Copulas and Coherence

Portfolio analysis in a non-normal world.

Kevin Dowd

Traditional portfolio theory (TPT) assumes we live in a world where risk factors are normally (or a little more generally, elliptically) distributed. Under these circumstances, we can take the portfolio standard deviation to be a reliable proxy for portfolio risk, and we can take a plot of portfolio expected return against portfolio standard deviation as the trade-off between expected return and risk. If we want, we could also express the risk side of this trade-off in terms of other risk measures such as the value at risk (VaR).

The normality (or ellipticality) assumed in TPT is multivariate. At its heart is the assumption that we can model the statistical dependence or comovement between random variables using linear (or Pearson) correlation coefficients. The validity of the linear correlation coefficient carries over to correlation matrices and variance-covariance matrices, and TPT is often expressed in terms of such matrices.[1]

Unfortunately, most real-world risk factors are not elliptically distributed, and under these circumstances we cannot rely on TPT to give us reliable guidance on portfolio management. Once we step outside the limiting confines of multivariate ellipticality, portfolio analysis runs into two difficult problems:

- We cannot rely on correlation-based (e.g., variance-covariance) approaches to model statistical dependence.
- Widely used risk measures such as the portfolio standard deviation or VaR become unreliable.

Kevin Dowd
is a professor of financial risk management at Nottingham University Business School, and director of research at Black Swan Risk Advisors.
kevin.dowd@nottingham.ac.uk

The first problem is a statistical one; it tells us we need an alternative way to model dependence that is valid in non-elliptical situations. The second problem is a matter of risk measurement theory; it tells us we need to be careful in our choice of risk measure.

MODELING DEPENDENCE

The use of linear correlations to model statistical dependence is so ingrained in financial risk management practice that many practitioners have little awareness of any alternatives. Its great advantages are its simplicity and that in an elliptical world the correlations tell us everything we need to know about dependence. Yet correlation is only one of many dependence measures.[2]

One limitation is that the correlation is not invariant to transformations of the underlying variables (e.g., the correlation between $\ln X$ and $\ln Y$ is not generally the same as that between X and Y). Another is that the correlation does not always go from -1 to 1 (e.g., two variables can always move in the same direction as each other, and have a correlation of less than 1).

More seriously, in some circumstances correlation is not defined. This would be the case if one of the variables concerned did not have a finite variance, or if the two variables were not cointegrated. In other circumstances the correlation tells us very little about dependence, especially dependence in the tails. The bottom line is that correlation will often be of limited or even no use in non-elliptical situations.

We need a better way to model dependence, and the answer is provided by copulas. The term *copula*, from the Latin, refers to connecting or joining together (as in the English word, *couple*). In statistical terms, a copula is a function that joins together a collection of marginal distributions—the statistical distributions for each random variable taken on its own—to form the multivariate distribution that describes how they all move together. And, viewed conversely, copulas enable us to take a multivariate density function and separate its dependence structure from the marginal distribution functions.[3]

Copula Theory

The key result is a theorem due to Sklar [1959]. Suppose for convenience that we are concerned with only two random variables, X_1 and X_2. If $F(x_1, x_2)$ (the probability that X_1 is less than or equal to some constant x_1 and that X_2 is less than or equal to some other constant x_2) is a joint distribution function with continuous marginals $F_{x_1}(x_1) = u_1$ and $F_{x_2}(x_2) = u_2$ (so $F_{x_1}(x_1)$ is the probability that $X_1 \leq x_1$, and this probability is set to $u_1 \ldots$), $F(x_1, x_2)$ can be written in terms of a unique function $C(u_1, u_2)$:

$$F(x_1, x_2) = C(u_1, u_2) \tag{1}$$

where $C(u_1, u_2)$ is known as the copula of $F(x_1, x_2)$.

The copula function describes how the multivariate function $F(x_1, x_2)$ is derived from or coupled with the marginal distribution functions $F_{x_1}(x_1)$ and $F_{x_2}(x_2)$, and we can interpret the copula as giving the dependence structure of $F(x_1, x_2)$.

This result is important, because it enables us to construct joint distribution functions from marginal distribution functions in a way that takes account of the dependence structure of our random variables. To model the joint distribution function, we specify our marginal distributions, choose a copula, estimate the parameters involved, and then apply the copula function to our marginals. Once we can model the joint distribution function (or copula), we can use it to estimate risk measures, identify the risk-return trade-off, and so on.

Advantages over Correlation

Copulas have a number of advantages over correlation as measures of dependence. The most obvious one is that copulas are universally valid, in the sense that they can be applied to non-elliptical as well as elliptical distributions. Given that most real-world multivariate distributions are non-elliptical, this means that copulas give us (and correlation-based approaches do not) a valid means of modeling the dependence structures in many empirical multivariate distribution functions.

Copulas are also extremely flexible, because they allow us to work at two separate levels. At the first level we focus on the marginal distributions, and fit the appropriate marginal distribution to each factor in turn. We then consider the dependence structure (i.e., the copula) to be fitted to the marginal distributions.

This separation makes for much greater flexibility because it means that we can in principle fit quite different marginal distributions to each risk factor (e.g., we can fit a normal distribution to one risk factor, a t to another, and so on). With correlation-based approaches, however, we have to fit the same marginals to each risk factor. This same separation also makes for greater flexibility in another respect; working with copulas gives us a much greater range of choice over the dependence structure that we can fit to our data.

In short, copulas provide a means of modeling dependence that is much more flexible and has a much greater range of applicability than traditional correlation-based approaches provide.

MEASURES OF FINANCIAL RISK

Granted that copulas enable us to estimate our risk measures properly, there remains the problem that measures such as the standard deviation or the VaR are unreliable in non-elliptical situations. The standard deviation gives us a reliable risk measure in an elliptical context because the higher moments (skewness, kurtosis) of elliptical distributions are pinned down by the distribution. For example, for a normal distribution the skewness is 0 and the kurtosis is 3.

For non-elliptical distributions, however, the higher moments are free parameters that can take many different values (e.g., distributions might be very skewed, or have greater than normal kurtosis) and can vary from one situation to the other. Two such distributions might have very different shapes and yet have the same standard deviation; in these circumstances the standard deviation does not suffice to give us a reliable indication of the shape of the distribution.

For its part, the VaR is also known to be a very unreliable risk measure in these same circumstances. Perhaps the biggest problem is that the VaR then fails to be subadditive.[4]

Subadditivity reflects the natural expectation that risks should diversify, or at least should not increase, when we put them together, and subadditivity is a basic requirement of any sensible risk measure. The failure of VaR to be subadditive is therefore a major drawback to the VaR's credibility as a risk measure. This lack of subadditivity also means that a VaR-based risk-return analysis can give extremely misleading results.[5]

The best response to these sorts of problems is to go back to first principles and set out the criteria that we would expect a good risk measure to satisfy. Various sets of criteria have been proposed, but perhaps the best (and certainly best-known) are the axioms of coherence set out by Artzner et al. [1999]. The most important of these axioms is that of subadditivity, and the VaR's lack of subaddivity means that the VaR is not coherent.

Risk measures that satisfy the coherence axioms—coherent risk measures—have a number of advantages over the VaR. For example, coherent risk measures tell us something about how bad bad states (or losses exceeding VaR) might be, while the VaR tells us nothing about them; coherent risk measures are more compatible with expected utility maximization, and are more reliable when used for risk-return decision-making; and coherent risk measures can be relied on to generate convex risk surfaces (while the VaR cannot), which ensures that portfolio optimization problems have a well-defined optimum and that standard programming techniques can find that optimum.

There are many coherent risk measures to choose from, but the best known is the expected shortfall (ES)—the average of the worst $p\%$ of outcomes, where p is 100% minus the confidence level expressed in percent. That is, if we are working at the 99% confidence level, the ES is the average of the worst 1% of losses.

RISK-RETURN ANALYSIS IN A NON-ELLIPTICAL WORLD: EXAMPLE

The answers to our earlier problems are now clear. Outside ellipticality, we should model dependence using a copula rather than a correlation-based approach, and we should use a coherent risk measure (such as the ES) instead of unreliable alternatives such as the standard deviation or the VaR. To obtain our trade-off between expected return and risk, we can estimate expected returns the old-fashioned way, but we take the risk side of this trade-off to be a coherent risk measure that we estimate using copulas.

It is helpful to illustrate the issues involved using an example. This will highlight both the mechanics of using copulas for portfolio analysis and the difference that copulas can make in our estimated risk-return trade-off.

Suppose we have two financial assets whose returns r_1 and r_2 have normal marginal distributions. The portfolio return is therefore $wr_1 + (1 - w)r_2$, where w is the proportion of the portfolio invested in Asset 1, and we wish to estimate the 99% ES of the portfolio over some period. We model the dependence between r_1 and r_2 and using a Clayton copula:[6]

$$C(u_1, u_2) = \max\left\{\left[u_1^{-\alpha} + u_2^{-\alpha} - 1\right]^{-\frac{1}{\alpha}}, 0\right\} \quad \text{for } \alpha > 0 \qquad (2)$$

where u_1 and u_2 are standard uniform random variables (i.e., are distributed evenly over the range [0, 1]). This is a convenient one-parameter copula, and we assume its parameter α is known.[7]

The easiest approach to estimate VaR and ES is

Monte Carlo simulation, based on simulations of u_1 and u_2 from the specified copula:

- We first simulate independent variables v_1 and v_2 from a standard uniform distribution.
- Set $u_1 = v_1$.
- We then simulate u_2, from the distribution of u_2, conditional on the realized value of u_1. This conditional distribution $C(u_2 | u_1)$ depends on the copula, and in the case of the Clayton copula, $C(u_2 | u_1)$ is given by:[8]

$$C_2(u_2|u_1) = (u_2^{-\alpha}+u_1^{-\alpha}-1)^{-\frac{\alpha+1}{\alpha}}u_1^{-(\alpha+1)} \tag{3}$$

We can then obtain u_2 by setting this expression to v_2 and solving it for u_2. This gives us:

$$u_2 = \left[v_1^{-\alpha}\left(v_2^{-\frac{\alpha}{\alpha+1}}-1\right)+1\right]^{-\frac{1}{\alpha}} \tag{4}$$

- Having obtained u_1 and u_2, we regard these as the values of the marginal cumulative density functions for r_1 and r_2. Now recall that r_1 is normal. Let its mean and standard deviation be μ_i and σ_i, and denote its cumulative density function as $\Phi(r_i; \mu_i, \sigma_i)$ for $i = 1, 2$. We then obtain our simulated values for r_1 and r_2 by inverting $\Phi(r_i; \mu_i, \sigma_i) = u_i$ using standard algorithms. This gives us $r_i = \Phi^{-1}(u_i, \mu_i, \sigma_i)$, and we plug these return values into $wr_i + (1 - w)r_2$ to obtain our simulated portfolio return.
- We repeat these steps many m times, and sort the m simulated portfolio returns from lowest to highest. We then estimate the 99% ES as the negative of the average of the worst $k = (1 - 0.99)\ m$ returns.

We can use this simulation routine to estimate the ES for any portfolio allocation (i.e., any value of w), and estimating the expected portfolio return (equal to $w\mu_1 + (1 - w)\mu_2$) is trivial. We then vary w from 0 to 1, estimate the associated expected portfolio returns and portfolio ES, and plot the former against the latter to give us the trade-off between expected return and risk.

An (empirically plausible) example of this trade-off is shown in the Exhibit. Note that this trade-off is correct in a dual sense. It is correct in the sense that the dependence between the two returns is modeled correctly, using the assumed (Clayton) copula; and it is correct in the sense that the risk measure used, the ES, is known to be well-behaved in non-elliptical situations.

EXHIBIT
Trade-Offs Between Risk and Expected Return

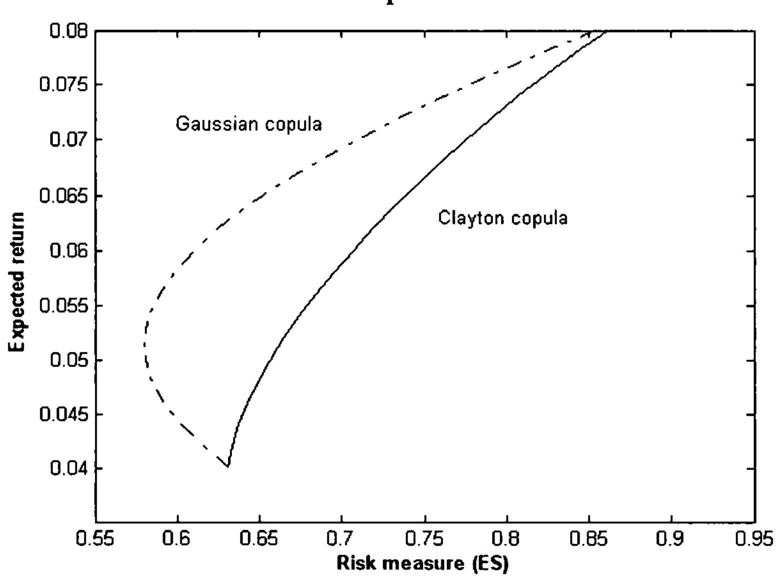

The return to Asset 1 has a normal marginal with mean 0.04 and standard deviation 0.25; the return to Asset 2 has a normal marginal with mean 0.08 and standard deviation 0.35. The Gaussian copula is predicated on a linear correlation coefficient of 0.4; this implies a Kendall's τ of 0.2594, and this value of τ implies that a Clayton copula has an α of 0.7000. The ES is predicated on a 99% confidence level and holding period equal to that on which the return parameters are predicated. Estimates are based on 50,000 simulation trials.

The Exhibit also shows the trade-off we would have obtained had we used a Gaussian copula with parameter values that are consistent with those of the Clayton copula. Given that our marginals are normal, this is equivalent to assuming that returns are multivariate-normal. Thus, the one curve gives us the true trade-off between risk and expected return based on the assumed copula, and the other gives us the trade-off that we would have obtained had we made the conventional assumption of multivariate normality (e.g., as in TPT).

The difference between the two curves is striking. In this example, the true trade-off is usually well to the right of the conventional TPT one, implying that the true risk measures can be considerably greater than conventionally estimated ones.

The lesson is very clear. Incorrectly

assuming ellipticality (or failing to use the appropriate copula) can lead to major errors in estimates of the trade-off between risk and expected return. This shows that copulas are practically and not just theoretically important for portfolio management.

ENDNOTES

The author thanks Carlos Blanco, John Cotter, and Peter Oliver for very helpful comments on an earlier draft of this article. Thanks also to the Economic and Social Research Council for a research fellowship on risk measurement in financial institutions (RES-000-27-0014).

[1]Elliptical distributions have "normal-like" density functions, and include distributions such as finite-variance t distributions. They can therefore be regarded as generalizations of the normal. Strictly speaking, TPT also presupposes that our positions are linear functions of normal (or elliptical) risk factors, but we take the linearity for granted in the present context.

The linear or Pearson correlation coefficient between two variables is defined as the covariance between the two variables divided by the product of their standard deviations.

[2]Alternatives include Kendall's τ, the Spearman rank correlation, Blomqvist's β, and the Gini coefficient. For more on these risk measures, see, e.g., Cherubini, Luciano, and Vecchiato [2004, chapter 3].

[3]For more on copulas and their financial applications, see Nelson [1999] and Cherubini, Luciano, and Vecchiato [2004].

[4]If $\rho(.)$ is a measure of risk, and A and B are any two risky positions, the risk measure $\rho(.)$ is said to be subadditive of $\rho(A + B) \leq \rho(A) + \rho(B)$.

[5]The root of the problem is that any VaR-based decision analysis ignores the size of prospective losses exceeding VaR. Hence, provided the probabilities of these losses are sufficiently low, a VaR-based risk-return analysis would tell us to take on the additional risk, however high those prospective losses might be and however small the extra expected return might be. This ignores the marginal trade-offs on which any "sensible" risk-return analysis must always depend. For more on this and other problems with the VaR, see, e.g., Artzner et al. [1999], Acerbi [2004], or Dowd [2005].

[6]The Clayton is also an empirically plausible copula, and both Cherubini and Luciano [2001] and Breymann, Dias, and Embrechts [2003] present evidence that it fits real-world return data well.

[7]Estimating the parameter of this type of copula is also straightforward. For more on this, see, e.g., Cherubini, Luciano, and Vecchiato [2004, chapter 5].

[8]This conditional simulation has a familiar parallel in the variance-covariance paradigm. To simulate correlated random numbers, we often start by simulating independent ones; we then apply a Choleski matrix to transform the independent random numbers into dependent ones. The Choleski approach is explained in Hull [2003, pp. 412-413].

REFERENCES

Acerbi, C. "Coherent Representations of Subjective Risk Aversion." In G. Szegö, ed., *Risk Measures for the 21st Century.* New York: John Wiley, 2004, pp. 147–207.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath. "Coherent Measures of Risk." *Mathematical Finance,* 9 (November 1999), pp. 203-228.

Breymann, W., A. Dias, and P. Embrechts. "Dependence Structures for Multivariate High-Frequency Data in Finance." *Quantitative Finance,* 3 (2003), pp. 1-14.

Cherubini, U., and E. Luciano. "Value at Risk Trade-Off and Capital Allocation with Copulas." *Economic Notes,* 30 (2001), pp. 235-256.

Cherubini, U., E. Luciano, and W. Vecchiato. *Copula Methods in Finance.* Chichester: John Wiley, 2004.

Dowd, K. *Measuring Market Risk,* 2nd ed. Chichester: John Wiley, 2005.

Hull, J. *Options, Futures, and Other Derivatives,* 5th ed. Upper Saddle River, NJ: Prentice-Hall, 2003.

Nelson, R.B. *An Introduction to Copulas.* Berlin: Springer-Verlag, 1999.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.

©Euromoney Institutional Investor PLC. This material must be used for the customer's internal business use only and a maximum of ten (10) hard copy print-outs may be made. No further copying or transmission of this material is allowed without the express permission of Euromoney Instituitonal Investor PLC. Copyright of Journal of Portfolio Management is the property of Euromoney Publications PLC and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.
The most recent two editions of this title are only ever available at http://www.euromoneyplc.com