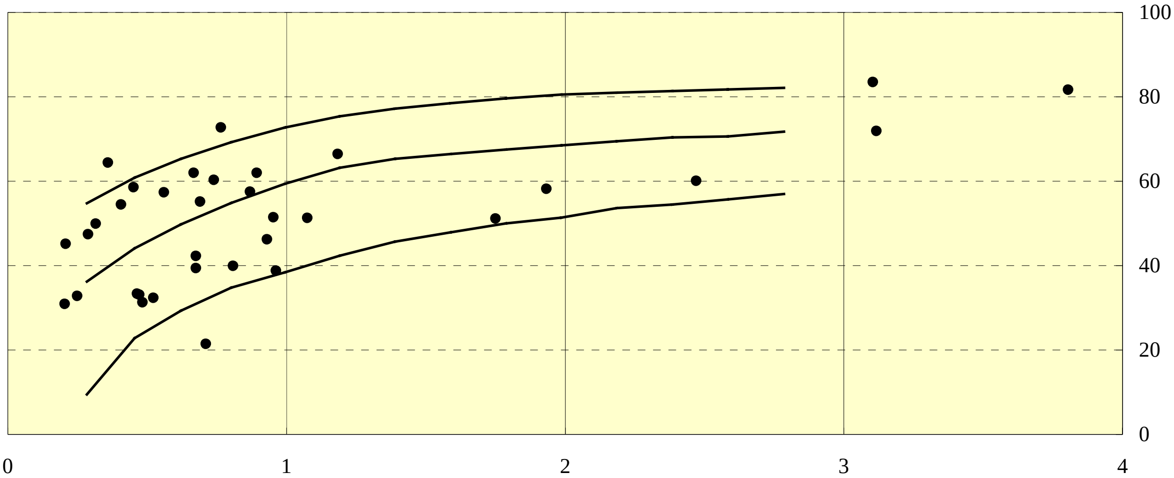

[1] The vertical axis shows the quarterly correlation of the Dax vs the FTSE. The horizontal axis shows the relative quarterly variances of Dax returns.

The equity data fit the pattern implied by the simple theory surprisingly well.[20] The observations are scattered fairly evenly around the median line, and only a few of the 34 observations lie outside the 90% confidence contours. While a more comprehensive test is beyond the scope of this article, our results suggest that one should not be too quick to conclude that fluctuations in correlations during periods of market volatility, including those observed in the second half of 1998, represent true changes in the distribution of asset returns. Rather, they may be nothing more than the predictable consequences of observing certain (low probability) draws from an unchanged distribution. This conclusion need not imply that "contagion" does not occur: rather, it suggests that if one defines contagion to mean elevated sample correlations between asset returns, then contagion can be a natural by-product of high sampling volatilities.

Implications

The statistical link between sampling volatilities and correlations of asset returns has important implications for the evaluation of portfolio risk by market participants and investors as well as for the supervision of financial firms' risk management practices.

Risk managers sometimes use data from a relatively short interval when calculating correlations and volatilities for use in risk management models. Some estimation methods are based on longer intervals of data, but they apply geometrically declining weights, thereby reducing the effective number of observations employed. The theoretical and empirical results presented here suggest that the use of relatively short intervals of data for estimating correlations and volatilities may be dangerous. If the interval happens to be atypically stable, then not only may the estimated volatilities be too low, but, perhaps more important, the estimated correlations between returns will be lower than average. As a result, assessments of market risk may overstate the amount of diversification in a portfolio, leading the investing firm to take on excessive risk. Conversely, if the interval of data employed is a relatively

[20] Our results are based on the volatility of asset returns with no distinction made between increases and decreases in asset prices. In a related study, Longin and Solnik (1998) find that measured correlations between equity returns in different countries behave as the theory would suggest when there are large positive stock market returns but are higher than the theory would suggest when there are large negative returns. We leave an examination of this issue for future research.

volatile one, then the resulting estimates of correlations will be atypically high and could lead the firm to take positions that are excessively risk-averse.

This does not necessarily imply that the use of longer time series produces more reliable calculations. Indeed, short intervals have some desirable features. Since financial markets can change over time, one may not want to depend on data from the distant past.[21] Moreover, the emphasis on recent data allows account to be taken of time-varying volatility, which appears to be a feature of actual returns. However, our results suggest that when determining the appropriate time interval to use, risk managers should not exclude periods of relatively high or low volatility. Such periods contain important information about the underlying relationship between asset returns.

Another way in which the link between in-sample volatility and correlation could cause problems for risk managers is in the calculation of worst case scenarios and in stress testing. Put simply, risk managers should not consider the possible effects of high return volatilities without also taking into account the higher correlations between asset returns that would generally accompany the elevated volatility (see Ronn (1995) for a related discussion). One way to do so would be to employ information from historical periods of high volatility in order to form estimates of correlations conditional on being in a period of heightened volatility.[22] These conditional correlations could then be used to evaluate the distribution of returns under a high volatility scenario. Put differently, the method used for stress testing a portfolio must not (inadvertently) exclude the empirical feature that periods of high volatility are also likely to be periods of elevated correlation.

Supervisors of financial institutions also need to be aware of the link between volatilities and correlations when assessing firms' risk management practices. For example, in evaluating such firms' internal models, supervisors need to keep in mind the difficulties noted earlier with relying on a relatively short interval of data for information on correlations and the need to form appropriate conditional correlations for stress tests.

References

Bankers Trust (1995): *Annual Report 1994*.

Boyer, Brian H, Michael S Gibson and Mico Loretan (1999): "Pitfalls in Tests for Changes in Correlations", *International Finance Discussion Paper No. 597R*, Federal Reserve Board, March.

Chase Manhattan Corp (1999): *Annual Report 1998*.

Committee on the Global Financial System (1999): *A Review of Financial Market Events in Autumn 1998*, Bank for International Settlements, October.

Drazen, Allan (1998): "Political Contagion in Currency Crises", mimeo, University of Maryland.

Eichengreen, Barry, Andrew Rose and C Wyplosz (1996): "Contagious Currency Crises", *NBER Working Paper No. 5681*, July.

Forbes, Kristin and Roberto Rigobon (1999): "No Contagion, Only Interdependence: Measuring stock market co-movements", *NBER Working Paper No. 7267,* July.

Gerlach, Stefan and Frank Smets (1995): "Contagious Speculative Attacks", *European Journal of Political Economy*.

---

[21] Similarly, if the assets under consideration are firm-specific (rather than indices), the behaviour of firms can change over time as managers or business strategies are changed, making older information less useful.

[22] Alternatively, firms might want to use actual data from earlier periods of high volatility to stress test their portfolios. For example, Chase Manhattan uses asset price movements during three historical episodes - the bond market sell-off in 1994, the 1994 Mexican peso crisis and the 1997 Asian markets crisis - as well as internally developed scenarios, when assessing the risk of its portfolio (Chase Manhattan (1999, p 37)).

Goldberger, Arthur S (1991): *A Course in Econometrics*, Harvard University Press.

Kodres, Laura and Matthew G Pritsker (1999): "A Rational Expectations Model of Financial Contagion", mimeo, Federal Reserve Board, May.

Longin, François and Bruno Solnik (1998): "Correlation Structure of International Equity Markets During Extremely Volatile Periods", mimeo, Group HEC, April.

Loretan, Mico and William B English (2000): "Evaluating 'Correlation Breakdowns' during Periods of Market Volatility", *BIS Conference Papers, Vol 8,* Bank for International Settlements, March.

RiskMetrics (1996): *RiskMetrics - Technical Document*, 4th Ed, JP Morgan/Reuters.

Ronn, Ehud I (1995): "The Impact of Large Changes in Asset Prices on Intra-Market Correlations in the Stock and Bond Markets", mimeo, Department of Finance, University of Texas at Austin, July.