Long Memory and Regime Shifts in Asset Volatility

Jonathan Kinlay,
Partner, Investment Analytics, jkinlay@investment-analytics.com

Long Memory

The conditional distribution of asset volatility has been the subject of extensive empirical research in the last decade. The overwhelming preponderance of evidence points to the existence of pronounced long term dependence in volatility, characterized by slow decay rates in autocorrelations and significant correlations at long lags (e.g. Crato and de Lima, 1993, and Ding, Granger and Engle, 1993). Andersen, et al, 1999, find similar patterns for autocorrelations in the realized volatility processes for the Dow 30 stocks—autocorrelations remain systematically above the conventional Bartlett 95% confidence band as far out as 120 days. Comparable results are seen when autocorrelations are examined for daily log range volatility, as the figure below illustrates. Here we see significant autocorrelations in some stocks as far back as two years.

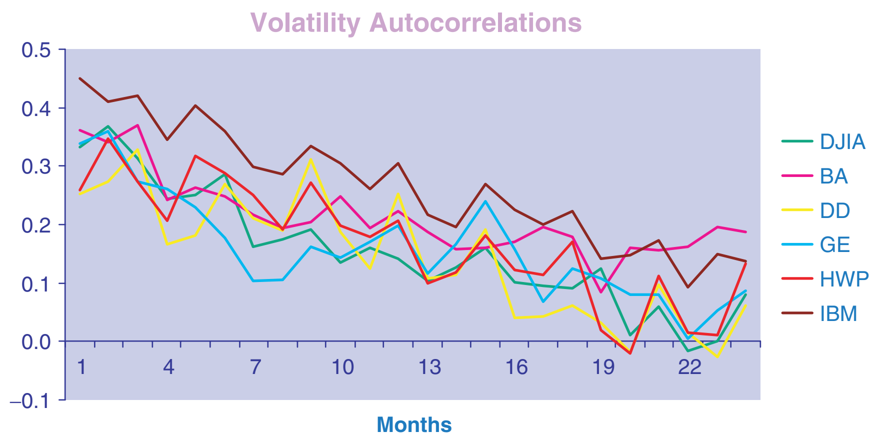

Long Memory Detection and Estimation

Among the first to consider the possibility of persistent statistical dependence in financial time series was Mandelbrot (1971), who focused on asset returns. Subsequent empirical studies, for example by Greene and Fielitz (1977), Fama and French (1988), Porteba and Summers (1988) and Jegadeesh (1990), appeared to lend support for his findings of anomalous behavior in long-horizon stock returns. Tests for long range dependence were initially developed by Mandelbrot using a refined version of a test statistic, the Rescaled Range, initially developed by English hydrologist Harold Hurst (1951). The classical rescaled range statistic is defined as

$$R/S(n) = \frac{1}{s_n}\left[\underset{1 \le k \le n}{\mathrm{Max}\sum_{j=1}^{k}(X_j - \bar{X}_n) - \mathrm{Min}\sum_{j=1}^{k}(X_j - \bar{X}_n)}\right]$$

Where s_n the sample standard deviation:

$$s_n = \left[\frac{1}{n}\sum_j (X_j - \bar{X}_n)^2\right]^{1/2}$$

The first term is the maximum of the partial sums of the first k deviations of X_j from the sample mean. Since the sum of all n deviations of the X_j's from their mean is zero, this term is always nonnegative. Correspondingly, the second term is always nonpositive and hence the difference between the two terms, known as the range for obvious reasons, is always nonnegative.

Madelbrot and Wallis (1969) use the R/S statistic to detect long range dependence in the following way. For a random process there is scaling relationship between the rescaled range and the number of observations n of the form:

$$R/S(n) \sim n^H$$

where H is known as the Hurst exponent. For a white noise process $H = 0.5$, whereas for a persistent, long memory process $H > 0$. The difference $d = (H - 0.5)$ represents the degree of fractional integration in the process.

Mandelbrot and Wallis suggest estimating the Hurst coefficient by plotting the logarithm of $R/S(n)$ against $\log(n)$. For large n, the slope of such a plot should provide an estimate of H. The researchers demonstrate the robustness of the test by showing by Monte Carlo simulation that the R/S statistic can detect long-range dependence in highly non-Gaussian processes with large skewness and kurtosis. Madelbrot (1972) also argues that, unlike spectral analysis which detects periodic cycles, R/S analysis is capable of detecting nonperiodic cycles with periods equal to or greater than the sample period.

The technique is illustrated below for the volatility process of General Electric Corporation, a DOW Industrial Index component. The estimated Hurst exponent given by the slop of the regression, approximately 0.8, indicates the presence of a substantial degree of long-run persistence in the volatility process. Analysis of the volatility processes of other DOW components yield comparable Hurst exponent estimates in the region of 0.76–0.96.

A major shortcoming of the rescaled range is its sensitivity to short range dependence. Any departure from the predicted behavior of the R/S statistic under the null hypothesis need not be the result of long-range dependence, but may merely be a symptom of short-term memory. Lo (1991) show that this results from the limiting distribution of the rescaled range:

$$\frac{1}{\sqrt{n}} R/S(n) \Rightarrow V$$

Where V is the range of a Brownian bridge on the unit interval.

Suppose now that the underlying process $\{X_j\}$ is short range dependent, in the form of a stationary $AR(1)$, i.e.

$$r_t = \rho r_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, \sigma^2), |\rho| \in (0, 1)$$

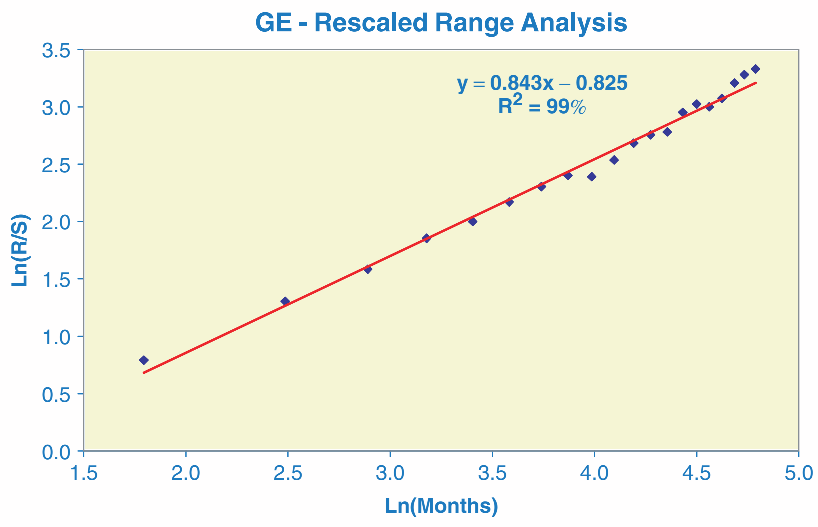

Estimation of Hurst Exponent for GE Volatility Process

The limiting distribution of $R/S(n)/\sqrt{n}$ is $V[(1+\rho)/(1-\rho)]^{1/2}$. As Lo points out, for some common stocks the estimated autoregressive coefficient is as large at 0.5, implying that the mean of $R/S(n)/\sqrt{n}$ may be biased upward by as much as 73%. In empirical tests, Davies and Harte (1987) show that even though the Hurst coefficient of a stationary Gaussian $AR(1)$ is precisely 0.5, the 5% Mandelbrot regression test rejects this null hypothesis 47% of the time for an autoregressive parameter of 0.3

To distinguish between long-range and short-term dependence, Lo proposes a modification of the R/S statistic to ensure that its statistical behavior is invariant over a general class of short memory processes, but deviates for long memory processes. His version of the R/S test statistic differs only in the denominator. Rather than using the sample standard deviation, Lo's formula applies the standard deviation of the partial sum, which includes not only the sums of squares of deviations for X_j, but also the weighted autocovariances (up to lag q):

$$\hat{\sigma}_n^2(q) = \frac{1}{n}\sum_{j=1}^{n}(X_j - \bar{X}_n)^2 + 2\sum_{j=1}^{q}\omega_j(q)\hat{\gamma}_j, \omega_j(q) = 1 - \frac{j}{q+1}, q < n$$

where the γ_j are the usual autocovariance estimators.

While in principle this adjustment to the R/S statistic ensures its robustness in the presence of short-term dependency, the problem remains of selecting an appropriate lag order q. Lo and MacKinlay (1989) have shown that when q becomes relatively large to the sample size n, the finite-sample distribution of the estimator can be radically different from its asymptotic limit. On the other hand, q cannot be taken too small as the omitted autocovariances beyond lag q may be substantial. Andrews (1991) provides some guidance on the choice of q, but since criteria are based on asymptotic behavior and little is known about the optimal choice of lag in finite samples.

Another method used to measure long-range dependence is the detrended fluctuation analysis (DFA) approach of Peng et al (1994) and further developed by Viswanathan et al (1997). Its advantage over the rescaled range methodology is that it avoids the spurious detection of apparent long-run correlation due to non-stationarities. In the DFA approach the integrate time series $y(t')$ is obtained:

$$y(t') = \sum_{T=1}^{t'} x(t).$$

The series $y(t')$ is divided into non-overlapping intervals each containing m data points and a least squares line is fitted to the data. Next, the root mean square fluctuation of the detrended time series is calculated for all intervals:

$$F(m) = \sqrt{\frac{1}{T}\sum_{t'=1}^{T}[y(t') - y_m(t')]^2}$$

A log-log plot of $F(m)$ vs the interval size m indicates the existence of a power-scaling law. If there is no correlation, or only short term correlation, then $F(m) \propto m^{1/2}$, but if there is long-term correlation then $F(m)$ will scale at rates greater than $\frac{1}{2}$.

A third approach is a semi-parametric procedure to ovbtain an estimate of the fractional differenceing parameter d. This techniques, due to Geweke and Porter-Hudak (GPH), is based on the slope of the spectral density around the angular frequency $w = 0$. The spectral regression is defined by:

$$\ln\{I(\omega_\lambda\} = a + b\ln\left\{4\sin^2\frac{\omega_\lambda}{2}\right\} + n_\lambda, \lambda = 1, \ldots, \upsilon$$

Where $I(w_\lambda)$ is the periodogram of the time series at frequencies $w_\lambda = 2\pi\lambda/T$ with $l = 1, \ldots, (T-1)/2$. T is the number of observations and υ is the number of Fourier frequencies included in the spectral regression. The least squares estimate of the slope of the regression line provides an estimate of d. The error variance is $\pi^2/6$ and allows for the calculation of the t-statistics for the fractional differencing parameter d. An issue with this procedure is the choice of υ, which is typically set to $T1/2$, with Sowell (1992) arguing that u should be based on the shortest cycle associated with long-run correlation.

The final method we consider is due to Sowell (1992) and is a procedure for estimating stationary ARFIMA models of the form:

$$\Phi(L)(1-L)^d(y_t - \mu) = \Theta(L)\varepsilon_t$$

Where Φ and Θ are lag polynomials, d is the fractional differencing parameter, μ is the mean of the process $y_t \sim N(\mu, \Sigma)$ and ε_t is an error process with zero mean and constant variance σ_e^2. We can use any set of exogenous regressors to explain the mean: $z = \mathbf{y} - \mu$, $\mu = f(\mathbf{X}, \beta)$.

The spectral density function is written in terms of the model parameter d, from which Sowell derives the autocovariance function at lag k in the form:

$$\gamma(k) = \frac{1}{2\pi}\int_0^{2\pi} f(W)e^{iwk}\,dw$$

The parameters of the model are then estimated by exact maximum likelihood, with log likelihood:

$$\log L(d, \phi, \theta, \beta, \sigma_e^2) = -\frac{T}{2}\log(2\pi) - \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{z}'\Sigma^{-1}\mathbf{z}$$

Structural Breaks

Granger and Hyung, 1999, take a different approach to the analysis of long term serial autocorrelation effects. Their starting point is the standard $I(d)$ representation of an fractionally integrated process y_t of the form:

$$(1-L)^d y_t = \varepsilon_t$$

where d is the fractional integration parameter and, from its Maclaurin expension

$$(1-L)^d = \sum_{j=0}^{\infty}\pi_j L^j, \pi_j = \frac{j-1-d}{j}\pi_{j-1}, \pi_0 = 1$$

The researchers examine the evidence for structural change in the series of absolute returns for the SP500 Index by applying the sequential break point estimation methodology of Bai (1997) and Bai and Perron (1998) and Iterative Cumulative Sums of Squares (ICSS) technique of Aggarwal, Inclan and Leal, 1999. Bai's procedure works as follows. When the break point is found at period k, the whole sample is divided into two subsamples with the first subsample consisting of k observations and the second containing the remaining $(T-k)$ observations. A break point is then estimated for the subsample where a hypothesis test of parameter consistency is rejected. The corresponding subsample is then divided into further subsamples at the estimated break point and a parameter constancy test performed for the hierarchical subsamples. The procedure is repeated until the parameter constancy test is not rejected for all subsamples. The number of break points is equal to the number of subsamples minus 1. Bai shows how the sequential procedure coupled with hypothesis testing can yield a consistent estimate for the true number of breaks.

Aggarwal, Inclan and Leal's (1999) approach uses the Iterative Cumulative Sums of Squares (ICSS) as follows. We let $\{\varepsilon_t\}$ denote a series of independent observations from a normal distribution with zero mean and unconditional variance σ_t^2. The variance within each interval is denoted by τ_j^2, $j = 0, 1, \ldots, N_t$, where N_t is the total number of variance changes in T observations and $1 < k_1 < k_2 < \cdots < k_{NT} < T$ are the set of change points.

So $\quad \sigma_t = \tau_j k_j < t < k_{j+1}$

To estimate the number of changes in variance and the point in time of the shift a cumulative sum of squares is used.

Let $C_k = \sum_{t=1}^{k} \varepsilon_t^2$, $k = 1, \ldots, T$ be the cumulative sum of the squared observations from the start of the series until the k^{th} point in time. Then define $D_k = (C_k/C_T) - k/T$.

If there are no changes in variance over the sample period, the D_k oscillate around zero. Critical values based on the distribution of D_k under the null hypothesis of no change in variance provide upper and lower bounds to detect a significant change in variance with a known level of probability. Specifically, if $\max_k \sqrt{(T/2)}|D_k|$ exceeds 1:36, the 95$^{\text{th}}$ percentile of the asymptotic distribution, then we take k^*, the value of k at which the maximum value is attained as an estimate of the change point.

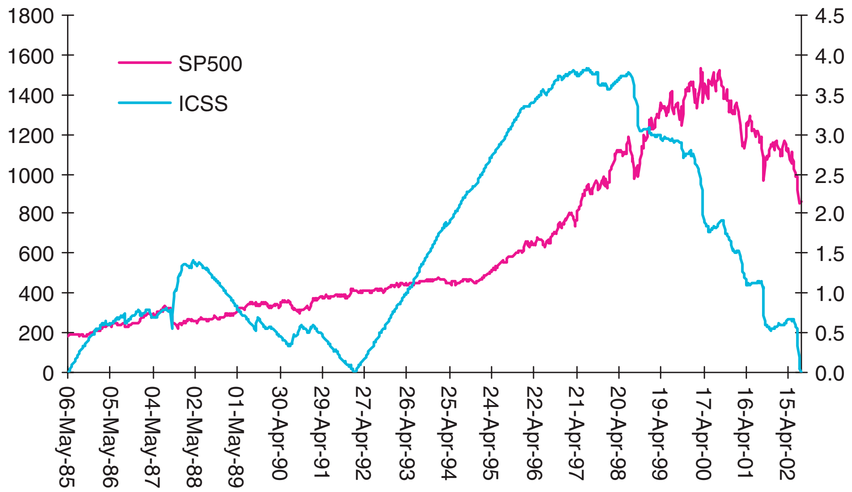

Testing for Structural Breaks in SP500 Index Returns Using Iterative Cumulative Sums of Squares

The figure below illustrates the procedure for a simulated GBM process with initial volatility of 20%, which changes to 30% after 190 periods, and then reverts to 20% once again in period 350. The test statistic $\sqrt{(T/2)}|D_k|$ reaches local maxima at $t = 189(2.313)$ and $t = 349(1.155)$, clearly and accurately identifying the two break points in the series.

A similar analysis is carried out for the series of weekly returns in the SP500 index from April 1985 to April 2002. Several structural shifts in the volatility process are apparent, including the week of 19 Oct 1987, 20 July 1990 (Gulf War), the market tops around Aug 1997, Aug 1998 and Oct 2000.

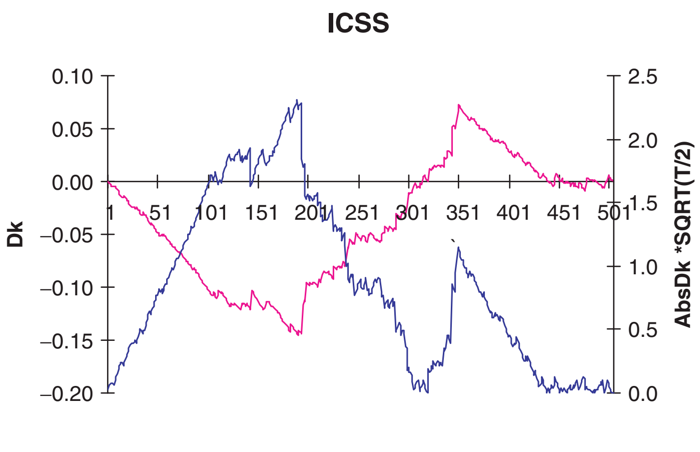

Testing for Structural Breaks in Simulated GBM Process Using Iterative Cumulative Sums of Squares

In their comprehensive analysis of several emerging and developed markets, Aggarwal et al identify numerous structural shifts relating to market crashes, currency crises, hyperinflation and government intervention, including, to take one example, as many as seven significant volatility shifts in Argentina over the period from 1985–1995.

It is common for structural breaks to result in ill-conditioning in the volatility processes distribution, often in the form of excess kurtosis. This kind of problem can sometimes be resolved by modeling the different regime segments individually. Less commonly, regime shifts can produce spurious long memory effects. For example, Granger and Hyung estimate the degree of fractional integration d in daily SP500 returns for 10 subperiods from 1928–1991 using the standard Geweke and Porter-Hudak approach. All of the subperiods have strong evidence of long memory in the absolute stock return. They find clear evidence of a positive relationship between the time-varying property of d and the number of breaks, and conclude that the SP500 Index absolute returns series is more likely to show the "long memory" property because of the presence of a number of structural breaks in the series rather than being an $I(d)$ process.

Stocks in Asian-Pacific markets typically exhibit volatility regime shifts at around the time of the regional financial crisis in the latter half of 1997. The case of the ASX200 Index component stock AMC is typical (see figure below). Rescaled range analysis of the entire volatility process history leads to estimates of fractional integration of the order of 0.2. But there is no evidence of volatility persistence is the series post-1997. The conclusion is that, in this case, apparent long

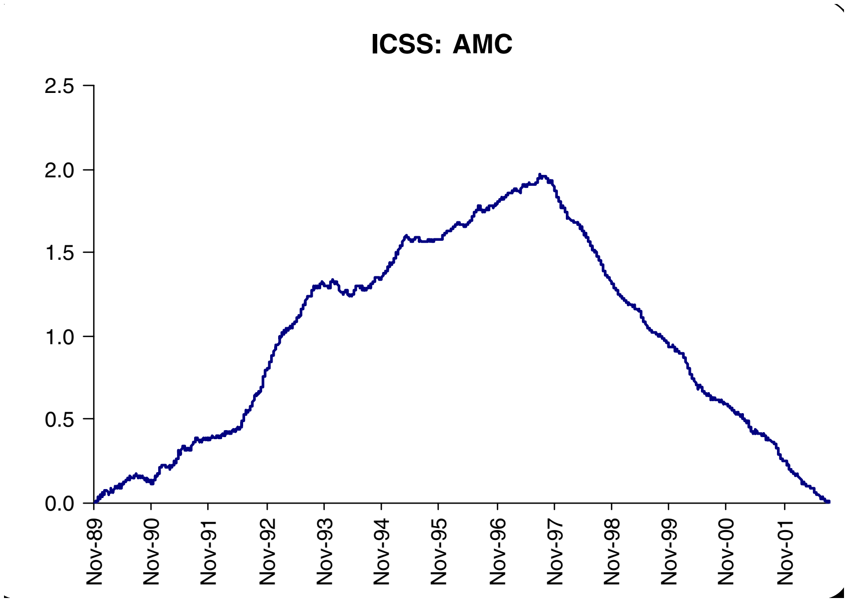

Structural Breaks in the Asian Crisis Period for ASX Component Stock AMC

memory effects are probably the result of a fundamental shift in the volatility process.

Conclusion

Long memory effects that are consistently found to be present in the volatility processes in financial assets of all classes may be the result of structural breaks in the processes themselves, rather than signifying long-term volatility persistence.

Reliable techniques for detecting regime shifts are now available and these can be used to segment the data in a way that reduces the risk of model misspecification.

However, it would be mistaken conclude that all long memory effects must be the result of regime shifts of one kind or another. Many US stocks, for example, show compelling evidence for volatility persistence both pre- and post-regime shifts. Finally, long memory effects can also result from the interaction of a small number of short-term correlated factors.

FOOTNOTES & REFERENCES

■ Aggarwal, R., Inclan, C., Leal, R., 1999, Volatility in Emerging Stock Markets, Journal of Financial and Quantitative Analysis, Forthcoming.
■ Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P. (1999), "The Distribution of Exchange Rate Volatility," Wharton Financial Institutions Center Working Paper 99-08 and NBER Working Paper 6961.
■ Bai, J., 1997, Estimating Multiple Breaks One at a Time, Econometric Theory, 13, 315–352.
■ Bai, J., Perron, P., 1998, Estimating and Testing Linear Models with Multiple Structural Changes, Econometrica, 66, 47-78.
■ Breidt, J.F., Crato, N., de Lima, P., 1998. The detection and estimation of long memory in stochastic volatility. Journal of Economics 83, 325–348.
■ Davies, R and Harte, D, 1987, "Tests for Hurst Effect", Biometrika, 74, 95–101.
■ Ding, Z., Granger, C.W.J., Engle, R.E., 1993. Long memory properties of stock market returns and a new model. Journal of Empirical Finance 1, 83–106.
■ Fama, E, and French, K, 1988, "Permanent and Temporary Components of Stock Prices", Journal of Political Economy, 96, 246–273.
■ Geweke, J., Porter-Hudak, S., 1983, The Estimation and Application of Long Memory Time Series Models, Journal of Time Series Analysis, 4, 15–39.
■ Granger, C.W.J., Hyung, N., 1999, Occasional Structural Breaks and Long Memory, University of California, Discussion Paper.
■ Greene, M., and Fielitz, B., 1997, Long-Term Dependence in Common Stock Returns, Journal of Financial Economics, 4, 339–349.
■ Jegadeesh, N, 1990, "Evidence of Predictable Behavior of Security Returns", Journal of Finance, 45, 881–898.
■ Lo, A, 1991, "Long-Term Memory in Stock Market Prices", Econometrica, 59, 1279-1313.
■ Lo, A and MacKinlay, 1988, "Stock Market Prices Do Not Follow Randow Walks: Evidence from a Simple Specification Test", Review of financial Studies, 1, 41–66.
■ Mandelbrot, B, 1971, "When Can Price Be Arbitraged Efficiently? A Limit to the Validity of the Random Walk and Martingale Models", The Review of Economics & Statistics, 53, 225–236.
■ Peng, C.K., S.V. Buldyrev, S. Havlin, M. Simons, H.E. Stanley, A.L. Goldberger, Phys. Rev. E 49 (1994) 1684.
■ Porteba, J, and Summers, L, 1988, "Mean Reversion in Stock Returns Evidence and Implications", Journal of Financial Economics, 22, 27–60.
■ Sowell, F., 1989. Maximum likelihood estimation of fractionally integrated time series models, Mimeo, Carnegie-Mellon University.
■ Sowell, F., 1992. Maximum likelihood estimation of stationary univariate fractionally integrated time series models. Journal of Econometrics 53, 165–188.
■ Teyssiere, G., 1996. Double long-memory financial time series, Paper No 348, Department of Economics, QMW, London.
■ Viswanathan, G.M. , S.V. Buldyrev, S. Havlin, H.E. Stanley, Biophys. J. 72 (1997) 866.