Educated estimates

Kostas Giannopoulos *and* **Brian Eales** *look at two different approaches to estimating risk, and ask the question: Which is best, exponential smoothing or Garch volatility?*

Volatility can be described as the uncertainty surrounding a forecast value. Given a series of returns with expected value $E(Y_t)$ and observed value Y_t at time t, the forecast error can be defined as:

$$\varepsilon_t = Y_t - E(Y_t) \tag{1}$$

If forecasts are ranked according to their variance forecast error, σ_t^2, the variance for all realisations of ε_t, it can be mathematically shown that those methodologies which dynamically update their forecasts of Y_t as new information becomes available (known also as conditional models) are superior to those methodologies which produce static forecasts (unconditional models).

Examples of conditional approaches are the autoregressive (AR) processes, autoregressive moving average (ARMA) processes, Bayesian and state space models, while the historical mean approach, and moving averages, provide examples of unconditional forecasting approaches. Exponential smoothing falls between these two approaches, since it updates forecasts as new information becomes available. The methodology does, however, lack other properties that the conditional models mentioned above have – efficiency, unbiasedness etc – because those models are estimated using advanced econometric techniques (eg maximum likelihood estimation, GMM) where certain statistical properties have to be satisfied.

Since the uncertainty of a forecast can be represented by volatility, a way exists of producing confidence bands in which forecast values can oscillate. In the case of those methods using forecasts generated in an unconditional framework, the confidence bands will remain a constant distance from each forecast.

Those forecasts generated in a conditional framework will have confidence bands which alter to reflect the dynamics of the situation at each point in time. Thus during quiet periods, confidence bands will narrow, whilst periods of rapid activity will see the bands widen: intuitively the conditional approach is to be preferred. In this article, the properties of two popular volatility models, Garch and ES, will be compared, and their forecasting power analysed using daily data on selected equity indexes.

Exponential Smoothing (ES)

Given a series of returns[1], Y, the ES model[2] for conditional variance is given as:

$$\sigma_t^2 = \lambda\sigma^2_{t-1} + (1-\lambda)Y^2_{t-1} \qquad 0 < \lambda < 1. \tag{2}$$

Today's volatility is σ_t. From **(2)** it emerges that the current level of volatility, σ_t, is function of yesterday's volatility and the square of yesterday's returns, Y_{t-1}.

Generalised Autoregressive Conditional Heteroskedasticity (Garch)

The representation in **(2)** is very similar to that of the popular Garch (1,1) model, introduced by Robert Engle (1982). The Garch model, in its simplest form, is given as:

$$\sigma^2_t = \omega + \alpha Y^2_{t-1} + \beta \sigma^2_{t-1} \tag{3}$$

so the current variance, depends upon yesterday's surprise, Y_{t-1} , and volatility, σ_{t-1}. These two models have many similarities, ie today's volatility is estimated conditionally upon the information set available at each period, t. Both the Garch (1,1) model in **(3)** and the ES model in **(2)** use last period's returns to determine current levels of volatility. Thus implying that today's volatility is known immediately after yesterday's market closure.

Since the latest available information Y_{t-1}, is weighted in a more effective way, it can be shown that both models will provide more accurate estimators of volatility than unconditional models, ie historical volatility.

However, several differences exist in the operational characteristics of the two models. The Garch model, for example, uses two independent coefficients, α and β, to set the impact that last period's errors and volatility in determining current volatility, while the ES model uses only one coefficient, λ, and forces the variables, Y^2_{t-1} and σ^2_{t-1}, to have a unit effect on the current period's volatility. Thus, a large shock will have longer lasting impact on volatility in model **(2)** than in **(3)**.

The terms α and β in Garch do not need to sum to unity and one parameter is not the complement of the other. Their estimation is achieved by maximising the likelihood function.[3] This is very important, since the values of α and β are critical in determining the current levels of volatility. Incorrect selection of the parameter values will adversely affect the estimation of volatility. The assumption that α and β (see model 3) sum to unity is, however, very strong and presents an hypothesis that can be tested rather than a condition to be imposed.

Furthermore, the Garch model has an additional parameter, ω, that acts as a floor and prevents volatility dropping below that level. In the extreme case the α and β equal zero, volatility is constant and equal to ω. The value of ω is estimated together with α and β using maximum likelihood estimation and the hypothesis $\omega = 0$ can be tested easily.

The absence of the ω parameter in the ES model allows volatility, after a few quiet trading days, to drop very low. Examples of this can be seen in Figures 1-6.

Garch modelling also it allows for many varied structures.

Fig 1-6: Exponential smoothing versus Garch volatility

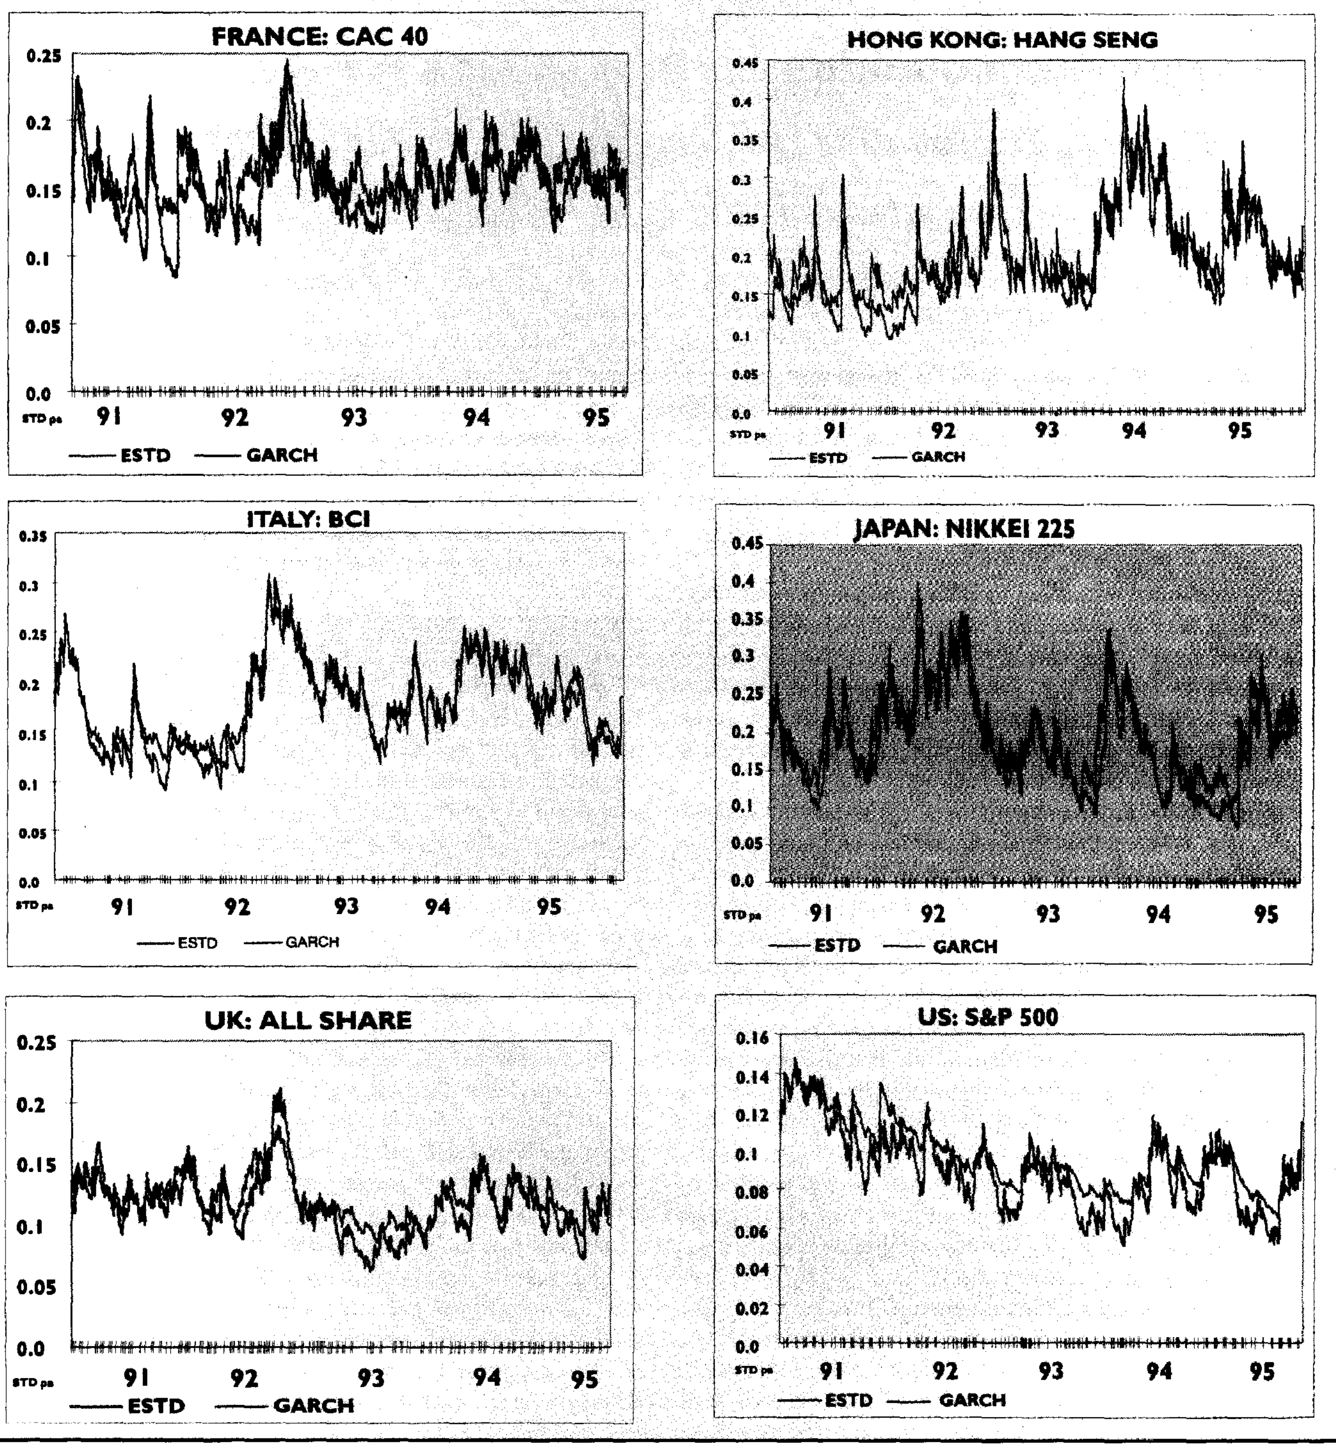

The Garch (1,1) model can be easily modified to allow for asymmetries in volatility. This type of modification enables negative returns to have a greater impact on the current estimate of volatility than positive returns. One such model specification is the A(G)ARCH of equation.

(4) $\sigma^2_t = \omega + \alpha\,(Y_{t-1} + \gamma)^2 + \beta\,\sigma^2_{t-1}$

The estimated value for γ is negative, which implies a negative correlation between price changes and volatility.

Testing the model fit

A good estimator of volatility needs to explain as much as possible of the daily variation of return (as daily variation is considered the square of return) and be unbiased (ie it should not systematically over- or under-estimate volatility).

One way to test the forecasting power and unbiasedness of a volatility model consists in estimating the following linear regression:

(5) $Y_t^2 = a + b\,\sigma_t^2 + v_t$

thus regressing the square of the forecast errors against a constant and the estimates of conditional variance. This procedure can be repeated for each of the volatility models (Garch and ES). For the volatility estimator to be unbiased, the constant term must not be statistically significantly different from zero, and the estimated slope coefficient should not be significantly different from one. The R-squared of this regression reports the power that the volatility model has to predict the next day's square returns.

A second test, based on the Ljung-Box (Q) statistic, investigates the forecast squared errors standardised (Y_t^2/σ^2). Large values of this statistic could support the hypothesis that the forecast errors are not independent, while low values of the Q-statistic would provide evidence that they are independent. Failure to accept the maintained hypothesis of independence would indicate that the estimated variance has not removed all the clusters of volatility. This would imply that the data still holds information that can be usefully translated into volatility.

An Empirical Investigation

Six equity indices were used to test the forecasting power of the two models. The dataset consists of daily closing quotes from the following countries, US (S&P500), Japan (Nikkei 225), UK (All Share), France (CAC-40), Hong Kong (Hang Seng), Italy (BCI) – over the period 3 January 1991 to 28 July 1995. The ES models for each index was estimated with λ set equal to 0.94[4].

The results for the exponential smoothing model using the square of yesterday's shocks (4) are presented in Table 1. In the columns labelled 'a' and 'b' the first set of values refer to the estimates of the regression coefficients used in the test. The values directly underneath them, in parentheses, are the estimated t-statistics constructed to test the hypothesis that a = 0, and b = 1. The R-squared result associated with the test is also reported.

One striking feature of the estimates of 'a' in these results is that in all the markets values are positive and significantly different from zero when using standard t-tables at the 95% level of confidence. Which suggests that the ES model systematically under-estimates the variance. This may be because the parameter ω, which acts as a floor in the Garch estimation of volatility, is missing from the model. When performing the test on the hypothesis that the slope coefficient is significantly different from unity, both the S&P 500 data and the CAC 40 data yield results that indicate yesterday's variance is being given too great a role in the production of today's variance. It is likely that this phenomena would be corrected if λ were set to a lower value, but finding an appropriate value is, of course, associated with more computational effort, and therefore, cost.

The last column reports the estimated Q-statistic, of order 6, on the standardised residuals of the Garch process. In the row below, the p-values of the reported statistics are presented. The Q-statistic is distributed as χ^2 with degrees of freedom set equal to the number of lags used to undertake the test. In this work six lags have been used.

At a 5% significance level, standard tables provide a critical value of 12.6. Clearly, in respect of the results presented in Table 1, at this level of significance the S&P 500, CAC 40, and the Italian index (BCI) all fail the test, suggesting that the model used does not offer a full explanation of the historical behaviour of volatility.

With the Garch model, the diagnostic tests displayed in Table 2 show a good model fit with each data series. In respect of the test for unbiasedness, at the 95% level of confidence, the maintained hypotheses cannot be rejected. This indicates that there is no systematic bias in the variance estimator. The Q-statistic also suggests that the Garch volatility has successfully removed all clusters present in the data. Looking at figures 1-6, there is a clear picture of the exponential smoothing volatility estimates under-estimating volatility when plotted against their Garch counterparts. One of the best visual matches of the different estimates is that obtained for the Italian index. But even here, the evidence of under-estimation is still strong.

There are obviously major differences between the two popular statistical methodologies. The empirical results obtained suggest that a good, and computationally inexpensive estimate of volatility (risk) can be obtained using the exponential smoothing approach. However, at times, especially following bouts of high activity in markets, more accurate measures can be obtained using Garch. End users of these methods must balance the cost and effort required in obtaining Garch estimates of volatility against the accuracy required. ■

TABLE 1. DIAGNOSTIC ON ES VOLATILITY MODEL

	a	b	R**2	Q-stat
US	12.70	0.75	0.317	38.12
	3.26	1.97		0.00
UK	19.99	0.77	0.336	2.34
	2.61	1.53		0.88
France	53.42	0.57	0.341	19.37
	3.57	2.59		0.01
Japan	40.65	0.88	0.300	5.89
	2.70	1.08		0.44
Italy	39.80	0.82	0.325	79.36
	3.02	1.52		0.00
Hong Kong	36.51	0.92	0.320	9.17
	2.19	0.66		0.16

TABLE 2. DIAGNOSTIC ON GARCH VOLATILITY ESTIMATORS

	a	b	R**2	Q-stat
US	-5.68	1.18	0.207	2.6
	0.85	0.89		0.86
UK	-14.38	1.25	0.345	3.65
	1.32	1.33		0.72
France	4.81	1.05	0.348	3.91
	0.22	0.25		0.69
Japan	3.04	0.99	0.326	4.88
	0.19	0.09		0.56
Italy	2.56	0.99	0.327	5.49
	0.16	0.00		0.48
Hong-Kong	12.19	1.07	0.349	11.90
	0.58	0.57		0.064

This article benefited greatly from suggestions made by members of the Risk Management team of Credit Suisse in Zurich. Kostas Giannopoulos is from University of Westminster, and Brain Eales is from London Guildhall University

FOOTNOTES

[1] It should be noted that if, in order to simplify computation, the base forecast is that prices are not going to change then returns, Y_t, and surprise, ε_t, will be equal.

[2] There are two versions of this model used to estimate volatility. One version follows the J. P. Morgan methodology of using today's market change as an input, the second, used by Engle in his article in Risk magazine (September, 1995), uses yesterday's market change as an input. In this study we will use the later version since uses the same set of information with the GARCH model.

[3] Maximisation of the likelihood function is the computational price that is involved in using GARCH. The ES model only involves one parameter whose optimum value can be obtained by selecting that estimate which generates the minimum sum of squared residuals.

[4] In view of the fact that volatility is the variance error associated with every forecast, Y_t is used in (1), which are the residuals obtained from the exponential moving average approach, as explained in the 2nd edition of J. P. Morgan's RiskMetrics™.