RiskMetrics Group™ Research Monthly

Much ado about correlation

Christopher C. Finger
chris.finger@riskmetrics.com

April 2007

Some years ago, it was all too common to treat correlations in a risk model as the can opener in the old economist joke[1]: we assumed we had them. Seemingly wiser now, we worry about correlation instability and correlation risk, and industry conferences still promote sessions on "coping with correlation breakdown".

But is this taking things a bit far? Before we "cope with" correlation risk, we should define it. Moreover, we should differentiate between correlations that actually move around and correlations that are fixed, but because they are by definition expectations, are difficult to estimate. Only with the problem (or problems) well stated should we delve into the modeling challenges that inevitably arise.

To state the problems, we first make the distinction between two types of correlation risk: in the first, changes in correlation impact future portfolio risk, but not the present value of the portfolio; in the second, changes in correlation impact the portfolio value directly, and also impact portfolio risk.

An analogy with volatility is useful here. In a simple portfolio, comprised only of linear positions, portfolio risk depends (among other things) on the volatility of these positions. But a volatility change does not directly impact the value of the portfolio. If a portfolio contains options, then the value of the portfolio does depend directly on volatility. So if we ask how volatility impacts our portfolio, our answer in the first case is that it changes the risk forecast, and in the second, that it impacts the portfolio value and consequently also the risk forecast.

The simple portfolio's risk, but not its value, depend on correlation. While standard options do not derive their prices from correlation, there are numerous financial products that do. For these products, we consider the analog of volatility risk: the risk of price changes due to changes in correlation. As we will see, this new notion of correlation risk brings up some modeling challenges, but these are more tractable than we might have feared.

Why worry?

In a risk modeling context, we consider the correlation that describes the level of dependence between risk factors. For the time being, we neglect the second form of correlation—that which impacts the pricing of some positions. To assess the risk of the portfolio, we must forecast correlations over the risk horizon of interest. Acknowledging that correlations are an important driver of risk, risk managers commonly request the ability to change (or "stress") them.

[1] See, for instance, http://davidwildasin.us/humor.html

Figure 1: Realized correlations for equity pairs, 50-day moving window. Actual data (left) and simulated, constant correlation data (right)

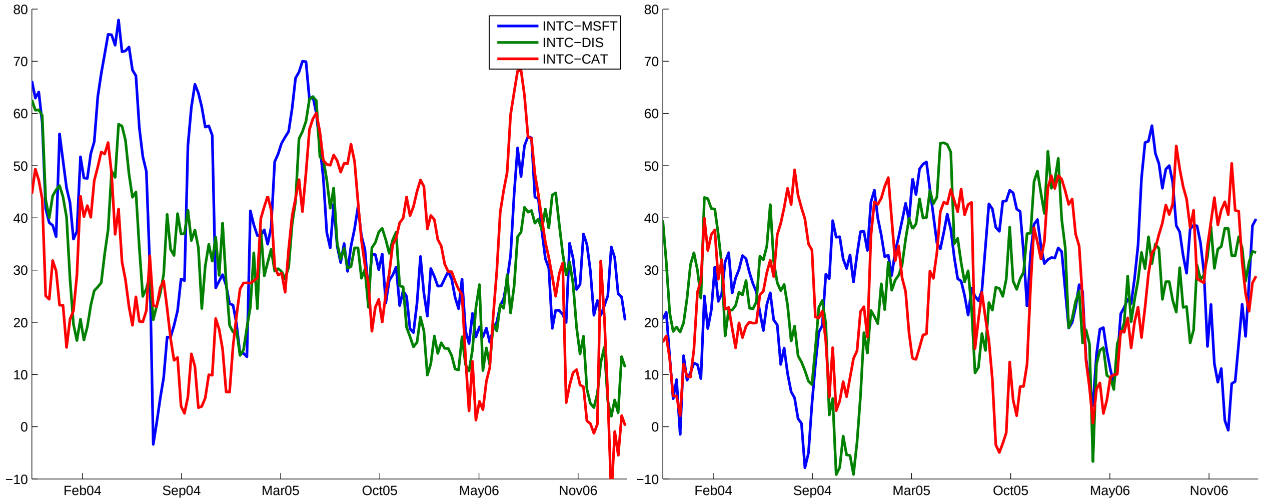

Often, the request to stress correlations arises from the conventional wisdom that correlations are notoriously unstable. Evidence of the instability is typically presented in the form of charts such as in Figure 1. In both charts, we see the correlation of three pairs of securities, as measured by the standard correlation estimate on a 50-day moving window. In both charts, we see a significant fluctuation: all three pairs on the right and two pairs on the left touch negative values at some point; and all pairs reach as high as 50%, with two pairs on the left reaching 70%.

We need interpret such charts carefully, however. Whether correlation is unstable is a point to debate; but that it is difficult to estimate is certain. Indeed, the three asset pairs on the right chart are simulated returns from return distributions assumed to have a constant correlation of 30%. Thus, all of the fluctuation in this chart is estimation error. The pairs represented on the left are three actual equity pairs whose long-term average correlations are approximately 30%. It is striking that the degree of fluctuation in the two charts is quite similar; using the right chart as a benchmark, there does appear to be some evidence of correlation instability in the actual data, but the evidence is by no means indisputable.

So what to take from this? On the one hand, the example in Figure 1 demonstrates that the more breathless pronouncements of correlation instability are likely overstated. On the other hand, the example should serve as a warning of the difficulties of arriving at a good estimate for correlation, even if that correlation is fixed.

Correlation stresses

So what are we to make of the multitude of requests to change the correlation parameters in a risk model? Does this address correlation risk? If we borrow our semantics from the volatility case, then no, since we are not assessing the amount we would (certainly) lose given a particular change in correlation as a market factor. We are instead assessing how much our *forecast* of what we could lose changes as we change the correlation as a risk model parameter.

In this sense, what we are concerned with is not the chance that correlation will change, but rather the hazard that our forecast of correlation is deficient. This deficiency could derive from a poor mode, or from a view that the period to come will be different from the recent past (even if our model has in fact been working well), or could simply be a matter of estimation error, such as what we observed in Figure 1. In all of these cases, it is certainly prudent to examine our risk forecasts under different correlation values.

Note that this is a different exercise from inspecting a crash scenario in which all securities, though some have appeared uncorrelated in the past, lose value at once. Such a test is an examination of a specific market event on the portfolio value, whereas what we aim to address is the impact of a different model parameter on the portfolio risk forecast.

We should also reiterate that the correlation here is a model parameter and an *expectation* of what is to occur over our risk horizon, but not a quantity we can ever directly observe. So in this context, there is in general no way that a specific view of correlation can prove to be correct; two factors moving oppositely on one day does not validate the view that they are negatively correlated. Being right about correlation means being right about a method for forecasting it, over a long enough period to establish something statistically.

Changing correlations is not something we can do arbitrarily. The correlation for a single pair must be between -100% and 100%, and the correlation matrix for a set of factors must be positive definite. Intuitively, this just means that a correlation structure must be logically consistent. It is logically impossible, for instance, that two securities are strongly negatively correlated to each other, and each negatively correlated to a third. (Three people cannot all walk in opposite directions on a sidewalk.)

A simple approach to stressing correlations is to use a forecasting model guaranteed to produce a positive definite matrix, but to be creative with the data we feed into the model. We could, for instance, examine our portfolio risk forecast using volatilities estimated from the most recent data, but correlations estimated from data from 1998. The drawback here is that we cannot specify any particular correlation values; we just choose a historical period that is representative of what we would like to test.

If we change the correlation values directly, we either need a mechanism to correct correlation matrices that we break inadvertently, or else a restriction on the changes such that our matrix is guaranteed to remain positive definite. Finger (1997) proposes a method in the second category, in which the user identifies blocks of factors, and within each block sets a desired new average correlation level. The method achieves these desired levels, and changes the correlations between blocks such that the overall correlation structure is still acceptable.

True correlation risk

Recalling our volatility analogy, we said that true volatility risk came from positions whose value (not simply risk) depend on volatility. What then are positions whose value depends on correlation? In general, any position whose value depends in a nontrivial way on more than one underlying factor will be exposed to the correlation across those factors. We discuss some examples of these positions below.

Synthetic collateralized debt obligations (CDOs) are derivatives whose underlying is the default loss on a specific portfolio of credits. The CDO pays out losses in a particular slice (or tranche); for example, one tranche might pay out for portfolio losses greater than 3%, and up to 7%, of the initial portfolio value. The expected loss on the portfolio is insensitive to correlation, but the likelihood that the loss falls into a particular tranche depends on the portfolio loss distribution, and therefore on the correlation between portfolio constituents.

Quanto options are options on an underlying in one currency that pay off in a second currency at a fixed exchange rate. The option is more valuable if it tends to be in-the-money in the same scenarios in which the exchange rate moves favorably to the option holder. It is thus sensitive to the correlation between the option underlying and the exchange rate.

Rainbow options and other portfolio derivatives. Rainbow options come (pardon the pun) in many colors, but typically depend on the performance of a basket of stocks. One example is a maximum option—a call option on the best performer in a basket of stocks. Such an option is less valuable if the correlation is strong (in which case there is little difference between the performance of the different stocks) and more valuable if the correlation is weak.

Spread and basis products. An option on the spread (or basis) between two prices (or rates) can be viewed as a correlation product, since the volatility of the spread will depend on the correlation between the two prices. Often though, these products are quoted directly in terms of spread or basis, and are considered as volatility, rather than correlation, exposures.

Correlation swaps. In exchange for a fixed payment, a correlation swap pays the realized correlation on a basket of securities. The realized correlation is calculated by applying the standard correlation estimate to the returns on the basket constituents, and then averaging the correlations across all constituent pairs. In the end, participants in a correlation swap are not just exposed to the theoretical (expected) correlation among items in the basket; even if they have perfect knowledge of the underlying process and correlation, they are exposed, since only a finite number of returns are observed, to the random fluctuations around the theoretical value.

That these products derive their prices from an expectation of correlation means that the market for these products can be seen as an expression of, or market for, correlation itself. And if we take positions in these products, then beyond being exposed to changes in whatever prices underlie the products, we are exposed to changes in the markets expectation of future correlation. It is this risk that we refer to as true correlation risk.

Of the five correlation product types above, three do not lead us to reliable risk factor time series. As we mentioned, spread options are not typically cast as correlation products. Rainbow options and quantos, while being true correlation products, trade typically as customized transactions. As such, they are priced at inception, and valued by the participants, but the lack of a benchmark product precludes the existence of any true correlation index.

Synthetic CDOs do in fact provide us with good correlation time series. There are a number of standard products, with tranches on index portfolios trading liquidly for at least three markets (US and European high grade, and US high yield). Moreover, the market largely utilizes a single pricing model, if only for communicating and quoting prices. The combination of a standard set of products, a liquid market, and a standard mechanism to express correlation leads to a reliable correlation benchmark.

Finally, correlation swaps give us an observable correlation factor, though somewhat indirectly. A direct indication of correlation would be to observe the price of a standard (in both the set of constituents and the time to maturity) correlation swap, and to infer the history of correlation expectations implied in these prices. Unfortunately, neither the standardization nor liquidity exists yet in these products to produce for us a correlation swap time series.

We can, however, infer such a series for equity indices from prices on options on the index and its individual constituents. The intuition for the inference comes from the decomposition of portfolio variance into two terms: the first is the sum of the individual constituent variances, and the second the sum over covariance terms. We can then suppose a single level of correlation for all pairs, and solve for the level implied by our volatilities. The implied correlation for the index can thus be thought of as the difference between the index volatility and the average constituent volatility.

Beyond allowing us to infer correlation, this decomposition motivates a particular strategy—the dispersion trade—the essence of which is a short position in an index option coupled with long positions in options on the index constituents. The goal of the dispersion trade is to isolate an exposure to expected correlation, either to express a view on future expectations, or to take advantage of the premium that the market seems to place on correlation.[2] We will examine the implied correlation for the Dow Jones Industrials basket of 30 stocks, using prices for one-month at-the-money options.

Correlation behavior

From the market then, we have two benchmark correlation series to consider: the CDX correlation,[3] which is inferred directly from the pricing of a correlation product, and the Dow Jones (DJX) correlation, which is derived from equity options based on a trade that mimics a correlation exposure. As these are new risk factors for us, our first task is to ask whether they behave similarly to our other risk factors, to assess the need for any particular new features in our risk model, and whether they behave similarly to each other, to determine whether a category of correlation risk factors is justified.

For comparison, we also examine the realized analog of each of these—that is, the average correlation, calculated over a rolling 50-day period, across all pairs of constituents of the respective indices.[4] And to keep us honest, we examine the behavior of similar realized correlation series for both a 30-stock

[2]Driessen et al (2005) show that while there appears to be little risk premium in individual equity options, there is a premium in index options. They demonstrate that the premium can be traced to a premium on correlations—that is, the market asks for compensation for greater correlation levels than are typically observed. The conclusion is that the premium in out-of-the-money index options is actually a premium for correlation “events”, where all constituents fall simultaneously.

[3]To be precise, the base correlation for the 7% detachment point, as implied by the prices of the 0–3% and 3–7% tranches on the US investment grade index

[4]For the CDX, though the implied (model) correlation should refer to assets, we take a simple approach and examine equity correlations. For both indices, we examine the history of the current constituents (Series 7 for the CDX), rather than tracking historical index changes.

Figure 2: Implied, realized, and simulated correlations for DJX (left plot) and CDX (right plot) indices

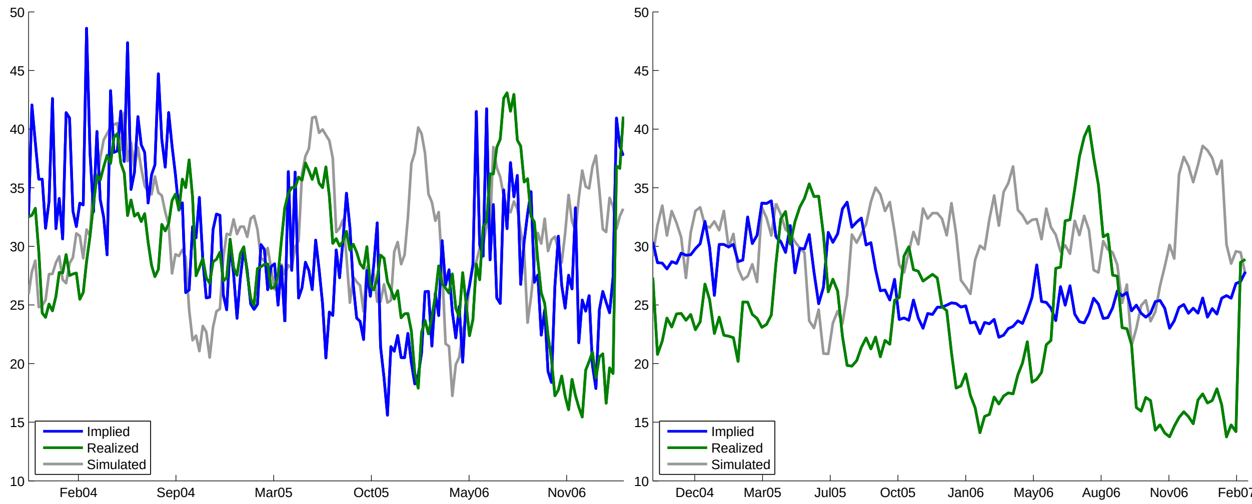

basket (to match the DJX) and a 125-stock basket (to match the CDX), using randomly generated returns assuming a constant correlation of 30%.

The correlation time series are presented in Figure 2. We observe that the implied and realized correlations have similar average levels, and move broadly in synch, though they do fluctuate in smaller ranges away from each other. One implication is that the trailing realized correlation is at least a good first guess for how to price a new correlation contract.

Our second observation is that unlike in Figure 1, here we never observe negative correlations. This is not a coincidence, as the required positive definiteness of the correlation matrix imposes a floor on the average correlation; the floor approaches zero as the size of the basket increases. Moreover, the standard CDO model does not even admit negative values for the correlation. From modeling point of view, this simplifies the problem, in that we do not have to concern ourselves with the correlation changing sign.

At a daily level, the DJX implied correlation appears significantly more volatile than its realized and simulated counterparts. Over longer timeframes, however, the implied and realized correlation achieve similar maxima and minima. In other words, the volatilities of the two series are comparable over long horizons, even if they are different at shorter horizons. The apparent disconnect here is explained by the tendency of the implied correlation toward mean reversion over short horizons.[5] Thus, sharp daily moves are likely with the implied correlation, but these are also likely to be compensated for in subsequent days. This behavior has implications for risk forecasts over longer horizons, as it suggests that a simple volatility scaling is inappropriate.

The CDX implied correlation is much less volatile, showing smaller swings than either the realized or simulated correlations. This is sensible, as the bas-

[5]The implied DJX correlation displays serial correlations of -16% and -15%, respectively, over lags of one and two days. For comparison, typical equity series display serial correlation no stronger than -2%.

[6]In fact, using a longer window for the realized correlations does produce similar daily volatility to the implied correlation, but the overall swings in this realized series are still large relative to what we observe with the implied correlation.

ket here is larger, and the maturity of the contract is longer than in the case of the DJX.[6] As with the DJX, the serial correlations are material, with any moves in the implied correlation likely to be followed by an offsetting change in subsequent days.

A final observation is that the realized and simulated correlations are difficult to distinguish, at least during calm markets. Thus, the fluctuations in realized correlations typically are not much more than what we would see at random were the true correlations constant. The events of February 2007, though, are enlightening: realized correlation for both indices jump significantly, and the magnitude of the jump is unlike anything we observe in the simulated series. The DJX implied correlation reacts similarly, while the CDX, with its much longer maturity, displays muted, if any, reaction.

What kind of change?

Our last question is how to best define daily changes in correlations. For equities, the price is just as likely to move by 5% (a relative change) regardless of its level, whereas a $1 move (an absolute change) is much more likely with the price at $100 than with the price at $10. It is most appropriate to model relative changes. For other quantities such as interest rates and credit spreads, the choice is less clear, and other notions of price changes could be more useful.

The choice of price change is relevant to any risk forecasting model, but is even relevant to the (supposedly assumption-free) historical simulation technique: historical changes applied to today's level should be just as likely to happen now as they were in the past, when the level could have been different.

For simple pairwise correlations (as in Figure 1), we need a definition that permits changes in sign, ruling out relative changes. Moreover, we need to account for changes being necessarily smaller in absolute size as the correlation approaches its bounds at 100% or -100%, ruling out absolute changes. Other change definitions might be more appropriate, but for our purposes here, we need only handle the average index correlations, which are essentially always positive, and for which the chance of approaching 100% is sufficiently remote for us to neglect.

One method to assess different definitions is to compare the magnitude of changes to the level of our risk factor. With a good change definition, the typical change magnitude should not depend on the level of the factor. This means that it is appropriate to apply historical returns to today's price level, and that we can seek to model (possibly time varying) volatility looking only at the price changes, and without introducing the complexity of a dependence on level.

As an indicator of change magnitude, we consider the trailing 25-day volatility. In Figure 3, we compare, for absolute and relative changes, this volatility to the current implied correlation level. For the DJX, the choice is clear: absolute changes show a trend toward higher magnitudes at higher correlation levels; relative change magnitudes are essentially invariant with correlation level. For CDX, neither choice is a bad one; the correlation for this index has varied so little that there is little difference in the behavior of the two change definitions.

Wrapping up

Our brief examination of correlation risk leaves us somewhat conflicted. On the one hand, it appears

Figure 3: Correlation versus volatility of correlation for DJX (blue) and CDX (green). Volatility using absolute changes (left) and relative changes (right)

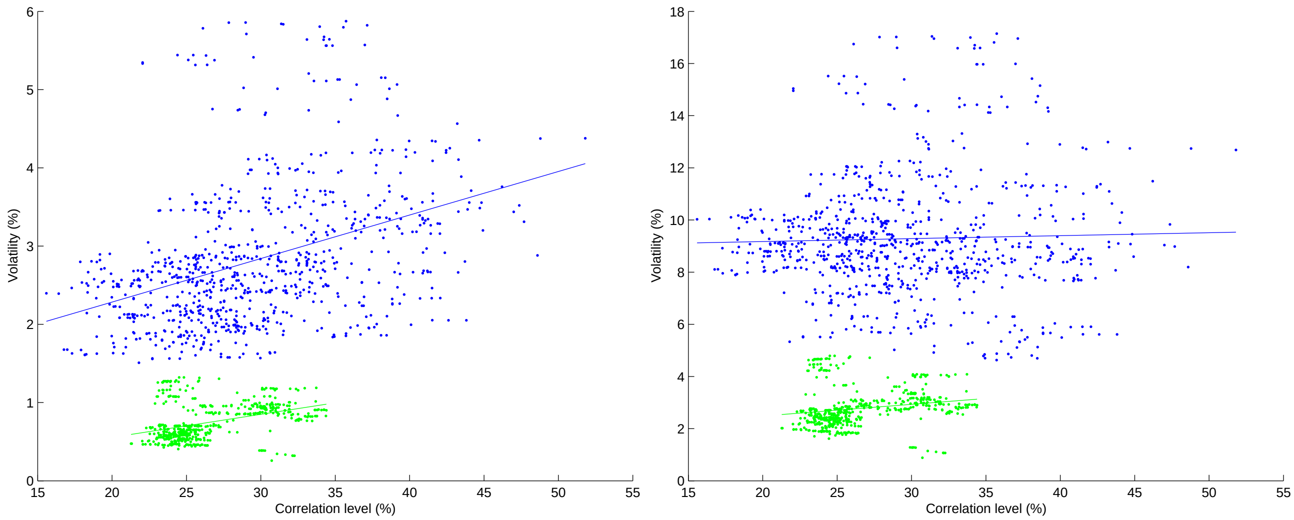

that much of the perceived risk in correlations, and much of the fluctuations in actual market correlations, reduces to an estimation problem. That correlation moves is less of a problem than that a good estimation of correlation is elusive. Thus, much of our challenge lies in assessing the impact of our estimation uncertainty.

On the other hand, the advent of a broad class of products whose values depend on correlation expectations gives us ample motivation to consider this as a new type of risk. This would appear to present a significant modeling challenge, as correlation can take on positive or negative values, and is bounded in absolute value by 100%.

But for now, the correlation risk factors for which we have meaningful benchmark time series are average correlations across large baskets of securities. For these, correlation is positive, and its bound at 100% appears irrelevant. Our results here indicate that at least one aspect of our standard model—the use of relative price changes—is appropriate. For medium risk horizons, we should account for the mean reversion that is apparent in both implied correlation series we examined. And most importantly, at least for short maturity correlation products, we should recognize the possibility that a "locking up" occurs, and the implied correlation jumps by more than any continuous model would forecast.

In all, these are sensible additions to our arsenal, but nothing earth shattering. We will be faced eventually with modeling risk driven by pairwise, rather than average, correlations. At that point, a different title, such as *The Tempest*, may be more apt.

Further reading

- Finger, C. (1997). A methodology to stress correlations. *RiskMetrics Monitor*, Q4, 3–11.
- Driessen, J., Maenhout, P., and Vilkov, G. (2006). Option-implied correlations and the price of correlation risk. European Finance Association 2005 Moscow Meetings.