Analysing correlations under stress

A structural assessment of co-variability during periods of stress can improve crisis management and contribute to better strategic planning, argues **David Rowe**

It is commonly recognised that correlations in a crisis vary significantly from those we observe under normal market conditions. Unfortunately, the analysis of such crisis behaviour often stops at a breezy claim that correlations all migrate to one or minus one. It certainly is true many correlations tend to shift to large absolute values in a stress situation, but this observation alone is of little value for diagnosing how a crisis might unfold. What we need to do is apply structural imagination to evaluate which pairs of variables will exhibit highly positive and which will exhibit highly negative correlation in a given crisis.

One place this type of thinking is badly needed is in the area of corporate strategic planning. Sam Savage, consulting professor at the Stanford University department of management science and engineering, and two colleagues have advanced an interesting approach to addressing this problem.[1] They point out that a common approach to authorising new investment projects is to rank them in terms of their expected returns and fund those where the expected return is highest. However, this approach effectively ignores the role played by risk in such allocation decisions.

The authors' proposed alternative is to apply a variation of the well-established financial tools associated with portfolio construction. Doing this in the context of real projects rather than financial portfolios, though, does present some operational challenges. Maintaining statistical consistency across projects demands that common sources of risk be modelled identically in each case. For example, if several oil exploration projects can be represented in terms of an uncertain distribution of physical output, this can be translated into a distribution of home currency revenue by applying relevant distributions for the oil price (and for specific exchange rates where appropriate) to each simulated level of production volume. Coherence across the firm, however, requires application of the same distributions for the price and foreign exchange variables to the analysis of each project. Applying non-consistent price and forex distributions will reduce the calculated aggregate revenue volatility of the combined projects by ignoring common movements in these external variables. One way of addressing this issue of statistical coherence is to create a central corporate authority charged with maintaining relevant and internally consistent simulated values for the required variables. The authors' preferred name for this certifying authority is the chief probability officer. Of particular interest is the potential this approach holds for capturing the reinforcing or offsetting impact of different variables in stress scenarios. Consider, for example, a decision to choose any two of three projects. One project is a gas extraction project in Nigeria and the other two are offshore oil projects in Norway. On the surface, the Norwegian projects appear to offer modest returns with little political risk. The Nigerian project offers attractive expected returns but carries significantly higher risk of loss from political upheavals of various kinds. Despite its high expected return, the Nigerian project might be rejected as too risky. A portfolio view, however, reveals an interesting point. A political upheaval that significantly affected that project would also reduce Nigerian production more generally. This would increase the price of oil globally and enhance the profitability of the Norwegian projects. Therefore, the Norwegian and Nigerian projects represent natural hedges for each other and this casts the portfolio decision in a very different light. The Nigerian investment is effectively less risky in the context of also investing in the Norwegian project.

Savage *et al* describe some differences that arise when applying their approach to real investment decisions rather than to financial portfolio construction. In general, real projects are of more or less fixed size and either have to be 'in' or 'out' of the portfolio. In addition, there is no direct analogue to the traditional risk-return trade-off. Instead, there are many potential trade-offs between pairs of metrics, including reserve impact versus revenue, short-term versus long-term benefits, and so forth.

One interesting cultural effect arose from using this approach to project funding decisions. Rather than the deliberation revolving around a list of specific projects to be included or excluded, the focus was on the set of feasible portfolios consistent with the investment budget constraint. While the framework was initially surprising to those who employed it, the discussion quickly shifted from 'where does my project rank on the priority list?' to 'how does my project contribute to the total portfolio?'. In essence, everyone was forced into a more global perspective and the process became a motivation to act as a more cohesive team in choosing the optimum portfolio of projects.

Applying modern portfolio insights to real investment allocation decisions will not be an easy transition, but including rigorous assessment of risk implications in such deliberations is long overdue. ■

David Rowe is president of David M Rowe Risk Advisory, a risk management consulting firm. Email: davidmrowe@dmrra.com

[1] See S Savage, S Scholtes and D Zweidler, Probability management; OR/MS today, *February 2006, available at* www.lionhrtpub.com/orms/orms-2-06/frprobability.html