Hedge fund risk and VAR uncertainty

External transparency of the risk of hedge funds continues to be a difficult issue. But even internally, traditional risk measures can fail to portray the full implication of highly leveraged positions. David Rowe argues that the parameter sensitivity of value-at-risk estimates can help

On first encounter, the term 'hedge fund' seems counter-intuitive. Hedging, after all, is a well-known method for reducing risk. Why then are hedge funds, with Long-Term Capital Management (LTCM) as exhibit one, such risky investment vehicles? The answer, of course, is that 'hedge fund' is something of a misnomer. These vehicles might better be called 'convergence-arbitrage investment funds'. Traditional pension and mutual funds sell shares to investors and distribute the proceeds among a variety of long asset positions. Some aggressive funds might borrow from banks or in the debt markets and invest the additional sources of funds in increased holdings of risky assets. This traditional source of leverage, however, is plainly visible in a fund's financial statements. More importantly, the extent of such leverage is limited by the willingness of lenders to provide such funding.

Hedge fund managers, on the other hand, don't hold a simple portfolio of long positions at all. Rather, they try to find anomalies between prices of two or more securities they believe will revert to a more normal pattern in a reasonable length of time. They then short the relatively 'rich' security and use the proceeds from the short sale to fund a long position in the relatively 'cheap' security. The expectation is that when spreads return to more normal levels they will be able to unwind the position at a net profit. Either both positions will show gains or the gains on one will more than offset the losses on the other.

Naturally, in active markets, such pricing anomalies tend to be small. As a result, reasonable returns can only be achieved if this is done on a leveraged basis. The gross long and short positions are a multiple, in the case of LTCM a multiple as high as 25, of the fund investors' capital.

Calibrating the degree of leverage

In analysing the demise of LTCM, Philippe Jorion presented an excellent analysis of how convergence-arbitrage traders think about the appropriate degree of leverage.[1] The first step is to decide on the target level of volatility of returns on the in-
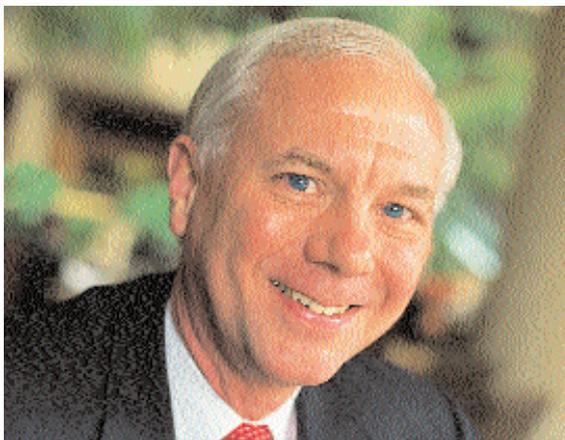
David Rowe is group executive vice-president for risk management at SunGard Trading and Risk Systems
e-mail: david.rowe@risk.sungard.com

vestor's capital. In the case of LTCM, this target was stated to be the volatility of an unleveraged portfolio of US equities, eg, the portfolio implied in the S&P 500. Given the estimated correlations between the components of the spread trades, the volatility of returns can be calculated conditional on the degree of leverage. Then it is straightforward to solve for the degree of leverage required to achieve the target volatility.

What can go wrong?

The problem is that empirical estimates of the market's behavioural parameters are unstable, reflecting changes in the underlying stochastic process. Correlation coefficients are especially prone to this type of second-order uncertainty. Moreover, trades based on extremely high correlations carry additional risks. First, the degree of leverage required to achieve a given expected return must be high to compensate for the comparative stability of the spread. Second, being bounded by 1.0, any significant shift in the correlation must be downward not upward. In this situation, reliance on continuance of recent market experience can be much riskier than traditional value-at-risk measures indicate. As Jorion also points out, optimising risk versus return based on a given set of market parameters, then mea-

suring VAR using those same parameters, leads to a serious optimisation bias. While this would be perfectly acceptable in a world driven by stable stochastic processes, it is fraught with danger when market parameters are just empirical estimates of a shifting underlying reality.

One logical response is stress testing. But the key is to focus on those stress scenarios of particular relevance to the current portfolio. One approach is to isolate the worst-case loss scenarios in a Monte Carlo simulation and examine which market data changes are associated with these results.

An alternative is to derive systematically the sensitivity of the VAR estimate to changes in the value of the parameters used to drive the simulation. Then those parameters to which the VAR result is highly sensitive can be examined more carefully, for example by reviewing a longer historical sample.

Jorion cites an example of the correlation between US Treasury yields and those on BAA-rated corporate bonds. Using a two-year moving estimate, this correlation was consistently in the 0.95 to 0.98 range from 1994 to 1998. Examination of the previous six years, however, shows it fluctuating in the range of 0.85 to 0.90 and even dropping as low as 0.75 at one point. Since this kind of credit spread is reportedly typical of the convergence trades entered into by LTCM, their risk profile would have been highly sensitive to a shift in this correlation parameter. Jorion calculates that if the operative assumption was a correlation of 0.97, a drop to 0.80 would have multiplied LTCM's estimated risk on such trades by over two-and-a-half.

Systematic reporting of the sensitivity of VAR results to underlying parameters would provide a valuable guide to which of these warrant more careful examination as to the reliability of their estimates. Human nature being what it is, this would not eliminate unexpected losses on this kind of trading, but it would likely reduce their frequency. ∎

[1] Jorion P, 2000, Risk Management Lessons from Long-Term Capital Management, *European Financial Management, September, pages 277–300*