Don't look back

Historical simulation may be a natural setting for scenario analysis, but it must take account of current market conditions, caution Giovanni Barone-Adesi, Frederick Bourgoin and Kostas Giannopoulos

Value-at-risk is becoming increasingly popular as a management and regulatory tool. But before this acceptance goes much further, we need to assess its reliability under financial market conditions. Most VAR models deal either with the non-normality of security returns or with their conditional heteroscedasticity, but not with both. We are developing a modified historical simulation approach that allows for both effects.

Historical simulation relies on a specific distribution (usually uniform or normal) to select returns from the past. These returns are applied to current asset prices to simulate their future returns. Once enough different paths have been explored, it is possible to determine a portfolio VAR without making arbitrary assumptions about the distribution of portfolio returns. This is especially useful where there are abnormally large portfolio returns.

It is well known that large returns cluster in time (see, for example, Mandelbrot, 1963, and Black, 1976). The resulting fluctuations in daily volatility make the confidence levels of some VAR calculations unreliable (Boudoukh *et al*, 1995). This is the case with those that ignore clustering, such as VAR measurements based on the standard variance-covariance matrix and Monte Carlo methods, which typically ignore current market conditions to produce flat volatility forecasts for future days. Moreover, the use of the covariance matrix of security returns or the choice of an arbitrary distribution in the Monte Carlo method usually destroys valuable information about the distribution of portfolio returns.

To make our historical simulation consistent with the clustering of large returns, we model the volatility of our portfolio as an asymmetric Garch (Generalised autoregressive conditional heteroscedasticity) process (Engle & Ng, 1993) that generalises the Garch model. This model allows positive and negative returns to have different impacts on volatility (known as the leverage effect, see Black, 1976). Past daily portfolio returns are divided by the Garch volatility estimated for the same date to obtain standardised residuals. These are independent and identically distributed (IID) and are therefore suitable for historical simulation.

To adjust them to current market conditions, we multiply a randomly selected standardised residual by the Garch forecast of tomorrow's volatility. In this way, a simulated portfolio return for tomorrow is obtained. This simulated return is used to update the Garch forecast for the following day, which is then multiplied by a newly selected, standardised residual to simulate the return for the second day. Our recursive procedure is repeated until the VAR horizon (ie, 10 days) is reached, generating a sample path of portfolio volatilities and returns. We repeat our procedure to obtain a batch of sample paths of portfolio returns. A confidence band for the corresponding portfolio values is built by taking the kernel (empirical) frequency distribution of values at each time. The lower 1% area identifies the worst case over the next 10 days.

To illustrate our procedure, we constructed a hypothetical portfolio, diversified across all 13 national equity markets in our data sample. To form our portfolio, each equity market is weighted in proportion to its capitalisation in the world index (MSCI) as at December 1995. The portfolio weights are reported in table A.

These weights are held constant for the entire 10-year period and multiplied by the 13 local index returns. So the portfolio returns are calculated again backwards to reflect the current weightings. Since the aim of market risk is to quantify eventual portfolio losses in a single currency, all local portfolio returns are measured in dollars. The descriptive statistics, together with the Jarque-Bera (1980) test for normality, are shown in table B, where the p-value indicates the probability that our portfolio returns are generated from a normal distribution.

Figure 1 shows the empirical distribution of the portfolio's returns. The rejection of normality in table A and the pattern of clustering visible in figure 1 leads us to model our portfolio returns, r_t, as a Garch process with asymmetries, with daily volatility, h_t, given by:

$$r_t = R_t = \mu + \varepsilon_t \quad (1a)$$

$$h_t = \omega + \alpha(\varepsilon_{t-1} + \gamma)^2 + \beta h_{t-1} \quad (1b)$$

The variance for small increments on the other end can be written as:

$$h_t^2 = c_t^2 \Delta t = O(\Delta t)$$

The daily return in equation (1a) is the sum of each expected value, μ, plus a random residual, ε_t. Because of the small, statistically insignificant value of μ^2, this term will be neglected in the calculation of daily volatilities.[1] Equation (1b) defines the volatility of ε_t, h_t, as an asymmetric Garch process. h_t is the sum of a constant, ω, plus two terms reflecting the contributions of the most recent "surprise", ε_{t-1}, and the last period's volatility, h_{t-1}. Finally, γ allows for the asymmetric response of the innovation on the volatility and is statistically significant.

Therefore, our portfolio volatility is modelled to depend on the most recently observed portfolio returns. The combination of asymmetric Garch volatility and portfolio historical returns offers us a fast and accurate

A. Portfolio weights

Country	Our portfolio	World index (Dec 1995)
Denmark	0.004854	0.004528
France	0.038444	0.035857
Germany	0.041905	0.039086
Hong Kong	0.018918	0.017645
Italy	0.013626	0.012709
Japan	0.250371	0.233527
Netherlands	0.024552	0.022900
Singapore	0.007147	0.006667
Spain	0.010993	0.010254
Sweden	0.012406	0.011571
Switzerland	0.036343	0.033898
UK	0.103207	0.096264
US	0.437233	0.407818

B. Descriptive statistics of the equally weighted portfolio

Mean (pa)	10.92%	Std dev (annual)	12.34%
Skewness	−2.82	Kurtosis	62.36
Normality	3,474.39	p-value	0

[1] *In fact, for stock prices, μ^2 is in the order of $\mu^2\Delta_t$:*

$$\mu^2 = c^2 \Delta t^2 = O(\Delta t^2)$$

1. World capitalisation weighted portfolio returns: Jan 1985–Feb 1996

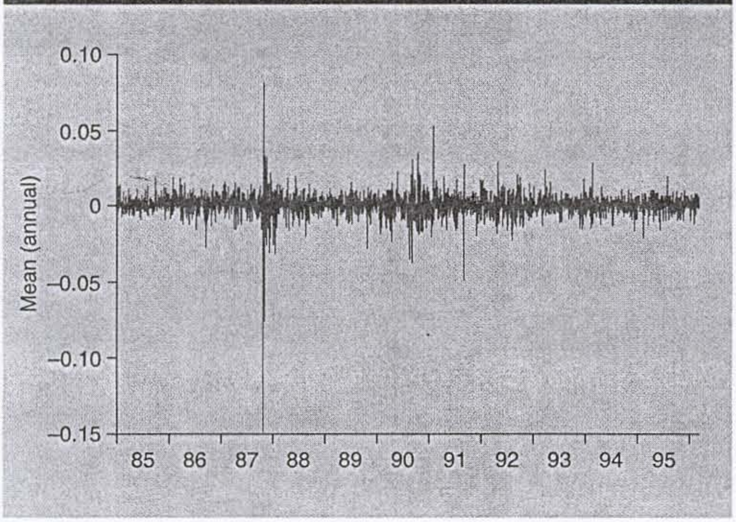

measure of the past, current and future volatilities of the current portfolio. We do not need the correlation matrix of security returns. Furthermore, our VAR method contains fewer "unpleasant surprises", since Garch models allow for fat tails on the unconditional distribution of the data.[2] The effects of our choice become apparent if we compare the returns in figure 1 with those in figure 2, where they have been scaled by their daily volatility, so that:

$$z_t = \frac{r_t}{\sqrt{h_t}} \tag{2}$$

Clustering of returns is reduced by volatility scaling, so the distribution of returns now appears to be more uniform, making the historical simulation more appropriate. However, the large number of returns still exceeding three standard deviations suggests that our scaling does not make returns normal. Our annualised portfolio volatility, shown in figure 3, varied from 7% to 21% over the 10-year period.

The scaled returns are the foundation of our simulation. To simulate portfolio returns over the next 10 days we select randomly 10 returns from figure 2 using the "bootstrap" methodology developed by Efron & Tibshirani (1993). We then construct iteratively the daily portfolio volatility that these returns imply according to equations (1a) and (1b) and use this volatility to rescale our returns. The resulting returns therefore reflect current market conditions, rather than the market conditions associated with returns in figure 1. In other words, we simulate future standardised residual returns as a random vector Θ of outcomes from a stationary distribution. The possible outcomes of the stationary distribution are the historical residuals, standardised by the corresponding daily volatility:

$$\varepsilon^* = \left\{\varepsilon^*_{ij} = \Theta\right\}, \quad \Theta = \{r_1, r_2, \ldots, r_T\}$$

[2] *For a Garch process with conditional normality, the excess of kurtosis of the unconditional distribution of the process is greater than three. See Bollerslev (1986)*

2. Portfolio stress analysis (standardised residuals): Jan 1986–Dec 1995

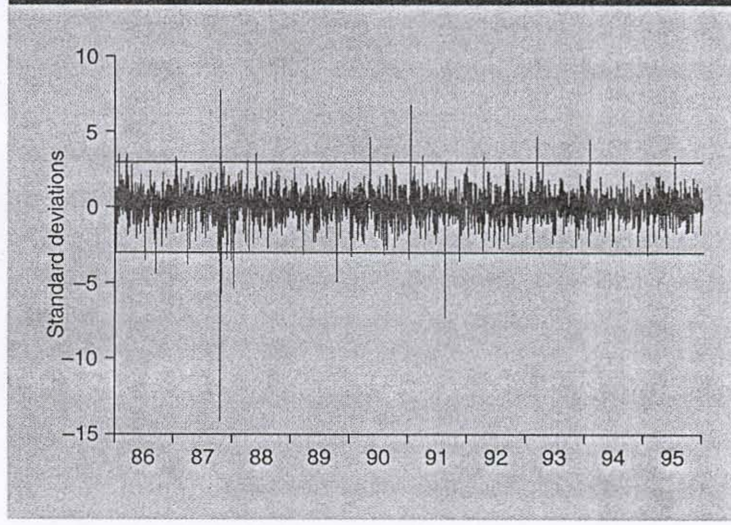

3. Annualised volatility of the portfolio: Jan 1986–Dec 1995

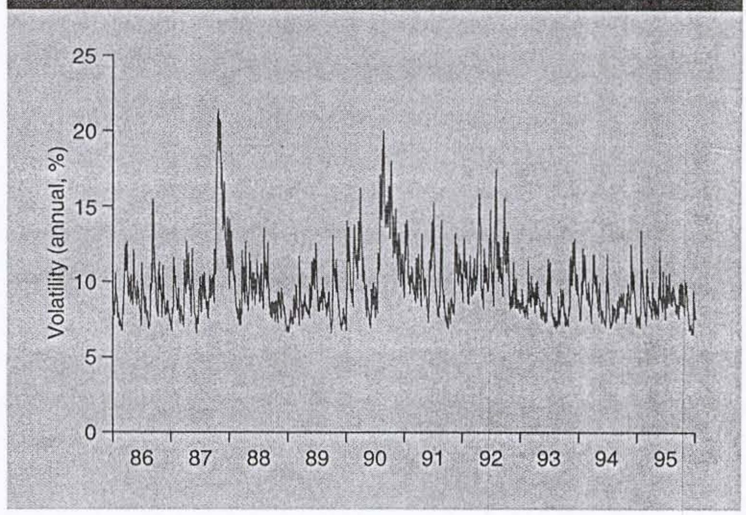

where $i = 1,\ldots,10$ days and $j = 1,\ldots,N$, where N is the number of simulation runs performed. The actual simulated returns are given by:

$$r^*_{t+i} = \varepsilon^*_{t+i}\sqrt{h^*_{t+i}} \tag{3a}$$

where h^*_{t+i} is a (simulated) volatility estimate obtained as:

$$h^*_{t+i} = \omega + \alpha\left(\varepsilon^*_{t+i-1} + \gamma\right)^2 + \beta h^*_{t+i-1} \tag{3b}$$

and:

4. Normalised estimated distribution of returns in 10 days v. the normal density (10,000 simulations)

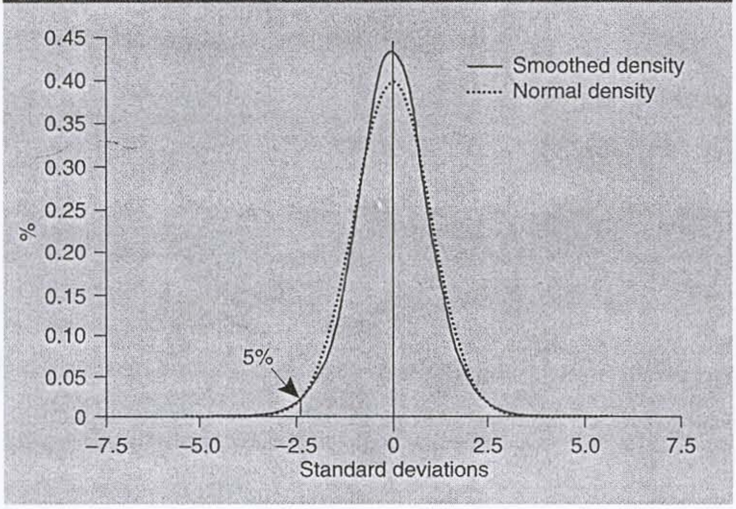

5. Estimated distribution of portfolio VAR in 10 days (10,000 simulations)

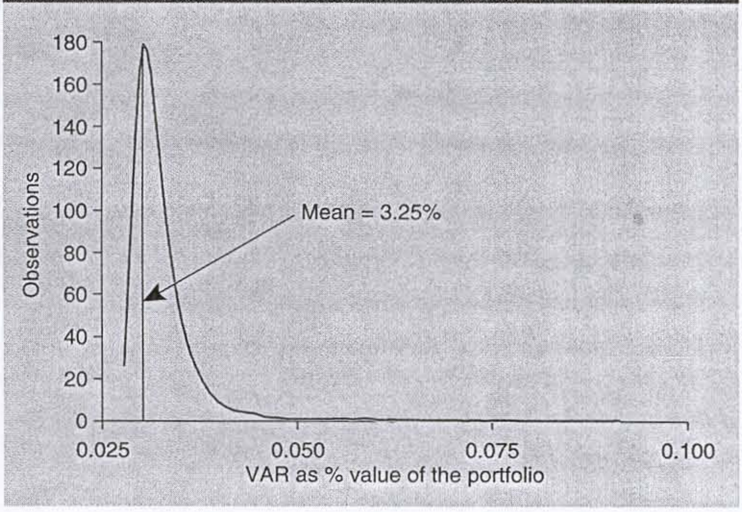

$$\varepsilon^*_{t+i} = z_t\sqrt{h^*_{t+i-1}}$$

z_t is a random standardised residual estimated as in equation (2), but rescaled to account for current market risk. In this way, we preserve the time-series properties of the data.

To obtain the distribution of our portfolio returns, we replicate the above procedure N = 10,000 times. The resulting (normalised) distribution is shown in figure 4. The normal distribution is shown as a dotted line for ease of comparison. We may extend our procedure to multiple assets, preserving the correlations of asset returns by taking returns in the same day for each asset as input to our simulation. Furthermore, unlike ordinary historical simulations, it is possible to preserve autocorrelation and lagged cross correlation patterns in the data by allowing past price changes to affect current returns.[3]

Not surprisingly, simulated returns on our well-diversified portfolio are almost normal, except for steeper peaking around zero and some clustering in the tails. The general shape of the distribution supports the validity of the usual measure of VAR for our portfolio. However, a closer examination of our simulation results shows how even our well-diversified portfolio may depart from normality. There are, in fact, several occurrences of very large negative returns, reaching a maximum loss of 9.52%. Our empirical distribution implies losses of 3.38% and 2.24% at confidence levels of 1% and 5% respectively.

The reason for this departure is the changes in portfolio volatility and thus in portfolio VAR, as shown in figure 5. The portfolio VAR over the next 10 days depends on the random returns selected in each simulation run. Its pattern is skewed to the right, showing how large returns tend to cluster in time. These clusters provide the base for realistic worst-case scenario analysis consistent with historical experience. To see the whole distribution of worst-case scenarios, we need simply to repeat our simulation and record the worst-case scenario of each run.

The worst-case scenario, as described in Boudoukh *et al*, is defined as the average of the outcomes in a given percentile. We have extended their approach by taking into account the effect of time-varying volatility. Of course, our method would produce more extreme departures from normality for less-diversified portfolios.

In conclusion, our simulation methodology allows for fast evaluation of VAR and worst-case scenarios for large portfolios. It takes into account current market conditions and does not rely on knowledge of either the correlation matrix of security returns or of the conditional distribution of the underlying process. Our methodology applies directly to asset returns that can be modelled as conditional heteroscedastic processes. Bonds and options may be included by expressing their values in terms of assets meeting our requirements, such as spot rates (for bonds) and underlying assets (for options). A full re-evaluation procedure for these assets can then be included at each step of our simulation (Barone-Adesi, Giannopoulos & Vosper, 1997). ■

[3] Only heteroscedasticity in this case. However, if appropriate, autoregressive and moving average returns can easily be inserted in equation (1a) to maintain any other properties

Giovanni Barone-Adesi is a professor at the University of Alberta, Canada and a visiting professor at City University Business School, London. Frederick Bourgoin is an analyst at Millennium Global Investments in London. Kostas Giannopoulos is a senior lecturer at the University of Westminster, London. The suitability of our model for stress testing is being evaluated at the London Clearing House (LCH).
We gratefully acknowledge LCH's continuing support for our research
e-mail: gbarone@gpu.srv.ualberta.ca
giannok@westminster.ac.uk

REFERENCES

Barone-Adesi G, K Giannopoulos and L Vosper, 1997
VAR without correlations for non-linear portfolios
Mimeo

Black F, 1976
Studies in stock price volatility changes
Proceedings of the 1976 Business Meeting of the Business and Economic Statistics Section, American Statistical Association, pages 177–181

Bollerslev T, 1986
Generalized autoregressive conditional heteroskedasticity
Journal of Econometrics 31, pages 307–327

Boudoukh J, M Richardson and R Whitelaw, 1995
Expect the worst
Risk September, pages 100–101

Christie A, 1982
The stochastic behaviour of common stock variance: value, leverage and interest rate effects
Journal of Financial Economics 10, pages 407–432

Efron B and R Tibshirani, 1993
An introduction to bootstrap
Monographs on Statistics and Applied Probability 57, Chapman & Hall

Engle R and V Ng, 1993
Measuring and testing the impact of news on volatility
Journal of Finance 48(5), pages 1,749–1,778

Jarque C and A Bera, 1980
Efficient test for normality, heteroscedasticity, and serial independence of regression residuals
Economics Letters 6, pages 255–259

Mandelbrot B, 1963
The variation of certain speculative prices
Journal of Business 36, pages 394–419