AJO

"There Was Nothing I Could Do—All The Correlations Had Suddenly Gone To One!"[1]

True long-term correlations don't change so much, but realized correlations over shorter periods can be all over the map. This is what risk is, so plan for it!

M. Barton Waring and Laurence B. Siegel[2]

After every market event, we hear some traders and money managers excuse their poor returns by arguing that they were helpless, that the correlations had "all gone to one," during some very short period that they are discussing. And similarly, we hear people saying that "diversification doesn't work," a claim often made with respect to U.S. investors' holdings of foreign equities, perhaps over a somewhat longer period of a couple of years. One guesses that what is really meant is that global diversification was expected to take advantage of correlations to reduce total portfolio risk and/or increase return—but instead for many years non-US benchmarks underperformed their US counterparts, leaving U.S. investors somewhat worse off, not better off.

What this language implies is that the correlations of returns in the markets where they are trading (securities, asset classes, whatever) are expected to be constant. If the correlation as measured over some time period is, say, 0.7, then our investor apparently expects it to be 0.7 at all times; diversification should "work" the same way all the time. So when something else happens, the argument seems to be that any resultant losses aren't the traders' or portfolio managers' fault—the world had moved right under his or her feet as the unthinkable happened and correlations all went to one, or in some other way diversification didn't work. And, by phrasing the argument in terms of changing correlations or of correlations "not working," a cloak of apparent quantitative sophistication is protectively wrapped around the speaker, daring challenge from those less statistically aware.

Now it is certainly possible for the long-term correlation of two or more traded assets to change, in some sort of regime change or other basic adjustment of the markets; it happens, but that's not what we're writing about. We don't think there are much real data about how often this happens, but it is clearly far rarer than the statistically normal variation of realized short-term correlations around the long-term average, "true," or population correlation. Just because the "true" correlation of two time series—should we be so prescient as to know it—of, let's say, monthly returns is 0.7 (or whatever), doesn't remotely suggest that the sample data observed in any sub-period will be correlated at the same level.

---

[1] Forthcoming in the *Journal of Portfolio Management*.

[2] M. Barton Waring is the retired CIO for investment policies and strategies at Barclays Global Investors (now BlackRock). Laurence B. Siegel is the Gary P. Brinson director of research at the CFA Institute Research Foundation. The authors may be reached at barton.waring@aya.yale.edu and lbsiegel@uchicago.edu.

What we are saying is no different in principle from the trivial observation that the realized return on an asset is going to differ, often substantially, from the expected return. If a stock has an expected long-term annual return of 8%, no one is surprised to see it return, say, +30% in one period and -20% in another. Thus also with correlations: they vary around a long-term central tendency.

Let's do a thought experiment relevant to our short term trader to explore this ordinary variability. Let's plot out a long series of monthly returns with known true annual expected values of 8%, and known annual standard deviations of 25%; perhaps these are a couple of stocks. And let's say that they have a known true correlation of 0.7. They are two volatile but correlated series of returns. Now take any three consecutive monthly returns at random from this long series and measure their sample correlation: Yes, some of the time it will be near the true value of 0.7; but more often it will be distant from that value! Sometimes it will be negative, even large negative. Sometimes it will be near zero. And yes, sometimes—actually a lot of times—it will be close to one!

A DEMONSTRATION USING SIMULATION

You don't have to leave this as a mental experiment. We've prepared an Excel™ spreadsheet that simulates two such series for you, and we've put it up on our websites. It calculates 100 years—1200 months—of random lognormal returns for two assets, with an overall 0.7 correlation (by virtue of a Cholesky decomposition), and then analyzes 3-month, 6-month, 1-year, 2-year, and 5-year rolling period correlations and plots their histograms.[3] You can tuck in whatever annual expected returns, standard deviations, and correlation values that you are interested in, and can re-randomize it endlessly by simply hitting the Excel F9 key. This technique allows the user to see many possible future histories, and the wide range of outcomes that might happen even when correlation is known with certainty—a luxury we don't enjoy in real life!

The bottom line is: the shorter your measurement period, the more likely it is that you will experience correlations that aren't merely somewhat different than your assumed true correlation, but that are *grossly* different. And *vice versa*: the longer your sub-period, the closer to "true" will be your realized correlations, and the more tightly they will be distributed—but even with longer sub-periods there can occasionally be quite long periods where realized correlations run well off of their

---

[3] This spreadsheet can be found at www.bartonwaring.com, under the Publications tab, and at http://www.larrysiegel.org, under the Spreadsheets tab. Think of it as generating repetitive single "runs" of a Monte Carlo simulation of two correlated returns series, giving the opportunity to examine the correlation characteristics of each run over various sub-periods.

We focused here on *rolling* sub-periods. We would use *independent* sub-periods if we were examining data histories, hoping to understand more about the underlying true correlations. But rolling periods suit this discussion better because when we complain about realized correlations we do it over random sub-period intervals, regardless of start date.

long-term expectation. The math behind this isn't necessary to the ordinary practitioner; we can use our simulation spreadsheet to convey the intuition more readily.[4]

And the effect is quite dramatic. For our short term thought experiment case, about 30% of the time the three-month rolling period realized correlation will be .95 or higher—nearly one! In fact, if you examine the scatter plot of 3-month rolling correlations, shown in Exhibit 1, you'll see that they tend to cluster towards one. The wide scatter is characteristic of two highly correlated (0.70 is pretty highly correlated) series when attempting to measure correlation with only three data points; the distribution of results is spread all the way to -1.0 on the left, a span of 1.7 correlation units, but it is compressed into the much smaller distance 0.70 to 1.0 on the right. The distribution of short-period correlations *has to be* concentrated to the right in the diagram; otherwise the average could not also be on the right (0.70). The consequence is that all correlations higher than the overall average correlation appear to cluster close to one. (There is more to this for such small *n*, involving the characteristics of the beta distribution.)

EXHIBIT 1: SIMULATED ROLLING 3-MONTH SHORT-TERM REALIZED CORRELATIONS, OVER 1200 MONTHS, WITH POPULATION CORRELATION OF 0.70

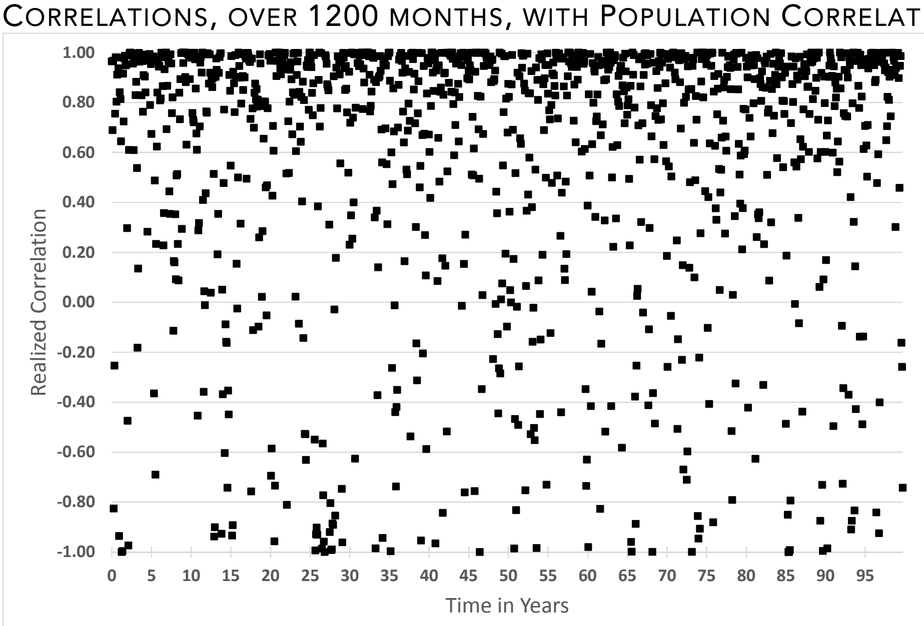

The extreme variability of the short-term correlations shouldn't really surprise anyone. Short-horizon measurements of correlation, being quite unstable, have very little value in revealing the underlying true correlation. Yes, our assumed high positive correlations, which are typical of actual data in finance, often, and quite ordinarily, "go to one"!

[4] A quick view of the math can be found at https://en.wikipedia.org/wiki/Pearson_product-moment_correlation_coefficient.

What can we take from this? Despite the high-sounding tone of blaming losses on "correlations suddenly going to one," it isn't really much of an excuse. *It is perfectly usual* for realized correlations to go to one (or any other distant-from-expected value) over short periods of time. Being a usual thing, it should be expected to happen; if it is expected to happen, then of course it should be planned for by the competent professional. It is a central part of what risk *is*, and as such the risk management tools of the trade should be employed to mitigate it: diversification, optimization, hedging, a lowering of the overall portfolio risk level.

Correlations over longer sub-periods

The variation in asset-class return correlation that occurs in the short run also occurs (and is complained about) over the longer run, perhaps even over years-long time spans. Exhibit 2 shows that the distributions of 2- and 5-year sample correlations around the true or population correlation are tighter than the three-month distributions seen in Exhibit 1, but there is still significant sample variation; of course the 5-year distribution is much tighter than the 2-year.

Exhibit 2: Simulated Distribution of Rolling 2- and 5- Year Sample Correlations, over 1200 months, with Population Correlation of 0.70

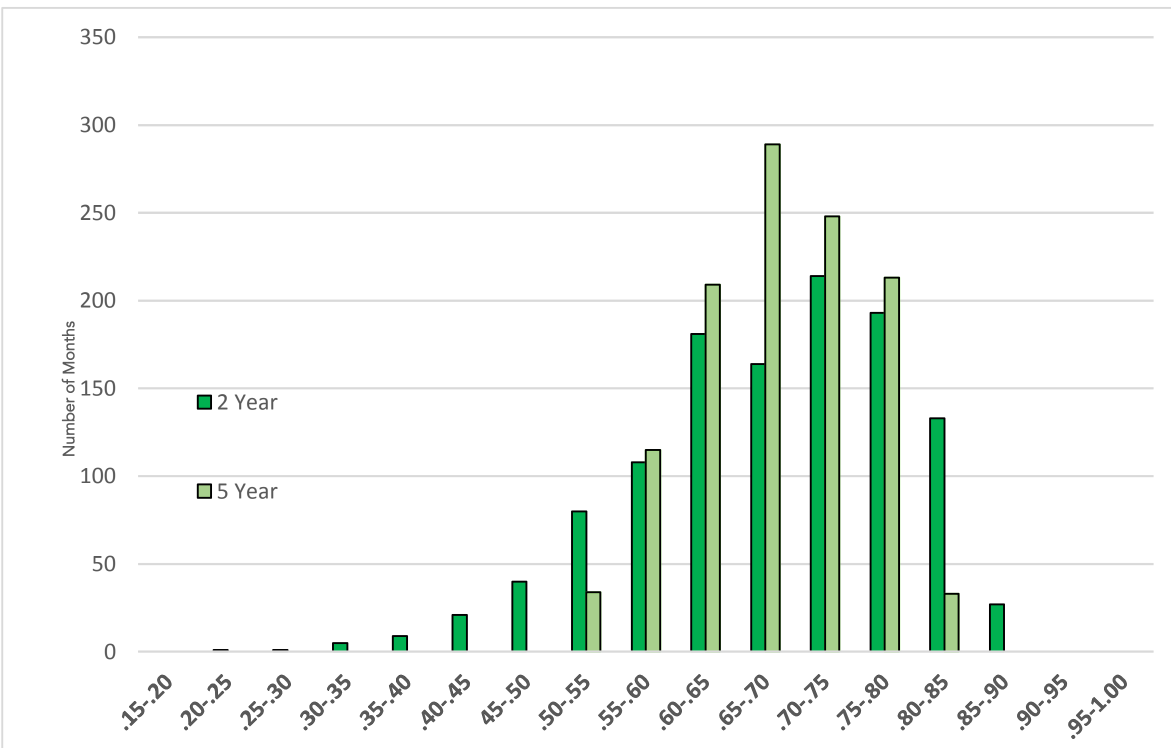

We hear that "correlations aren't working," for example, that non-U.S. equities are not helping U.S. equities in some combination of risk and return terms, over periods this long, including 1995-1999 and 2012-2014. These are periods during which U.S. equities went up more than non-U.S. equities.

With any two asset classes, one will always out-return the other (except by incredible coincidence). You do not know in advance which is which. All that investors are really saying when they complain about diversification into non-U.S. stocks not "working" is that U.S. stocks did better over some specific period. Exhibit 3 shows the rolling 36-month correlations of U.S. and non-U.S. stocks from 1979 to 2014: The correlations weren't significantly higher during the periods (1995-1999 and 2012-2014) when U.S. stocks outperformed. What both our shorter term and somewhat longer term investors are really complaining about is their own inability to forecast markets, a fault from which all of us suffer!

EXHIBIT 3: ROLLING 36-MONTH CORRELATION OF S&P 500 AND MSCI EAFE TOTAL RETURNS, FEBRUARY 1988-JANUARY 1990 TO JULY 2013-JUNE 2015

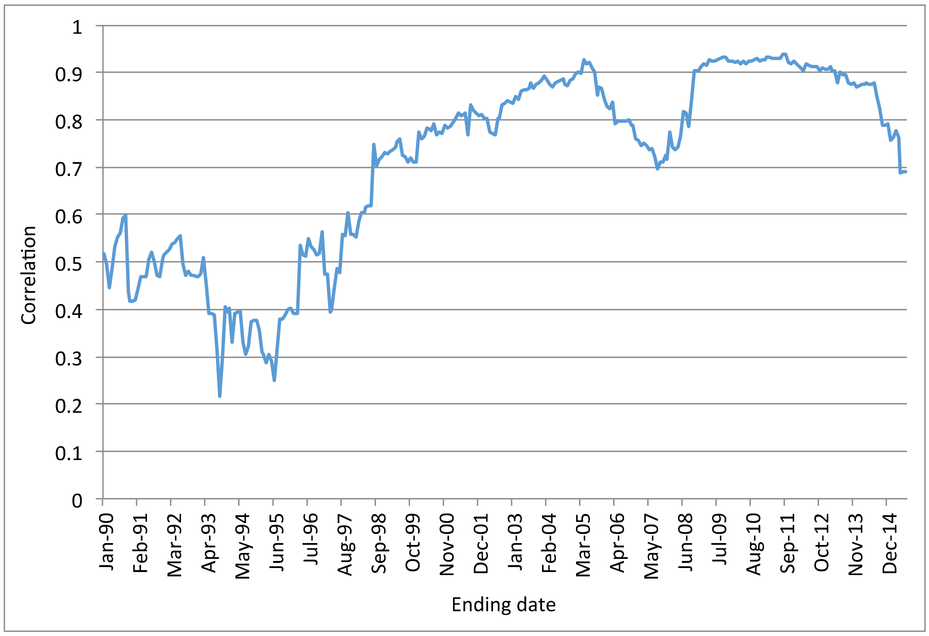

Source: Constructed by the authors using data from Morningstar.

Note that the U.S./non-U.S. equity correlation over the whole period was 0.71, but highly variable over shorter (36-month) periods.

DO COMMON FACTORS DRIVE CORRELATION?

Our simulation data are sterile by design—the correlation is there by virtue of a statistical method, and variations in it are purely random. In the real world, however, there are dozens or maybe hundreds of common factors that come into play, maybe only occasionally but maybe more often, and one can think of the correlation as the sum of the many smaller correlations from these common factors. If correlation measured over a short period reveals a spike upward it may be random, or it may be because the two assets are reacting to the same common factor, say, a slowing in the global growth rate.

In this narrow sense, perhaps it is fair to say that correlations really do "go to one" in stressful times if that factor emerges. But slow down—the effects of common factors, along with a lot of completely random events, were taken into account when we made our original and hopefully unbiased estimate of correlation: The fact that the correlation is sometimes very high (and sometimes not) is one of the reasons that the overall or expected correlation is high in the first place! In any single event it is quite unlikely that we have the data to say with confidence that some common factor suddenly overwhelmed all the other common factors and drove the correlation to one—so that assertion adds little insight. Estimating long-term correlations, *ex post* or *ex ante*, is difficult enough—estimating short-term correlations is almost literally impossible.

Not surprisingly, then, those portfolio managers who have diversified their equity holdings to include non-U.S. equities are in fact doing the right thing, assuming their original equilibrium estimate of expected correlation was unbiased and fair. Correlation benefits will likely work out over longer periods of time—but not with absolute certainty, only with a very high probability.

Conclusion

Correlations are never expected to be a constant across sub-periods. It's a mistake to think that they should be. So don't let yourself be surprised by the ordinary! If fairly estimated, they will reduce risk consistent with the standard math over long periods, but don't expect risk reduction/diversification benefits over every sub-period.