Editorial

Cross-sectional volatility is a recent concept that is now widely used in the asset management industry. Although commonly used, it does not appear to be particularly well understood — the formal definition involves taking a universe of stock returns at a point in time, treating each return as a data point, and computing the usual formula for sample variance; cross-sectional volatility results from taking the square root of this quantity. Clearly, if all stocks have the same return, then cross-sectional volatility is zero. Otherwise, it is a positive number measuring dispersion in returns at a point in time.

As a statistical measure, its properties remain rather mysterious since it depends upon the cross-sectional population means, population variances and population covariances of the stock universe, and we can imagine situations where the volatilities and the means may differ very little, but the covariances differ tremendously. In such a case, the result would be high cross-sectional volatility.

Practitioners use this measure to capture investment opportunity. Broadly speaking, it is believed that if cross-sectional volatility is high, so too are returns to active management. Cross-sectional volatility was undoubtedly high during the technology bubble, and recently has been much lower. A topic of great interest, to me, at least, is the formal links between the statistical properties of this risk measure and the practitioner interpretation. I believe that these are not fully understood.

Stephen E. Satchell
Editor