

## **History Debunked**

*Carol Alexander, University of Sussex, UK*

Correlation is one of the most widely used statistics in financial analysis - and it is also one of the most frequently misunderstood. So this article attempts to put the record straight about the use of standard correlation measures and to describe some of the newer measures of 'association' between financial variables. Broadly speaking, the basic methods of risk analysis use *constant parameter* assumptions - constant volatilities and correlations; constant betas and deltas - and these assumptions greatly simplify the theoretical models. However, not only are such assumptions very unrealistic, they can lead to all sorts of contradictions.

A case in point is the development of *capital at risk* models since the capital adequacy directive of the European Union was issued in August 1993. There is no single capital at risk model recommended by the BIS, and a wide variety of results are being obtained from institutions which have developed their own (see Lillian Chew's article in *RISK*, September 1994). Thomas Wilson (*RISK*, November 1994) describes many of the better known variants of capital at risk models, all of which use *historical* correlation at some stage. But one of the points of this article is to show how very misleading 'historical' measures of correlation can be, and the diverse results being obtained from capital at risk models might be attributable to the different uses of historical correlations in these models.

Another point of this article is to recommend, instead of historical correlation, the use of certain *dynamic* measures of association in risk management applications and in particular in capital at risk models (see Box 2). But first let's re-examine the standard ways in which correlation is measured:

## Historical Correlation

The most basic method of estimating correlations between returns on financial time series is termed *historical* (backward looking) or *actual* (forward looking). There are (at least) four problems with this measure:

- As it is a purely statistical measure, it contains little of the information about investors expectations which is encapsulated in derivatives prices;
- Since historic correlations do not change if the sample is re-ordered (in pairs) the information contained in the dynamic ordering of observations is completely ignored by historic correlation. As such we cannot expect it to be either a good measure of current correlation, or a good predictor of future correlation<sup>1</sup>;
- It is a contradiction in terms to calculate a *time-varying* series for historic/actual correlation - or for that matter historic/actual volatility. 'Historic' and 'actual' series are sample estimates of the *unconditional moments of stationary series*, which are constant by definition (see below). Indeed, assuming stability of the underlying data generation processes, all the variation in historic/actual correlations (and volatilities) comes from sampling error alone!
- Historic correlation is a function of maturity: if a highly significant event occurs on one particular day, it will remain in the volatility or correlation series for exactly n days (n being the maturity of the volatility/correlation) and on the n+1th day a structural change will appear in the series, even though nothing occurred on that day. Figure 1 shows historical correlations between the US and UK FT indices in the period following the October '87 crash. This event remains in the 30-day series for exactly 30 days, the 60-day series for exactly 60 days and so on ... but nothing particular happened on the 31st day after the crash!

---

<sup>1</sup>Historic volatility also ignores the information from dynamic ordering of observations within the samples, so it is not surprising that it is such a terrible predictor of future volatility.

## Implied Correlation

In some cases it is possible to calculate *implied* correlations from implied volatilities (for example in currency markets); or to invert pricing formulae of second order correlation products (such as quantos) and hence estimate implied correlations; or to approximate implied correlation using constancy assumptions, (for example in equity indices). Problems with the implied measure of correlation include:

- It isn't always possible to obtain even approximate implied correlation series;
- When approximations must be used, such as the approximate equity implied correlations described in Michael Kelly's recent article (*RISK*, August 1994), these approximations can lead to implied correlations greater than one in absolute value;
- Implied correlations are model dependent, and as such they inherit the usual problems of bias and negative autocorrelation in Black-Scholes implied volatility series;

## GARCH correlation

GARCH models of volatility and correlation provide *instantaneous* measures, and are a correct statistical measure to use for a time-varying series. GARCH is a flexible and transportable forecasting tool: having estimated the model, the GARCH series can very easily be forecast for  $n$  days ahead ( $n = 1, 2, 3, \dots$ ) and the forecasts aggregated to produce 10-day, 30-day, 60-day or any horizon forecasts which are then directly comparable to implied counterparts (see Box 1).

Problems peculiar to GARCH correlation include:

- the need for very parsimonious parameterizations to avoid convergence difficulties associated with flat likelihood functions;
- lack of robustness to starting values (even in some well known econometric packages) so that certain GARCH models should not be relied upon.

## The Importance of Stationarity.

However the limitations of each of the correlation models just described are quite insignificant when compared to a methodological defect which is commonly encountered in the computation of correlations. For correlations to be statistically meaningful it is necessary that the two underlying series be *jointly covariance stationary*<sup>2</sup>. Now, whilst it is usually reasonable to expect individual asset returns to be covariance stationary<sup>3</sup>, it is *not* always the case that *joint* stationarity will be automatically achieved. For example, we may expect that term structure returns will be jointly stationary because there are strong links between interest rates, but why should an interest rate or equity index be jointly covariance stationary with an exchange rate, such as is assumed in the pricing of second order correlation products? Stable pricing of first order products also requires joint covariance stationarity, for example between two FX rates, or two equity indices, and this assumption may not be empirically justified.

Non-sensical results will obtain if correlations are calculated on two series which are not jointly covariance stationary. Figure 2 illustrates two correlation series of daily returns calculated from a bivariate normal GARCH(1,1) model (see Box 1). In (a) the two returns series are jointly covariance stationary, and in (b) they are not. It does not take much to observe that predictions based on the first figure are likely to be relatively stable, whilst those based on the second will be very suspect indeed! Thus if historical correlations are calculated on non-jointly covariance stationary series they will *appear* to be very unstable, because theoretically they do not exist!<sup>4</sup>

---

<sup>2</sup>Let  $\{x_1, \dots, x_T\}$  be a time series, so  $x_t$  is a random variable which denotes the value of  $x$  at time  $t$ . Then  $x$  is said to be *covariance stationary*, or stationary in the *weak* or *wide* sense if

- $E(x_t)$  is the same for all  $t$  ( so  $E(x_t)$  is a finite constant);
- $V(x_t)$  is the same for all  $t$  ( so  $V(x_t)$  is a finite constant);
- The correlation between  $x_t$  and  $x_s$  is the same for all  $t$  ( so  $COV(x_t, x_s)$  depends only on the distance apart of the two points,  $t-s$ ).

Now consider two time series  $x$  and  $y$ , both of which are individually covariance stationary. Then they are *jointly covariance stationary* if

- The correlation between  $x_t$  and  $y_s$  is the same for all  $t$  ( so  $COV(x_t, y_s)$  depends only on the distance apart of the two points,  $t-s$ ).

<sup>3</sup>Hence calculation of historical/GARCH volatilities on asset returns is not usually subject to this methodological pitfall.

<sup>4</sup>This type of problem is often encountered in the naive application of time series analysis to financial data. For example calculating the historic volatility of a random walk is also meaningless since, theoretically, it is infinite.

### **Losing Information About the Trends.**

Even though GARCH is a vast improvement upon historical correlation as a statistical measure, *all* measures of correlation - implied and statistical - ignore valuable information: since they must be calculated on stationary returns series, correlations contain no information about the trends in the data. All the underlying price or rate series are *individually* converted into non-trended series, so that any information about *common* stochastic trends is automatically lost before correlations are calculated. Thus any correlation series - implied, historic/actual or GARCH - ignores the possibility that trends might also be associated.

In a recent article (RISK, February 1994) I describe a measure of association between time series which, following the classic paper by Robert Engle and Clive Granger in 1987, has taken the international economic community by storm. This measure, which is relatively new to the financial world, is called *cointegration* and is just one of many *common features* between time series, which are now being pioneered by Robert Engle. Cointegration occurs in a set of two or more time series if (a) each series is a random walk, and (b) some weighted sum of these series is stationary<sup>5</sup>. Thus although each series is unpredictable when taken by itself, the stationarity property of cointegrated series means that they cannot drift too far apart from each other. Cointegrated series may show little correlation in the short term, but they are 'tied together' in the longer term and it is this property which opens up dynamic forecasting opportunities.

One of the most appealing aspects of cointegration is that it loses none of the dynamic information, nor information about the trend in data. In fact, since cointegration analysis is carried out on the basic random walks which underlie returns data, one could describe cointegration as the multivariate analysis of random walks<sup>6</sup>. Whilst it makes no sense to build a univariate model

---

<sup>5</sup>This description is greatly simplified - see Engle and Granger's standard reference (footnote 11) for the general definition of cointegration.

<sup>6</sup>A time series  $x$  follows a random walk (without drift) if the best prediction of tomorrow's value is the value today. This implies a statistical model of the form

$$x_t = x_{t-1} + \varepsilon_t$$

where  $\varepsilon_t$  is a stationary error term. Thus  $\Delta x_t = \varepsilon_t$  is stationary, but  $x$  itself is not stationary since it has infinite (unconditional) variance ... this means that the 'walk' could go anywhere, given enough time. Adding a constant to

of a random walk, since the best prediction of any future value is simply the value today, there are enormous gains to be made from using cointegration to model the common stochastic trend of several random walks.

### Using and Modelling Cointegration

Uses of cointegration to model the yield curve, and to detect market inefficiencies are described in the RISK article mentioned above. Pick up any recent issue of the Journal of Futures Markets and you will probably find an article which employs cointegration to investigate the price discovery relationship between cash and derivatives markets<sup>7</sup>. Cointegrated volatilities - or *copersistence in variance* as Tim Bollerslev and Robert Engle name it<sup>8</sup> - are fundamental to the success of the GARCH dynamic hedge for currency portfolios<sup>9</sup>. Multivariate methods for forecasting currency volatility can be based on cointegration, equity index arbitrage can be made more efficient and dynamic hedging is more effective when a combination of cointegration and GARCH techniques is employed<sup>10</sup>. In many of these applications the elements of the cointegrating vector have natural interpretations, such as portfolio weights or option and factor deltas. And cointegration modelling by the Engle-Granger methodology requires nothing more than an ordinary least squares algorithm on the computer<sup>11</sup>

---

the right hand side of the above equation yields the random walk 'with drift' model, so that  $x$  'leans' in a direction given by the constant and the time series will have a trend in it. As before, the random walk 'with drift' is only stationary after first differencing.

<sup>7</sup>For example see Wang, G.H.K. and Jot Yau (1994) "A time series approach to testing for market linkage: unit root and cointegration tests" *Journal of Futures Markets* **14** : 4 pp457-474.

<sup>8</sup>Bollerslev, T. and R.F. Engle (1993) "Common persistence in conditional variances" *Econometrica* **61**:1 pp167-186.

<sup>9</sup> Ghose, D. and K.F. Kroner (1993) "Common persistence in conditional variances: implications for optimal hedging". Department of Economics, University of Arizona, Tucson, USA.

<sup>10</sup>These and other applications of cointegration are covered in the RISK conference on Correlation.

<sup>11</sup>Engle, R.F. and C.W.J. Granger (1987) "Co-integration and error correction: representation, estimation, and testing". *Econometrica* **55**:2 pp 251-76.

Johansen's methodology for investigating cointegration in a multivariate system<sup>12</sup> is commonly regarded as superior to the Engle-Granger method, particularly when the number of variables is greater than two. The Johansen tests, which are based on the eigenvalues of a stochastic matrix, can be viewed as a canonical correlation problem similar to that of principal components. In fact cointegration can be thought of as a generalisation of principal components analysis<sup>13</sup>. Although the Johansen tests are more difficult to program than the Engle-Granger method, they suffer less small sample bias and their power function has better properties. The objectives of the two tests also differ - the Johansen tests seek the linear combination which is most stationary, whilst the Engle-Granger tests, being based on ordinary least squares, seek the linear combination having minimum variance, and if that combination is stationary the variables are taken to be cointegrated. However for financial applications of cointegration there are several reasons why one could happily stick to the Engle-Granger methodology which, as already mentioned, is very straightforward to compute:

- the Engle-Granger small sample bias is not necessarily going to be a problem in financial analysis since sample sizes are generally quite large and the cointegrating vector is super consistent;
- there is often a natural choice of dependent variable in the cointegrating regressions (for example, in equity index arbitrage);
- in risk management applications it is generally the Engle-Granger criterion of *minimum variance*, rather than the Johansen criterion of *maximum stationarity*, which is paramount.

Here is a simplified algorithm for the Engle-Granger approach to cointegration that can be performed on a good spreadsheet .

---

<sup>12</sup>Johansen, S (1988) "Statistical analysis of cointegration vectors". *Journal of Economic Dynamics and Control* **12** pp 231-54.

<sup>13</sup>The connection between these two methodologies is that a principal component analysis of cointegrated variables will yield the common stochastic trend as the first principal component. But the outputs of the two analyses differ: principal components gives two or three series which can be used to approximate a much larger set of series (such as the yield curve); cointegration gives all possible stationary linear combinations of a set of random walks. Also

**Step 1: Augmented Dickey Fuller tests for stationarity:** Take first differences of  $x$ , and call this series  $\Delta x$  (so we often perform the test on  $\log(x)$  because the first differences are then approximately equal to returns); Do an ordinary least squares (OLS) regression of  $\Delta x$  on the following three variables: a constant, one lag of  $\Delta x$  and one lag of  $x$ ; Look for autocorrelation in the residuals. If residuals are autocorrelated then include one, two, three, four, ... however many lags of  $\Delta x$  it takes to remove this autocorrelation. (In financial data it is not often necessary to include more than one or two lags); The ADF statistic for the null hypothesis  $H_0$  'x is stationary only after first differencing' versus the alternative hypothesis  $H_1$  'x is already stationary' is the t-ratio on the coefficient of lagged  $x$  in the ADF regression; Compare the ADF statistic with the 5% critical value of -2.9. If it is negative and less than -2.9 then you can reject  $H_0$  at 5%: otherwise go back to step 1 but read  $\Delta x$  for  $x$  (so the ADF regression involves  $\Delta^2 x$  and  $\Delta x$  and the ADF statistic is the t-ratio on the lag of  $\Delta x$ ); Repeat until  $H_0$  is rejected. If  $\Delta x$  is stationary (so  $H_0$  is not rejected by the first ADF test on level  $x$ , but is rejected second time round, by the ADF test on first differences) then  $x$  is a random walk.

**Step 2: Testing for cointegration** Use ADF tests, as described in step 1, to ensure that each of the variables used in cointegration testing is a random walk; Choose any one of the variables as the dependent variable in an OLS regression on a constant and all the other variables; Do the OLS regression (called the 'cointegrating regression') and take the residuals; Do an ADF test for stationarity of these residuals, exactly as explained in step 1 above, but using -3.6 as the 5% critical value; The more stationary the residuals from the cointegrating regression, the stronger the cointegration between the variables; If the variables are cointegrated, then the residuals are the 'cointegrating vector' (i.e. a linear combination of the variables which is stationary) so the coefficients in the cointegrating vector are taken as the OLS estimated coefficients.

**Step 3: Building an error correction model (ECM)** If two or more time series are cointegrated then the usual sort of vector autoregression on first differences (which follows the Box-Jenkins ARIMA methodology of the late 1970's) is not correct. Instead we must model the relationship between the variables using what is called an error correction model (ECM) and it is in this framework that we can test for the direction(s) of causal flow between the variables. For two cointegrated variables  $x$  and  $y$  the ECM takes the form of two regression equations, one for  $\Delta x$  and one for  $\Delta y$ , which can be independently estimated by OLS. These regressions, which involve only lagged stationary variables, including the cointegrating vector  $z$ , are:  $\Delta x$  regressed on a constant, lagged  $\Delta x$ , lagged  $\Delta y$  and one lag of  $z$  - and  $\Delta y$  regressed on a constant, lagged  $\Delta x$ , lagged  $\Delta y$  and one lag of  $z$ . Thus if the logs of asset prices are investigated for cointegration, the resulting ECM gives us a forecasting model for asset returns.

---

principle components can be used to analyse *any* set of variables, whereas cointegration is confined to random walks.

### Cofeatures and Granger causality

Having pioneered both GARCH and cointegration, Robert Engle has turned attention to a new 'common features' or *cofeatures* approach to measuring association between financial time series<sup>14</sup>. Cointegration occurs when the series share a common stochastic trend, but serial correlation and ARCH volatilities are examples of other features which may be common to a set of time series<sup>15</sup>. Finding evidence of a common feature indicates the existence of (at least) one Granger causal flow in the system - that is, turning points in one of the variables will precede turning points in the others<sup>16</sup> - and hence can lead to forecasting and risk management opportunities which have hitherto been largely unexplored. For example a measure of market to market risk can be based on the Granger causality statistic from rolling bivariate error correction models<sup>17</sup>.

There are several problems with the use of correlation to measure association between variables - not the least of which is that correlations can appear to be very unstable because theoretically they do not always exist - and even when they do exist they ignore much valuable information. But now that the financial world is becoming more aware of the potential of cointegration and other 'common feature' methodologies for measuring association between variables, we should expect to see an increasing number of applications.

---

<sup>14</sup>The classic reference for this is Engle, R.F. and S. Kozicki (1993) "Testing for Common Features". *Journal of Business and Economic Statistics*, **11**, pp 369-95 (with discussions).

<sup>15</sup>Alexander, C.O. (1995) "Common volatility in the foreign exchange market" in *Applied Financial Economics*, **5**, No. 1.

<sup>16</sup>Granger, C.W.J. (1986) "Developments in the study of cointegrated economic variables". *Oxford Bulletin of Economics and Statistics* 42:3 pp 213-27.

<sup>17</sup>Alexander, C.O. (1994) "Cofeatures in international bond and equity markets" *University of Sussex Discussion paper in Economics No. 94/01*.

## BOX 1: GARCH MODELS

Recall the idea of *linear regression* is to find the 'line' of best fit to a series. For example we could use an *autoregressive* model to describe current returns  $r_t$  as functions of past returns, and when just one lag is used we have the AR(1) model

$$r_t = \phi_1 + \phi_2 r_{t-1} + \varepsilon_t$$

Such an equation is called the *conditional mean* equation because the fitted values give a *time-varying* series which can be used as an estimate of the conditional mean (as opposed to the *unconditional* mean, which is a constant by definition).

The idea of GARCH is to add a second equation to this conditional mean equation which describes the *conditional variance* of the series. For example in the GARCH(1,1) model we have

$$s_t^2 = a_0 + a_1 e_{t-1}^2 + b_1 s_{t-1}^2 \quad a_0 > 0, a_1, b_1 \geq 0$$

where  $s_t^2$  denoted the conditional variance of  $r_t$ . Its fitted values are square rooted and annualised in the usual way to give the GARCH volatility of the series. Since by definition GARCH volatility is non-constant, it is the correct statistical measure to use for time-varying volatility, but it is only valid when  $r_t$  is a stationary series.

When two returns series  $r_{1t}$  and  $r_{2t}$  are known to be jointly stationary we can generalise the univariate GARCH model above to a stable bivariate GARCH model, which could take the following form:

*Conditional mean equations*

$$r_{1,t} = \phi_{11} + \phi_{12} r_{1,t-1} + \varepsilon_{1,t}$$

$$r_{2,t} = \phi_{21} + \phi_{22} r_{2,t-1} + \varepsilon_{2,t}$$

*Conditional variance equations*

$$s_{1,t}^2 = a_{10} + a_{11} e_{1,t-1}^2 + b_{1,1} s_{1,t-1}^2$$

$$s_{2,t}^2 = a_{20} + a_{21} e_{2,t-1}^2 + b_{2,1} s_{2,t-1}^2$$

$$s_{12,t}^2 = a_{30} + a_{31} e_{1,t-1} e_{2,t-1} + b_{3,1} s_{12,t-1}^2$$

Now the conditional variance is a symmetric 2x2 matrix with diagonal elements equal to the conditional variances  $\mathbf{s}_{1t}^2$  and  $\mathbf{s}_{2t}^2$  and off-diagonals equal to the conditional covariance  $\mathbf{s}_{12,t}$ . The positivity constraints on the parameters in univariate GARCH model are generalised to ensure *positive definiteness* of this conditional variance matrix, then GARCH volatilities are obtained by annualising the square rooted fitted variance series as before, and GARCH correlations are given by  $\hat{\mathbf{s}}_{12,t} / \sqrt{\hat{\mathbf{s}}_{1,t}^2 \hat{\mathbf{s}}_{2,t}^2}$ .

One of the big advantages of GARCH models over neural networks and other embedding based predictor models is that GARCH forecasts are independent of a training data set, so they are easily transportable. To calculate an n-period GARCH forecast, which will be directly comparable to n-period implied or historic correlations (or volatilities), you need to use a recursive formula such as:

$$\hat{\mathbf{s}}_{t+1}^2 = \hat{\mathbf{a}}_0 + \hat{\mathbf{a}}_1 \mathbf{e}_t^2 + \hat{\mathbf{b}}_1 \mathbf{s}_t^2 : \hat{\mathbf{s}}_{t+s}^2 = \hat{\mathbf{a}}_0 + (\hat{\mathbf{a}}_1 + \hat{\mathbf{b}}_1) \hat{\mathbf{s}}_{t+s-1}^2$$

which yields s-step ahead conditional variance forecasts for s = 1 to n. These are aggregated and then converted to annualised volatilities in the usual way. Similarly conditional covariance forecasts are obtained and converted to conditional correlation forecasts, which are averaged for direct comparison with historic and implied counterparts. Another way in which GARCH models may be used for forecasting is to use the standard errors on each coefficient to obtain time-varying confidence intervals, such as those used in Saloman Brothers' GIFT (Garch Index Forecasting Tool).

The above example is a 'bivariate GARCH(1,1)-AR(1) model with a diagonal vech parameterizations', which is the most basic formulation of a GARCH correlation model. Modifications can be made in both conditional mean and conditional variance equations:

- The conditional means can be extended to include error correction terms, a modification which is often necessary when using GARCH for dynamic hedging since it will be more effective when asset and hedge are cointegrated;

- Alternatively it may be necessary to reduce the number of parameters in the model if convergence problems arise from flat likelihoods - a common problem in the maximum likelihood estimation of GARCH models. In this case the lagged returns might be excluded from the conditional mean equations;
- Although correlation is a two-dimensional concept, it may be necessary to estimate multivariate GARCH models with more than the two dimensions illustrated to ensure positive definiteness of the  $n \times n$  conditional covariance matrices. However, bivariate GARCH estimates may be robust to higher dimensions (particularly with the diagonal vech conditional variance parameterizations) in which case the  $n \times n$  covariance matrices can be inferred from a series of estimated  $2 \times 2$  matrices;
- The basic diagonal vech parameterizations has substantial cross equations restrictions - for example, the conditional variances has no lagged effect on the conditional covariances - but is commonly used because we need a very parsimonious parameterizations of the model to avoid convergence problems in maximum likelihood estimation. A more general parameterisation may be used, such as the 'BEKK' parameterizations<sup>18</sup> which ensures positive definiteness with the minimum number of parameters. However, there are 11 parameters in the conditional variance equations (as compared to 9 in the diagonal vech model) and in higher dimensions the number of extra parameters required rapidly increases.

---

<sup>18</sup>See Engle, R.F. and Kroner, K "Multivariate simultaneous generalized ARCH" USCD discussion paper, June 1993.

BOX 2: A DELTA NORMAL GARCH CAPITAL AT RISK MODEL

Capital at risk is defined as *the maximum possible loss from market movements over 24hrs* so the relationship between capital at risk (CAR) and the variance of the portfolio is often defined as

$$\text{CAR} = \alpha \sqrt{V(\Delta P)} \quad (1)$$

where  $\Delta P$  denotes the daily change in price of the portfolio and  $\alpha$  gives the required confidence level<sup>19</sup>.

The 'historical' method for evaluating CAR which was originally recommended by the BIS is very computationally intensive. So a wide variety of different CAR models are being developed, producing some very different results. Part of the problem is that when correlations are not ignored, *historical* correlation matrices are being used. But if CAR models based on historical correlations perform as well as forecasts based on historical correlations, it is not surprising that everybody is getting different answers ... all variation in the historical model can only be attributed to sampling errors, and these sampling errors are compounded in a variety of ways in different CAR models!

For example, in the asset-normal method consider a portfolio of  $n$  assets  $A_1, \dots, A_n$  in proportions  $w_1, \dots, w_n$ . If we write these as  $n \times 1$  column vectors  $\mathbf{A}$  and  $\mathbf{w}$  we have the dot product  $P = \mathbf{w}'\mathbf{A}$ , whence  $V(\Delta P) = \mathbf{w}'\mathbf{V}(\Delta\mathbf{A})\mathbf{w}$  where  $\mathbf{V}(\Delta\mathbf{A})$  is the  $n \times n$  historical covariance matrix of asset returns. Compare this with delta-normal methods which are based on a factor delta decomposition of the portfolio  $\Delta P = \boldsymbol{\delta}' \cdot \Delta \mathbf{F}$  where  $\mathbf{F}$  is a vector of risk factors and  $\boldsymbol{\delta}$  is a vector of net deltas with respect to these risk factors (so  $\delta_i = \sum w_j \partial A_j / \partial F_i$ ). From this representation we can calculate the variance of the portfolio innovations as  $V(\Delta P) = \boldsymbol{\delta}' \mathbf{V}(\Delta \mathbf{F}) \boldsymbol{\delta}$  where  $\mathbf{V}(\Delta \mathbf{F})$  is the historical covariance matrix between factors. Although historical *variances* (the diagonal elements of this matrix) are not necessarily used - in 'factor-push' methods for instance - the off-diagonal historical *correlations* are. It is clear that error in the use of historical correlations will be

---

<sup>19</sup>In 'asset-normal' CAR models we assume that  $\Delta P$  is normally distributed, so that  $\alpha = 1.96$  for a 95% level of confidence.

compounded so as to produce very different results from the two types of CAR calculations, especially when assets and/or factors are highly correlated.

A dynamic correlation approach to the problem can be based on either asset-normal or factor-delta representations of a portfolio, by using GARCH dynamic correlations in the resulting formulae for the variance of portfolio innovations: for example, using a factor-delta representation we have

$$V_t(\Delta P_t) = \sum \delta_i^2 V_t(\Delta F_{it}) + \sum \sum \delta_i \delta_j \text{COV}_t(\Delta F_{it}, \Delta F_{jt}) .$$

Clearly the factor-delta representation is preferable since it reduces the dimensionality of the GARCH covariance matrix, but when there are more than two or three risk factors the problem of parsimonious parameterizations will have to be dealt with, particularly when risk factors are cointegrated (since each of the conditional mean equations should contain an extra cointegrating vector term).

The use of principal components analysis will help enormously to reduce the dimensionality of the required GARCH models<sup>20</sup>. For example, in fixed interest portfolios: these might be well approximated by the first two principal components of a yield curve (the 'trend' and 'tilt'); similarly large international swaps portfolios can be approximated using the first principal component (i.e. the trend) of each term structure as risk factors. So our starting point is a principle components representation of the portfolio as  $P = \boldsymbol{\phi}'\mathbf{PC}$  , where  $\mathbf{PC}$  is a vector of  $n$  principal components of the portfolio (risk factors), and  $\boldsymbol{\phi}$  is a vector of net weights on these components (i.e. the products of principal component factor weights with the factor deltas of the portfolio). The number of components is chosen as small as possible to capture 'sufficient' variation in the risk factors (or portfolio assets if the asset normal model is used). Although the principal components are, by construction, orthogonal and therefore uncorrelated, their first differences *are* correlated. Hence the time-varying portfolio variance is

$$V_t(\Delta P_t) = \sum \phi_i^2 V_t(\Delta PC_{it}) + \sum \sum \phi_i \phi_j \text{COV}_t(\Delta PC_{it}, \Delta PC_{jt}) \quad (2)$$

with, in general, non-zero covariance terms.

---

<sup>20</sup>See Thomas Wilson's excellent article 'Debunking the Myths', *RISK*, April 1994.

Now a dynamic CAR model can be obtained from *either*  $\alpha \sqrt{V_t}(\Delta P_t)$  where  $\alpha$  is the appropriate critical value, *or* by using time-varying upper and lower confidence limits for the GARCH portfolio variance. Also, N-period forecasts of CAR can now be obtained using the methods described in Box 1.