1
Correlation Economics

1.1 Introduction

Today there are almost three thousand stocks listed on the New York Stock Exchange. NASDAQ lists another three thousand. There is yet another collection of stocks that are unlisted and traded on the Bulletin Board or Pink Sheets. These U.S.-traded stocks are joined by thousands of companies listed on foreign stock exchanges to make up a universe of publicly traded equities. Added to these are the enormous number of government and corporate and municipal bonds that are traded in the United States and around the world, as well as many short-term securities. Investors are now exploring a growing number of alternative asset classes each with its own large set of individual securities. On top of these underlying assets is a web of derivative contracts. It is truly a vast financial arena. A portfolio manager faces a staggering task in selecting investments.

The prices of all of these assets are constantly changing in response to news and in anticipation of future performance. Every day many stocks rise in value and many decline. The movements in price are, however, not independent. If they were independent, then it would be possible to form a portfolio with negligible volatility. Clearly this is not the case. The correlation structure across assets is a key feature of the portfolio choice problem because it is instrumental in determining the risk. Recognizing that the economy is an interconnected set of economic agents, sometimes considered a general equilibrium system, it is hardly surprising that movements in asset prices are correlated. Estimating the correlation structure of thousands of assets and using this to select superior portfolios is a Herculean task. It is especially difficult when it is recognized that these correlations vary over time, so that a forward-looking correlation estimator is needed. This problem is the focus of this book. We must "anticipate correlations" if we want to have optimal risk management, portfolio selection, and hedging.

Such forward-looking correlations are very important in risk management because the risk of a portfolio depends not on what the correlations were in the past, but on what they will be in the future. Similarly, portfolio choice depends on forecasts of asset dependence structure. Many aspects of financial planning involve hedging one asset with a collection of others. The optimal hedge will also depend upon the correlations and volatilities to be expected over the future holding period. An even more complex problem arises when it is recognized that the correlations can be forecast many periods into the future. Consequently, there are predictable changes in the risk–return trade-off that can be incorporated into optimal portfolios.

Derivatives such as options are now routinely traded not only on individual securities, but also on baskets and indices. Such derivative prices are related to the derivative prices of the component assets, but the relation depends on the correlations expected to prevail over the life of the derivative. A market for correlation swaps has recently developed that allows traders to take a position in the average correlation over a time interval. Structured products form a very large class of derivatives that are sensitive to correlations. An important example of a structured product is the collateralized debt obligation (CDO), which in its simplest form is a portfolio of corporate bonds that is sold to investors in tranches that have different risk characteristics. In this way credit risks can be bought and sold to achieve specific risk–return targets. There are many types of CDOs backed by loans, mortgages, subprime mortgages, credit default swaps, tranches of CDOs themselves, and many other assets. In these securities, the correlations between defaults are the key determinants of valuations. Because of the complexity of these structures and the difficulty in forecasting correlations and default correlations, it has been difficult to assess the risks of the tranches that are supposed to be low risk. Some of the "credit crunch" of 2007–8 can probably be attributed to this failure in risk management. This episode serves to reinforce the importance of anticipating correlations.

This book will introduce and carefully explain a collection of new methods for estimating and forecasting correlations for large systems of assets. The book initially discusses the economics of correlations. Then it turns to the measurement of comovement and dependence by correlations and alternative measures. A look at existing models for estimating correlations—such as historical correlation, exponential smoothing, and multivariate GARCH—leads to the introduction (in chapter 3) of the central method explored in the book: dynamic conditional correlation. Monte Carlo and empirical analyses of this model document its performance. Successive chapters deal with extensions to the basic model, new

Table 1.1. Correlations of large-cap stocks from 1998 to 2003.

	IBM	MCD	GE	Citibank	AXP	WMT	SP500
IBM	1.000	0.192	0.436	0.419	0.387	0.283	0.600
MCD	0.192	1.000	0.308	0.238	0.282	0.303	0.365
GE	0.436	0.308	1.000	0.595	0.614	0.484	0.760
Citibank	0.419	0.238	0.595	1.000	0.697	0.439	0.740
AXP	0.387	0.282	0.614	0.697	1.000	0.445	0.715
WMT	0.283	0.303	0.484	0.439	0.445	1.000	0.584
SP500	0.600	0.365	0.760	0.740	0.715	0.584	1.000

estimation methods, and a technical discussion of some econometric issues. Many empirical studies are documented in particular chapters, including stock–bond correlations, global equity correlations, and U.S. large-cap stock correlations. Finally, in a chapter called "Anticipating Correlations," these methods are used to forecast correlations through the turbulent environment of the summer and autumn of 2007.

The methods introduced in this book are simple, powerful, and will be shown to be highly stable over time. They offer investors and money managers up-to-date measures of volatilities and correlations that can be used to assess risk and optimize investment decisions even in the complex and high-dimensional world we inhabit.

1.2 How Big Are Correlations?

Correlations must all lie between −1 and 1, but the actual size varies dramatically across assets and over time. For example, using daily data for the six-year period from 1998 through 2003 and the textbook formula

$$\hat{\rho}_{x,y} = \frac{\sum_{t=1}^{T}(x_t - \bar{x})(y_t - \bar{y})}{\sqrt{\sum_{t=1}^{T}(x_t - \bar{x})\sum_{t=1}^{T}(y_t - \bar{y})}}, \tag{1.1}$$

it is interesting to calculate a variety of correlations. The correlation between daily returns on IBM stock and the S&P 500 measure of the broad U.S. market is 0.6. This means that the regression of IBM returns on a constant and S&P returns will have an R^2 value of 0.36. The systematic risk of IBM is 36% of the total variance and the idiosyncratic risk is 64%.

Looking across five large-capitalization stocks, the correlations with the S&P 500 for the six-year period range from 0.36 for McDonald's to 0.76 for General Electric (GE). These stocks are naturally correlated with each other as well, although the correlations are typically smaller (see table 1.1).

Table 1.2. Correlations of small-cap stocks from 1998 to 2003.

	PVA	NSC	ARG	DRTK	MTLG	SP
PVA	1.000	0.159	0.050	0.063	0.014	0.185
NSC	0.159	1.000	0.253	0.006	0.034	0.445
ARG	0.050	0.253	1.000	0.068	0.081	0.326
DRTK	0.063	0.006	0.068	1.000	0.025	0.101
MTLG	0.014	0.034	0.081	0.025	1.000	0.080
SP	0.185	0.445	0.326	0.101	0.080	1.000

A more careful examination of the correlations shows that the highest correlations are between stocks in the same industry. American Express (AXP) and Citibank have a correlation of almost 0.7 and GE has a correlation with both that is about 0.6. During this period GE had a big financial services business and therefore moved closely with banking stocks.

Examining a selection of small-cap stocks, the story is rather different. The correlations with the market factor are much lower and the correlations between stocks are lower; table 1.2 gives the results. The largest correlation with the market is 0.45 but most of the entries in the table are below 0.1.

Turning to other asset classes let us now examine the correlation between the returns on holding bonds and the returns on holding foreign currencies (see table 1.3). Notice first the low correlations between bond returns and the S&P 500 and between currency returns and the S&P 500. These asset classes are not highly correlated with each other on average.

Within asset classes, the correlations are higher. In fact the correlation between the five- and twenty-year bond returns is 0.875, which is the highest we have yet seen. The short rate has correlations of 0.3 and 0.2, respectively, with these two long rates. Within currencies, the highest correlation is 45% between the Canadian dollar and the Australian dollar, both relative to the U.S. dollar. The rest range from 15% to 25%.

When calculating correlations across countries, it is important to recognize the differences in trading times. When markets are not open at the same times, daily returns calculated from closing data can be influenced by news that appears to be on one day in one market but on the next day in the other. For example, news during U.S. business hours will influence measured Japanese equity prices only on the next day. The effect of the news that occurs when a market is closed will be seen primarily in the opening price and therefore is attributed to the following daily return. To mitigate this problem, it is common to use data that is more time aggregated to measure such correlations.

Table 1.3. Other assets.

	T3M	T5YR	T20YR	CAD/USD	GBP/USD	AUD/USD	JPY/USD	SP500
T3M	1.000	0.329	0.206	0.011	0.076	0.025	0.031	−0.031
T5YR	0.329	1.000	0.875	−0.0007	0.136	0.007	0.005	−0.057
T20YR	0.206	0.875	1.000	0.007	0.103	−0.002	−0.049	−0.016
CAD/USD	0.011	−0.0007	0.007	1.000	0.117	0.415	0.145	0.015
GBP/USD	0.076	0.136	0.103	0.117	1.000	0.253	0.224	−0.018
AUD/USD	0.025	0.007	−0.002	0.415	0.253	1.000	0.269	0.040
JPY/USD	0.031	0.005	−0.049	0.145	0.224	0.269	1.000	−0.003
SP500	−0.031	−0.057	−0.016	0.015	−0.018	0.040	−0.003	1.000

Notes: "T3M" denotes three-month Treasury Bill returns; "T5YR" denotes five-year Treasury bond returns; "T20YR" denotes twenty-year Treasury bond returns; "CAD/USD" Canadian dollar/U.S. dollar returns; "GBP/USD" U.K. pound/U.S. dollar returns; "AUD/USD" Australian dollar/U.S. dollar returns; "JPY/USD" Japanese yen/U.S. dollar returns; "SP500" denotes Standard & Poor's 500 index of equity returns.

Cappiello et al. (2007) analyze weekly global equity and bond correlations. The data employed in their paper consist of FTSE All-World indices for twenty-one countries and DataStream-constructed five-year average maturity bond indices for thirteen, all measured relative to U.S. dollars. The sample is fifteen years of weekly price observations, for a total of 785 observations from January 8, 1987, until February 7, 2002. Table 1.4 shows a sample of global equity and bond correlations. The bond correlations are above the diagonal and the equity correlations are below the diagonal.

The equity correlations range from 0.23 to 0.73 with about a third of the sample above 0.5. The highest are between closely connected economies such as Germany, France, and Switzerland, and the United States and Canada. The bond return correlations are often much higher. France and Germany have a correlation of 0.93 and most of the European correlations are above 0.6. The U.S. correlation with Canada is 0.45, while the correlations with other countries hover around 0.2. Japanese correlations are also lower. Cappiello et al. also report correlations between equities and bonds that vary greatly. Many of these are negative. Typically, however, the domestic equity- and bond-return correlations are fairly large. This is partly due to the fact that both returns are denominated in U.S. dollars.

1.3 The Economics of Correlations

To understand the relative magnitude of all these correlations and ultimately why they change, it is important to look at the economics behind movements in asset prices. Since assets are held by investors in anticipation of payments to be made in the future, the value of an asset is intrinsically linked to forecasts of the future prospects of the project or firm. Changes in asset prices reflect changing forecasts of future payments. The information that makes us change these forecasts we often simply call "news." This has been the basic model for changing asset prices since it was formalized by Samuelson (1965). Thus both the volatilities of asset returns and the correlations between asset returns depend on information that is used to update these distributions.

Every piece of news affects all asset prices to a greater or lesser extent. The effects are greater on some equity prices than on others because their lines of business are different. Hence the correlations in their returns due to this news event will depend upon their business. Naturally, if a firm changes its line of business, its correlations with other firms are likely to change. This is one of the most important reasons why correlations change over time.

Table 1.4. Global equity and bond correlations.

	Canada	Denmark	France	Germany	Ireland	Japan	Sweden	Switzerland	U.K.	U.S.
Canada	—	0.094	0.097	0.068	0.134	0.007	0.160	−0.019	0.167	0.452
Denmark	0.279	—	0.907	0.909	0.839	0.418	0.696	0.800	0.650	0.195
France	0.462	0.496	—	0.927	0.832	0.428	0.679	0.823	0.673	0.267
Germany	0.399	0.399	0.729	—	0.826	0.464	0.657	0.866	0.656	0.221
Ireland	0.361	0.449	0.486	0.515	—	0.359	0.664	0.710	0.699	0.212
Japan	0.230	0.288	0.340	0.321	0.279	—	0.294	0.475	0.343	0.038
Sweden	0.497	0.478	0.577	0.639	0.474	0.317	—	0.553	0.566	0.173
Switzerland	0.391	0.521	0.641	0.722	0.528	0.343	0.549	—	0.589	0.126
U.K.	0.463	0.460	0.594	0.562	0.634	0.350	0.539	0.585	—	0.249
U.S.	0.692	0.299	0.465	0.432	0.392	0.223	0.490	0.398	0.495	—

Notes: equity correlations appear above the diagonal and bond correlations appear below; the figures are for the period 1987–2002.

A second important reason is that the characteristics of the news change. News that has the same qualitative effect on two companies will generally increase their correlation. The magnitude of this news event will determine whether this is an important change in correlations. Consequently, correlations often change dramatically when some factor becomes very important having previously been dormant. An example of this might be energy prices. For years, these fluctuated very little. However, in 2004 prices more than doubled and suddenly many firms and countries whose profitability depended on energy prices showed fluctuations in returns that were more correlated than before (some of these are naturally negative). Thus when the news changes in magnitude it is natural that correlations will change.

Since asset prices of firms are based on the forecasts of earnings or dividends and of expected returns, the movements in prices are based on the updates to these forecasts, which we call firm news. For each asset, we will focus on two types of news: news on future dividends or earnings, and news on future expected returns. Both types of news will depend upon news about energy prices, wage rates, monetary policy, and so forth. Correlations are then based on the similarities between the news for different firms. In particular, it will be shown below that it is correlation between the firm news processes that drives correlation between returns.

To apply this idea to the correlations described in tables 1.1–1.4, it is necessary to show how the underlying firm news processes are correlated. Stocks in the same industry will have highly correlated dividend news and will therefore be more highly correlated than stocks in different industries. Small-cap stocks will often move dramatically with earnings news and this news may have important idiosyncratic components. Consequently, these stocks are naturally less correlated than large-cap stocks. Large-cap stocks will have rather predictable dividend streams, which may respond directly to macroeconomic news. These companies often have well-diversified business models. Hence, volatilities of large-cap stocks should be less than those for small-cap stocks and correlations should be higher. Index returns will also respond to macroeconomic news and hence are typically more correlated with large-cap stocks than with small-cap ones.

For equities, news about the expected return is essentially news about the relevant interest rate for this asset. It will be determined largely by shifts in macroeconomic policy, which determine short rates, and by the risk premium, which in turn will be influenced by market volatility. These effects are presumably highly correlated across stocks within the domestic market. There may be fluctuations across sectors and companies as the actual risk premium could vary with news, but one would expect that this factor would be quite correlated.

The net effect of these two news sources for equities will be a return correlation constructed from each of the basic correlations. The bigger the size of a news event, the more important its influence on correlations will be. Thus when future Federal Reserve policy is uncertain, every bit of news will move prices and the correlations will rise to look more like the correlation in required returns. When the macroeconomy is stable and interest rates have low volatility, the correlation of earnings news is most important. For government bonds there is little or no uncertainty about dividends, but news about the future short-term interest rate is a key determinant of returns. Bonds of all maturities will respond to news on monetary policy or short-term interest rate changes. When this is the major news source, the correlations will be quite high. When there are changes in risk premiums, it will again affect all fixed-income securities, leading to higher correlations. However, when the premium is a credit risk premium, the effect will be different for defaultable securities such as corporate bonds or bonds with particularly high yields. In this case, correlations might fall or even go negative between high-risk and low-risk bonds. Because equities as well as bonds are sensitive to the expected-return component of news, they will be positively correlated when this has high variance. When it has low variance, we might expect to see lower or negative correlations between stocks and treasuries, particularly if good news on the macroeconomy becomes bad news on interest rates because of countercyclical monetary policy.

Exchange rates respond to both domestic and foreign news. If all exchange rates are measured relative to the dollar, there is a natural common component in the correlations. Similarly, international equity and bond returns may be measured in dollar terms. This will increase the measured correlations. Countries with similar economies will have correlated news processes because the same events will affect them. Index returns such as those exhibited in table 1.4 will show more highly correlated returns as the idiosyncratic shocks will be averaged away. Bond returns across countries will generally be highly correlated as the market is truly global; the currency of denomination may be important in flexible exchange rate systems.

1.4 An Economic Model of Correlations

Many of these results are complex and interrelated. Because they have been described in words, the overall simplicity of the argument may not be apparent. In order to put these results in a quantitative context, a mathematical derivation of the correlation of returns needs to be developed. The goal is to show how correlations in returns are based on correlations in the news. We first express continuously compounded returns, r, in terms of the price per share, P, and the dividend per share, D:

$$r_{t+1} = \log(P_{t+1} + D_{t+1}) - \log(P_t). \tag{1.2}$$

Applying the Campbell and Shiller (1988a,b) or Campbell (1991) log-linearization, this can be approximately written as

$$r_{t+1} \approx k + \rho p_{t+1} + (1-\rho) d_{t+1} - p_t, \tag{1.3}$$

where lowercase letters refer to logs and k is a constant of linearization. Essentially this is a type of series expansion of the log of a sum, which is approximately the weighted average of the logs of the components. The approximation is good if the ratio of the two components is small and relatively constant. These conditions are typically satisfied for equity prices. The parameter ρ is essentially the discount rate and is very slightly below 1.

Solving this equation for p_t, assuming that stock prices do not diverge to infinity, gives

$$p_t = \frac{k}{1-\rho} + (1-\rho) \sum_{j=0}^{\infty} \rho^j d_{t+1+j} - \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{1.4}$$

Taking expectations of both sides with respect to the information at time t gives the same dependent variable, since p_t is in the information set:

$$p_t = \frac{k}{1-\rho} + (1-\rho) \sum_{j=0}^{\infty} \rho^j E_t(d_{t+1+j}) - \sum_{j=0}^{\infty} \rho^j E_t(r_{t+1+j}). \tag{1.5}$$

Similarly, taking expectations with respect to information at time $t-1$ gives the one-step-ahead predictor of prices. The difference between the log of the price expected today and that expected at $t-1$ is simply the surprise in returns. Hence,

$$r_t - E_{t-1}(r_t) = E_t(p_t) - E_{t-1}(p_t) \tag{1.6}$$

and

$$r_t - E_{t-1}(r_t) = (1-\rho) \sum_{j=0}^{\infty} \rho^j (E_t - E_{t-1})(d_{t+1+j}) - \sum_{j=0}^{\infty} \rho^j (E_t - E_{t-1})(r_{t+1+j}). \tag{1.7}$$

The unexpected returns have two components: surprises in future dividends and surprises in future expected returns. Often it is convenient to summarize this expression by the relation

$$r_t - E_{t-1} r_t = \eta_t^{\mathrm{d}} - \eta_t^{\mathrm{r}}. \tag{1.8}$$

These two innovations comprise the news; it is the new information that is used to forecast the discounted future dividends and expected returns.

Each of these innovations is the shock to a weighted average of future dividends or expected returns. Consequently, each is a martingale difference sequence. From (1.7) it is clear that even a small piece of information observed during period t could have a large effect on stock prices if it affects expected dividends for many periods into the future. But it could have a relatively small effect if it only affects dividends for a short period. In the simplest finance world, expected returns are constant, so the second term is zero. However, if there is some predictability in expected returns, either because risk premiums are predictable or because the risk-free rate is changing in a predictable way or because markets are not fully efficient, then the second term may be very important.

The conditional variance of an asset return is simply given from (1.8) as

$$V_{t-1}(r_t) = V_{t-1}(\eta_t^{\mathrm{d}}) + V_{t-1}(\eta_t^{\mathrm{r}}) - 2\,\mathrm{Cov}(\eta_t^{\mathrm{d}}, \eta_t^{\mathrm{r}}). \tag{1.9}$$

Each term measures the importance of today's news in forecasting the present value of future dividends or expected returns. If d is an infinite-order moving average, possibly with weights that do not converge (like a unit root),

$$d_t = \sum_{i=1}^{\infty} \theta_i \varepsilon_{t-i}^{\mathrm{d}}, \tag{1.10}$$

then

$$\eta_t^{\mathrm{d}} = \varepsilon_t^{\mathrm{d}}(1-\rho)\sum_{j=0}^{\infty} \rho^j \theta_{j+1} \tag{1.11}$$

and

$$V_{t-1}(\eta_t^{\mathrm{d}}) = V_{t-1}(\varepsilon_t^{\mathrm{d}})\left[\sum_{j=0}^{\infty} \theta_{j+1}\rho^j(1-\rho)\right]^2. \tag{1.12}$$

In this model, time variation arises only from substituting volatility in the innovation for dividends. If there is no predictability in expected returns, then this is also the conditional variance of returns. The longer the memory of the dividend process, the more important this effect is and the greater the volatility is. Of course if the dividend process has time variation in the moving-average coefficients, this would be another time-varying component.

The conditional covariance between two asset returns can be expressed in exactly the same terms:

$$\begin{aligned}\mathrm{Cov}_{t-1}(r_t^1, r_t^2) = {} & \mathrm{Cov}_{t-1}(\eta_t^{\mathrm{d1}}, \eta_t^{\mathrm{d2}}) + \mathrm{Cov}_{t-1}(\eta_t^{\mathrm{r1}}, \eta_t^{\mathrm{r2}}) \\ & - \mathrm{Cov}_{t-1}(\eta_t^{\mathrm{d1}}, \eta_t^{\mathrm{r2}}) - \mathrm{Cov}_{t-1}(\eta_t^{\mathrm{r1}}, \eta_t^{\mathrm{d2}}).\end{aligned} \tag{1.13}$$

In the simple case where expected returns are constant and dividends are fixed-weight moving averages, as in (1.10), and denoting the parameters for each asset as (ρ^1, θ^1) and (ρ^2, θ^2) respectively,

$$\begin{aligned}\operatorname{Cov}_{t-1}(r_t^1, r_t^2) \\ = \operatorname{Cov}_{t-1}(\varepsilon_t^{\mathrm{d}1}, \varepsilon_t^{\mathrm{d}2})(1-\rho^1)(1-\rho^2)\bigg[\sum_{j=0}^{\infty}\theta_{j+1}^1(\rho^1)^j\bigg]\bigg[\sum_{j=0}^{\infty}\theta_{j+1}^2(\rho^2)^j\bigg].\end{aligned} \tag{1.14}$$

Comparing equations (1.12) and (1.14) makes it clear that the conditional correlation is given simply by

$$\operatorname{corr}_{t-1}(r_t^1, r_t^2) = \operatorname{corr}_{t-1}(\varepsilon_t^{\mathrm{d}1}, \varepsilon_t^{\mathrm{d}2}). \tag{1.15}$$

Returns are correlated because the news is correlated. In fact, in this simple case they are equal.

More generally, the relation (1.8) can be used to form a general expression for the covariance matrix of returns. Letting $\boldsymbol{r}$ now represent a vector of asset returns and $\boldsymbol{\eta}$ the vector of innovations due to dividend or expected returns, the equation becomes

$$\boldsymbol{r}_t - E_{t-1}\boldsymbol{r}_t = \boldsymbol{\eta}_t^{\mathrm{d}} - \boldsymbol{\eta}_t^{\mathrm{r}}, \tag{1.16}$$

where the use of bold emphasizes that these are now vectors. The covariance matrix becomes

$$\operatorname{Cov}_{t-1}(\boldsymbol{r}_t) = \operatorname{Cov}_{t-1}(\boldsymbol{\eta}_t^{\mathrm{d}}) + \operatorname{Cov}_{t-1}(\boldsymbol{\eta}_t^{\mathrm{r}}) - \operatorname{Cov}_{t-1}(\boldsymbol{\eta}_t^{\mathrm{d}}, \boldsymbol{\eta}_t^{\mathrm{r}}) - \operatorname{Cov}_{t-1}(\boldsymbol{\eta}_t^{\mathrm{r}}, \boldsymbol{\eta}_t^{\mathrm{d}}). \tag{1.17}$$

Thus correlation will result either from correlation between dividend news events or correlations between risk premiums or expected returns. Most news that is relevant for the future profitability and hence for the dividends of one company will also contain information that is relevant for many other companies. This could be expressed in a factor model, although from the definition of the innovations in (1.8) it is clear that there are many dynamic assumptions implicit in such a representation. Similarly, one might express the covariances of required return in a factor model. Presumably, the factors might include short rates, market risk premiums, credit risk premiums, and possibly other factors. As these are covariances they can easily be more important at some times than at others, so the correlations will sometimes look more like dividend correlations and sometimes more like expected-return correlations. Notice also the cross terms, which could be important but are not well-understood.

Using monthly data Campbell and Ammer (1993) find that the biggest component of the unconditional variance of stock returns is the expected return $V(\eta_t^r)$. Ammer and Mei (1996) find that this is also the biggest component of the correlation between U.K. stocks and U.S. stocks; the covariance between dividend innovations is also significant. The high correlation from news on expected returns makes it critical to understand the source of movements in risk premiums.

1.5 Additional Influences on Correlations

While this analysis has focused on fundamental news as the source of correlations across assets, there are additional considerations that should be mentioned. When returns on two assets are measured over time periods that are not identical, the correlations will be understated. These are called nonsynchronous returns. This affects correlations between assets traded in markets with different trading hours. The correlation between the U.S. market and the Japanese market when measured on a daily closing basis will be much lower than when contemporaneous returns are measured. This is because the closing time in Japan is before the U.S. market opens so some news events will affect the United States one day after they affect Japan. (See Burns et al. (1998), as well as Scholes and Williams (1977) and Lo and MacKinlay (1990a), for a discussion of this and for econometric approaches to the problem.) Burns et al. suggest "synchronizing" the data first. This is applied in Michayluk et al. (2006), where it is demonstrated that synchronized returns on real-estate portfolios are more correlated than they appear using close-to-close returns.

Nonsynchronous returns are at the heart of the late-trading scandal for mutual funds, since late trades allow the investor to observe the news but trade on pre-news prices. To a lesser extent, the same thing happens even with indices that close at last-trade prices for all the components. In this case, some components of the index have stale prices, so the full effect of correlated news will not be seen until the next day. The same effect is present when examining correlations within the day; stale prices will reduce the correlations. Thus a stylized fact is that correlations at high frequencies are often estimated to be smaller than those at low frequencies. This is called the Epps effect after an early paper by Thomas Epps (1979).

Finally, there is much discussion about how correlations between returns can arise through correlated trading or correlated positions. If many portfolios have similar positions, then a news shock to one asset could lead all of the managers to take a similar action with respect to other assets. Such action would lead to correlated order flow and very probably to correlated movements in returns. These correlations might be interpreted as responses to supply and demand effects rather than fundamental news. However, microstructure theory would interpret the order flow as a response to private information that becomes public as trades reveal the information. Thus even correlations that move in response to order flow can be interpreted as being based on news.

In the summer of 1998 following the default of Russian bonds and the decline of the hedge fund Long-Term Capital Management (LTCM), correlations behaved very strangely. It has been argued that many banks and hedge funds held similar positions and as they were unwinding these positions, asset prices of conventionally uncorrelated assets began moving together. Thus trading, and not fundamental news, moved the correlations. Similarly, in August 2007 hedge fund deleveraging is often interpreted as having led to large shifts in correlations. In both cases, however, a more general interpretation is that these trades are informative of the economic conditions in the hedge funds and of the likelihood of future orders. Hence they move prices and correlations.

Internationally, such events are called “contagion.” When one emerging market has a financial crisis, often this affects many emerging markets even though they are not economically connected. The link is hypothesized to run through portfolios. It is not clear how important these episodes are to understanding correlations. If there are correlations that are not fundamentally due to news effects, then prices are temporarily or even permanently being driven from their equilibrium value. One could imagine hedge fund strategies to profit from such activities if they were regular and systematic. In fact, hedge funds do play an important role in index reconstitution and other non-information-based trading. Thus contagion effects may well have a basis in news, if only in news about the average investor’s tolerance for risk.

Whether trades move markets or whether news moves markets is somewhat of a semantic point. It is generally thought that private information motivates trades and these trades reveal the information. In any case, the concept that news moves asset prices carries with it the idea that the type of news and its intensity will influence correlations. For practical financial decision making, it is necessary to ascertain what types of news are moving the market and forecast how these are likely to evolve in the future.