

# Non-parametric VaR techniques. Myths and Realities.

By Giovanni Barone-Adesi<sup>(1)</sup> and Kostas Giannopoulos<sup>(2)</sup>

November 2000

(1) Università della Svizzera Italiana and City Business School,  
email : [Giovanni.Barone-Adesi@lu.unisi.ch](mailto:Giovanni.Barone-Adesi@lu.unisi.ch)

(2) Westminster Business School, email: [giannok@wmin.ac.uk](mailto:giannok@wmin.ac.uk)

VaR (Value at Risk) estimates are currently based on two main techniques, the variance-covariance approach or simulation. Statistical and computational problems affect the reliability of these techniques. We illustrate a new technique, filtered historical simulation, that is designed to remedy some of the shortcomings of the simulation approach. We compare the estimates it produces with traditional bootstrapping estimates.

## **1 Introduction**

Early VaR (Value-at-Risk) techniques were linear multipliers of variance-covariance estimates of the risk factors. This class of market risk techniques soon became very popular, mainly because of their link to Modern Portfolio Theory. However, during worldwide market crises users noticed that early models failed to provide good VaR estimates. In addition, variance-covariance VaR techniques require a large number of data inputs; all possible pairwise covariances of the risk factors must be included in a portfolio. To process all the necessary information demands much computer power and time. Factorisation methods provide only partially satisfactory answers.

The early VaR models are also referred as **parametric** because of the strong theoretical assumptions they impose on the underlying properties of the data<sup>1</sup>. One such assumption is that the density function of risk factors influencing asset returns must conform to the multivariate normal distribution<sup>2</sup>. The empirical evidence however, indicates that speculative asset price changes, especially the daily ones, are rather non-normal. Excess kurtosis will cause losses greater than VaR to occur more frequently and be more

---

<sup>1</sup> Parametric VaR models are based on strong theoretical assumptions and rules. They impose that the distribution of the data (daily price changes) conforms to a known theoretical distribution. The most popular of these models is the exponential smoothing (ES), see RiskMetrics®.

<sup>2</sup> The normality assumption is frequently used because the normal distribution is well described; it can be defined using only the first two moments and it can be understood easily. Other distributions can be used, but at a higher computational cost.

extreme than those predicted by the Gaussian distribution. Many risk managers remember the large losses they faced during the Mexico (1996), Asian (1997) and Russian (1998) market crises. During these periods negative returns several standard deviations beyond the 2.33 threshold predicted by the normal distribution<sup>3</sup> were recorded within only few days. A large number of markets were crashing together; the correlation forecasts used to calculate VaR failed to predict such a synchronous crash. That resulted in further VaR failure.

The problems of earlier models spurred the search for better estimates of VaR. A number of recent VaR techniques are based on non-parametric or mixture of parametric and non-parametric statistical methods. The family of Historical Simulation (HS) models belongs to the former group. The Filtered Historical Simulation (FHS) as developed by Barone-Adesi *et al* (1998) and Barone-Adesi *et al* (1999, 2000) belongs to the second group. In this paper we analyse the assumptions upon which these models are based. In addition we compare the VaR estimates produced by the above models in linear and non-linear portfolios.

## **2 Literature review**

Regulators require that financial institutions backtest their internal VaR models, e.g. see Basle Committee on Banking Supervision (1995). Although the popularity and the use of HS has increased during the last few years, reports of backtests conducted by users are not publicly available. Some researchers however used smaller portfolios to backtest HS. Van den Goorbergh and Vlaar (1999) used rolling windows of different lengths (250, 500, 1000 and 3038 days) over a fifteen year period to backtest HS daily data on the AEX (Dutch equity index). They found that the failure rate, i.e. the probability that actual losses exceed VaR, is often exceeding the corresponding left tail probabilities. Van den

---

<sup>3</sup> At 99% probability.

Goorbergh and Vlaar found that results are sensitive to the selection of the window length<sup>4</sup>.

In another study Vlaar (2000) investigates the accuracy of various VaR models on Dutch interest rate based portfolios. He concluded that HS produced satisfactory results only when a long history is included in the data sample.

Brooks and Persaud (2000) investigated the sensitivity of VaR models due to changes in the sample size and weighting methods. They used a set of equally weighted portfolios each containing two asset classes selected from a set of national equity indices, bond futures and FX rates<sup>5</sup>. They found strong evidence that VaR models could produce very inaccurate estimates when the “right” historical data sample length is not selected.

Perhaps the most comprehensive model comparison study published to date has been carried out by Hendricks (1996). Hendricks used 4,255 daily observations of eight FX rates against the US dollar and several performance criteria to study the performance of twelve VaR models. The twelve VaR models were grouped in three categories, equally weighted moving average, exponentially weighted moving average and historical simulation approaches. He found that none of the twelve approaches is superior to others in every criterion. Furthermore, Hendricks reported that risk measures from the various VaR approaches for the same portfolio on the same date could differ substantially.

Differences in the accuracy across models were also sensitive to the choice of the level of probability used in the VaR calculation. When a 95% probability was used in the VaR calculation, Hendricks found that the three approaches produce accurate risk measures. However, when a 99% probability was used there was a large discrepancy in the risk calculation between the three approaches. In general the three approaches predict only between 98.2 percent and 98.5 percent of the outcomes. Hendricks failed to single out

---

<sup>4</sup> They strongly criticised the use of HS in predicting extreme events (far left on the tail) when the window is not of a substantial length.

<sup>5</sup> The indices used were S&P500 and FTSE100. The bond portfolios were the 30-year US Treasury bonds; UK Gilt and long term German Bund. The FX rates used were dollar/Swiss franc, dollar/Dmark and dollar/Yen. The US Treasury bond rates were obtained from near-month futures prices on the 30-year interest rate. No information was given for the other two bonds.

any VaR approach and he predicted that a more accurate a VaR model may be born by combining the best features of each single approach.

Pritsker (2000) reviews the assumptions and limitations of Historical and Weighted Historical simulation (Boudoukh *et al* 1998). He points out that both methods associate risk with only the lower tail of the distribution. In an example he showed that after the crash of 1987 the estimated VaR of a short equity portfolio, as computed by Historical Simulation or Weighted Historical Simulation, did not increase. The reason is that the portfolio recorded a huge profit during the day of the crash. Pritsker goes further by formulating some interesting properties of the Historical and Weighted Historical simulation. He showed that if the portfolio's return follow a GARCH(1,1) process, then at a 1-day VaR horizon and 99% confidence level, the historical simulation and weighted historical simulation method fail to detect increases in VaR about 31% of the time. In a simulated example he showed that the VaR on a short equity portfolio did not increase the days after the crash of 1987.

Barone-Adesi et al (1999) carry out an extensive backtest analysis for the FHS model. They used economic and statistical criteria to analyse the breaks on 100,000 daily portfolios held by financial institutions. The portfolios consisted of interest rate futures and options on futures trading at LIFFE during 1996 and 1997 as well purely plain vanilla swaps and mixed portfolios invested on futures, options and swaps. Overall their findings sustain the validity of FHS as a risk measurement model.

### **3 Historical Simulation**

Historical simulation (bootstrapping) is being increasingly used in the risk management industry. It consists of generating scenarios by sampling historical returns associated with each risk factor included in the portfolio. The aggregate value of all linear and derivative

positions produces a simulated portfolio value. The procedure is repeated many times using all past returns<sup>6</sup>.

Historical simulation does not require any statistical assumption beyond stationarity of the distribution of returns or, in particular, their volatility<sup>7</sup>. Since the estimated VaR is based on the empirical distribution of historical returns of each individual risk factor, it reflects a more realistic picture of a portfolio's past risk. At least one year of recent daily returns must be used in historical simulation. A longer period is more appropriate when available, but availability of historical data is often problematic for the universe of (linear) contracts or other risk factors. In its simplest form historical simulation can be shown by the following example. Given a data set of historical returns  $\Theta$  we draw an element  $e^*$

$$e^* = \{e_1^*, e_2^*, \dots, e_T^*\} \quad e_i^* \in \Theta \quad \text{where } i = 1, \dots, T \text{ refers to past days}$$

to form a simulated price for asset Y:

$$Y_{T+1}^* = Y_T + Y_T e^* \quad (1)$$

The process in (1) is repeated and the simulated price series  $Y^*$  is recursively updated up the last day of the VaR horizon. This sequence of simulated prices for day T+1, T+2, ..., T+N forms a simulated pathway or scenario for the risk factor Y. Here is a simple example:

Assuming that  $Y_T$  is 100 and the vector of  $e^*$ 's =

-0.01053	-0.00759	-0.00408	0.00474	0.00093	0.00921	0.01712	-0.00443	0.01342	-0.00304
----------	----------	----------	---------	---------	---------	---------	----------	---------	----------

---

<sup>6</sup> Past returns are drawn with or without replacement.

<sup>7</sup> As we are going to see later this becomes a limitation which under certain market conditions results in underestimation of VaR.

The sequence for simulated prices,  $Y_{T+1}^*, \dots, Y_{T+10}^*$  for the first simulation run will be

day	1	2	3	4	5	6	7	8	9	10
price	98.94707	98.19566	97.79456	98.25811	98.34967	99.25578	100.9552	100.5078	101.8569	101.547

Derivative securities are priced from the knowledge of the simulated price of the underlying asset by their full re-evaluation under a pricing model at each node. Derivative values can be aggregated to form a simulated value of all positions on each contract.

Historical simulation relies on a uniform distribution to sample innovations from the past<sup>8</sup>. These innovations are applied to current asset prices to simulate their future evolution. Once a sufficient number of different paths have been explored, it is possible to compute a portfolio's risk without making arbitrary distributional assumptions. That makes historical simulation is especially useful in the presence of abnormally large historical losses. The final outcome of historical simulation is the simulated distribution of portfolio values at the desired horizon. VaR is a percentile of that distribution.

Historical simulation has a number of advantages. It is easy to understand and to implement. It uses the empirical distribution of past returns to generate realistic future scenarios. Non-linear positions can be re-priced under each scenario<sup>9</sup>. Furthermore, historical simulation does not require the computation of a covariance matrix. In fact statistical dependencies across simultaneous asset returns can be kept by taking returns from the same day in the historical record for all the assets at each node in the simulation.

---

<sup>8</sup> Historical simulation is also known as bootstrapping simulation. For a detailed discussion about this simulation technique see Efron and Tibshirani (1993).

<sup>9</sup> This is a more efficient treatment of derivative security risk when compared to models that use linear approximations. But as we will see later it still presents the weakness of assuming that implicit volatility is constant.

Historical simulation's ability to predict future losses may be however undermined if the distribution of any risk factors is not i.i.d.<sup>10</sup>. Using a constant volatility model to calculate VaR when the distribution of returns is not stationary, as it is the case with most daily financial time series, could be very misleading<sup>11</sup>. In fact the probability of having a large loss is then not equal across different days. During days with higher volatility we would expect larger than usual losses. That contrasts with historical simulation, where the volatility of N days is proportional to the square root of time, i.e. constant volatility is assumed over any period.

Today there is a large body of evidence suggesting that the distribution of speculative price changes is fat-tailed with changing conditional moments. Empirical studies have found that there is a tendency, known as volatility clustering, for large price changes to be followed by more large changes. An increase in portfolio volatility during the risk measurement horizon affects the portfolio VaR<sup>12</sup>. Unfortunately HS does not take into account such a change.

---

<sup>10</sup> Independent and identically distributed. When returns are i.i.d. and the moments of the distribution are known, any inferences made about potential portfolio losses will be accurate and unchanging over time. Stationarity implies that the probability of occurrence of a specified loss is the same for each day. Independence implies that the size of price movement in one period will not influence the movement of any successive prices. These properties combined with the normality assumption simplify the VaR calculation for any holding period and probability. Under these hypotheses VaR for longer periods can be found by multiplying the VaR estimate at a shorter, *i.e.* daily, horizon by the square root of the number of days  $t$  in the period of interest.

<sup>11</sup> Historical simulation assumes that past and present moments of the density function of returns of a specific risk factor are constant and equal. For a detailed discussion of the properties of historical simulation see Pritsker 2000.

<sup>12</sup> For some types of investments two or three consecutive adverse price changes may be sufficient to ruin the investor, *i.e.* contingent claims and investments with leverage. In a leveraged portfolio investors borrow at a fixed rate and invest in a risky asset, usually equity. Leverage increases positive expected return because leverage magnifies volatility. The larger is the leverage factor, the higher the gains will be in case of positive returns. But large adverse returns increase the downside risk and chances of a disaster. Thus the VaR is far more important in the management of leveraged investments.

## 4 Filtered Historical Simulation

The above limitations of HS, a limited set of outcomes and unresponsiveness to changes in market volatility, can be overcome with the use of Filtered Historical Simulation (FHS), by Barone-Adesi et al (1999). Filtered historical simulation is a generalised historical simulation. It has all the positive properties and overcomes most of the historical simulation weaknesses. In filtered historical simulation the stationarity assumption is relaxed; historical returns<sup>13</sup> are first standardised by volatility estimated on that particular day (hence the name of filtered),  $\mathbf{h}_t = \frac{\mathbf{e}_t}{\sqrt{\hat{h}_t}}$ . This filtering process yields

approximately i.i.d. returns (residuals) suited for historical simulation.

Before filtered returns are used as innovations they are scaled (multiplied) by the current conditional forecast of volatility; thus they reflect current market conditions:

$$Y_{T+n}^* = Y_{T+n-1}^* + Y_{T+n-1}^* \mathbf{h} \sqrt{h_{T+n}} \quad (2)$$

where  $\mathbf{h}^* = \{\mathbf{h}_1^*, \mathbf{h}_2^*, \dots, \mathbf{h}_T^*\}$   $\mathbf{h}_i^* \in \Theta$  where  $i = 1, \dots, T$  refers to

past days and  $\sqrt{h_{T+n}}$  is the simulated conditional volatility for VaR day  $n$  and is estimated recursively by a time-series model, such as the one illustrated in equation (3),

$$h_{T+n} = \mathbf{w} + \mathbf{a}(\mathbf{z}_{T+n-1}^*)^2 + \mathbf{b} h_{T+n-1} \quad (3)$$

where  $\mathbf{z}_{T+n}^* = \mathbf{h}^* \sqrt{h_{T+n}}$  is computed as in (2).

---

<sup>13</sup> Since returns are usually serial correlated in practice we replace  $Y_t$  with  $\mathbf{e}_t$ , where  $\mathbf{e}_t = Y_t - E[Y_t | \Phi_{t-1}]$

The “unconditional” or “unfiltered” historical simulation described earlier is a special case of FHS, which holds when returns are i.i.d<sup>14</sup>. A major advantage of FHS over HS is that the filtering process increases the range of outcomes beyond the historical record through a change of scale. In other words, FHS provides a systematic approach to generate extreme events not present in the historical record, completing the tails of the distribution. FHS requires therefore a shorter historical record than HS to simulate the tails of the distribution of returns. Because this process is essentially an extrapolation, its validity must be carefully tested. Barone-Adesi *et. al.* (2000) present evidence of the adequacy of this procedure for risk management.

The FHS approach can be adapted to stress testing because it simulates the whole distribution of security returns. It is not limited to the observed returns as regular bootstrapping is. Therefore it is possible to sample from more extreme points in the tails of the multivariate distribution by increasing the number of simulation runs. As an example, applying FHS over a 10-day horizon using a database of 500 daily returns, the number of different possible pathways is  $500^{10}$  for any given set of initial conditions. Therefore it is possible to generate an almost arbitrarily large number of points of the simulated distribution. This is especially important for portfolios of derivatives that may experience more stress at points not on the tails of the distribution.

Finding the most stressful conditions for a given portfolio may be slow, because the probability of each pathway is the inverse of the number of pathways. Of course the problem of finding the most stressful conditions is easier for linear portfolios, for which the most stressful conditions are given by the largest negative or positive returns, that can be pre-selected in the simulation procedure.

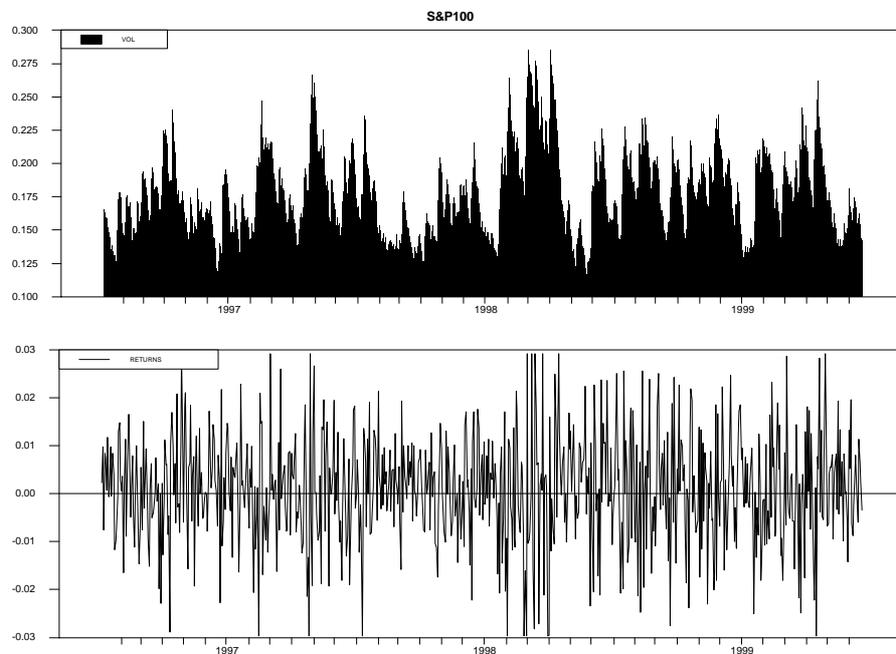
---

<sup>14</sup> The algorithm can easily be expanded to generate parallel pathways on a multi-asset environment. This does not require the use of a correlation matrix. For a full description see Barone-Adesi *et. al.* (1999).

## 5 HS or FHS ? An empirical investigation

In this example we investigate the difference in VaR estimates between HS and FHS. Three hypothetical portfolios are selected. The first portfolio consists of a long position on the S&P100 while each of the other two consists of a short European Call option on the same underlying. The historical data set used in the simulation covers the period from 1 January 1997 until 26 November 1999. Our scope is to compare VaR estimates for the three portfolios using the two alternative methodologies. In figure 1 are shown the daily returns and conditional volatilities from 2 January 1997 until 18 December 1999 for the S&P100. The conditional volatility of the returns, shown on the top graph, is rather non-constant, a fact that contradicts the assumption upon which historical simulation is based.

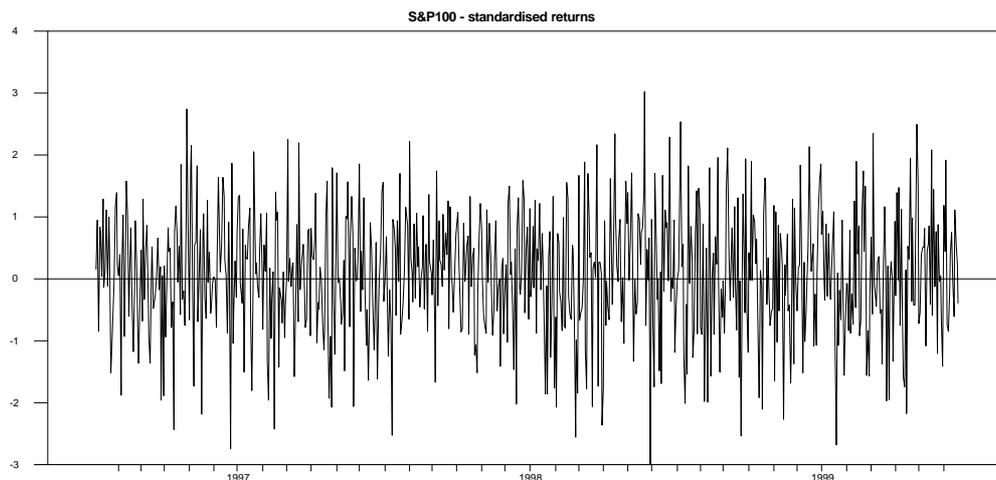
**Figure 1 conditional Volatility and returns for the S&P100**



The value of the linear portfolio, the S&P100 index, at the close of business on the 26 November is \$753.56. The first step in our simulation consists in generating the simulated scenarios for the S&P100 index using the two alternative methods. To form the innovations,  $e$ , in (1), we select a random return<sup>15</sup> from the historical data set<sup>16</sup>. By updating (1) we generate multiperiod scenarios. A total of 5000 simulated multiperiod scenarios, for 1, 2,...20 business days ahead, are generated.

By dividing each of the daily returns with that day's volatility we form the set of innovations,  $\eta^*$ , utilised by the FHS simulation. These are shown in figure 2.

**Figure 2** Standardised Returns



Using the above data set of standardised residuals and the FHS algorithm we generate 5000 multiperiod scenarios for days 1,2,.. 20.

---

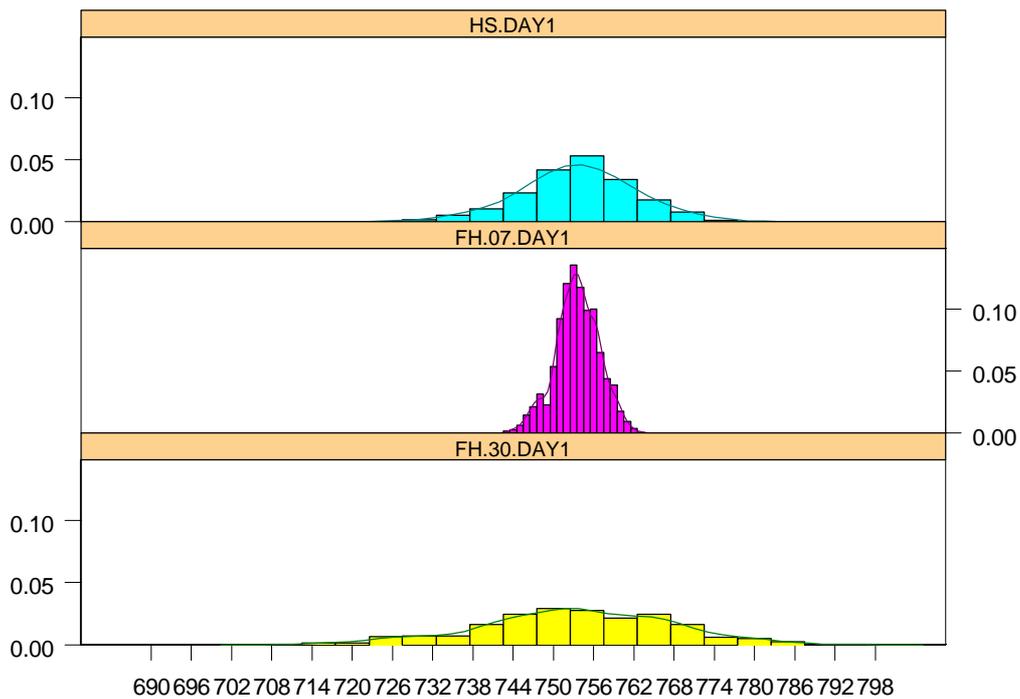
<sup>15</sup> These are rather zero mean returns.

<sup>16</sup> The data are drawn with replacement. There are 770 observations in the data set. At each simulation run on average each observation will be drawn 7 time circa.

To investigate any possible discrepancy on VaR between the two, HS and FHS, methods, we carry out the FHS using two alternative risk scenarios on the last trading day. Under the first set of simulations we set the volatility on 26 November 1999 to a low, but not unrealistic, level of 7% p.a. On the second sets of simulations we assume that there was a market crisis and the market volatility on the last trading date jumped to 30% p.a.

Figure 3 shows the histogram for the simulated prices for day one as produced in each of the three simulations. The top histogram displays the distribution of returns from the HS, while the other two histograms report returns from the FH under each the two, low and high, volatility regimes prevailing on the last trading day.

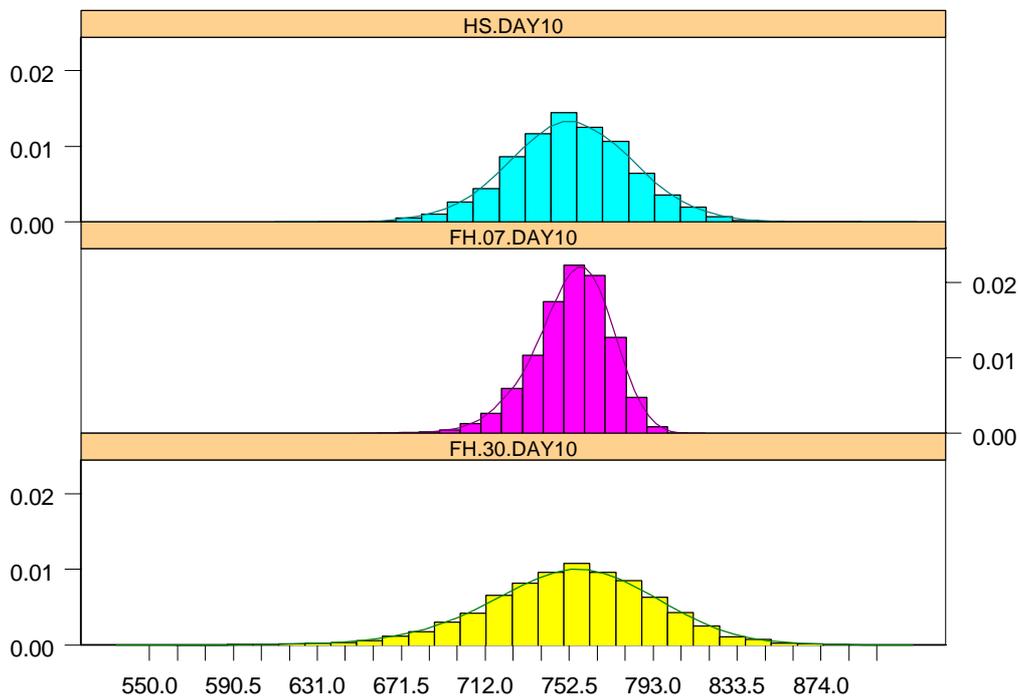
**Figure 3 S&P100 simulated prices for day 1**



The histograms for each set of simulated prices corresponding to a ten day horizon are reported on figure 4. As one should expect, in volatile period, e.g. when the volatility on

the last trading date is 30% p.a., the forecast interval is wide and in quiet periods is narrow. This is reflected in the forecast produced by the FHS. On the contrary, HS makes interval forecasts that are static, taking no notice of the last trading date's risk level.

**Figure 4 S&P100 simulated prices for day 10**



The VaR at 99% is the lower percentile in the set of simulated prices. Table 1 reports the various VaR estimates for days 1, 5, 10, 20. The second column from the left shows the VaR at 99% as estimated by the HS. The other two columns report the FHS under two different risk states prevailing at the close of business on 26 November 1999. The historical volatility in the historical data sample was 19.5% p.a.<sup>17</sup>.

<sup>17</sup> By calibration the average GARCH volatility is equal to the historical one.

**Table 1 VaR estimates (at 99%) on a linear portfolio**

Horizon (days)	HS	FHS @7%	FHS @30%
1	23.87	8.30	35.67
5	57.96	30.08	72.91
10	73.66	50.47	102.04
20	97.52	87.63	129.23

As expected the VaR estimates produced under the HS occur between the other two estimates as computed by the FHS algorithm, e.g. using two different volatility values on the close of business on 26 November 1999. The one-day VaR is proportional to the initial volatility used in this simulation process. Due however, to the mean-reverting volatility process that characterises the parameterisation of conditional volatility on all GARCH models, we observe that the VaR's estimated with the FHS algorithm have a tendency in the long run to converge towards the HS VaR estimates. The ratio of one-day VaR estimates computed by FHS with initial volatility of 7% p.a. and that computed by the HS is 0.35. This is very close to the ratio between the initial volatilities used by the two methods. On a ten day VaR the same FHS estimates closes the gap, with the ratio rising to 0.68, and on a twenty day VaR the ratio increases to 0.90. We draw similar conclusions by analysing the ratio between the HS and that of FHS that has an initial volatility of 30% p.a. This observation indicates that the bias present in the HS estimates becomes smaller at longer horizons.

The discrepancies in VaR between the two methods, HS and FHS are obvious in the above example. The magnitude in the discrepancies, however, increases when the portfolio contains predominantly non-linear instruments. In the next example we will use two alternative portfolios, each containing a single, short, European Call option. The option in the first portfolio has a strike of \$678.00 and is in the money. The second Call is chosen to be out-of-the-money with a strike of \$828.00. Both options expire in twenty business days. We used the Black and Scholes model (1973) to find their market value<sup>18</sup> on 26 November 1999. The market value for the "in-the-money" Call is estimated at

---

<sup>18</sup> Assuming that the implied volatility equals the historical volatility on 26 November 1999, i.e. 19.5% p.a.

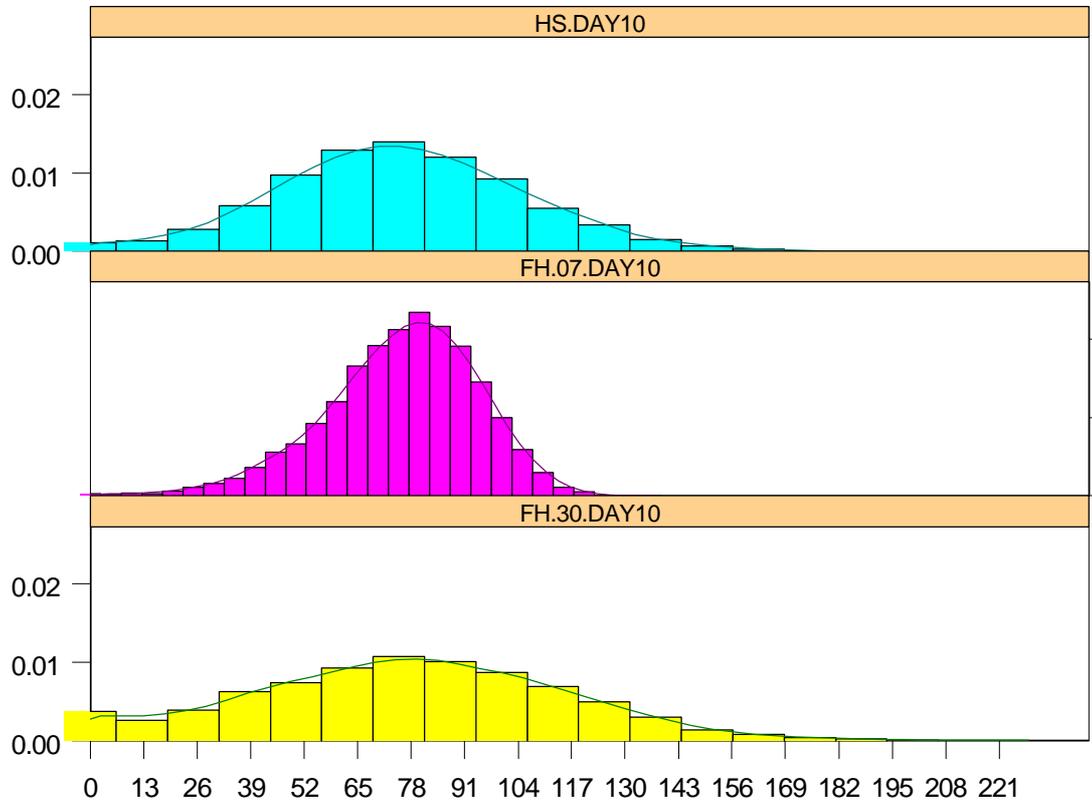
\$77.54 and the value for the “out-of-the-money” Call is estimated at \$0.85. These numbers are the values of the two portfolios at the end of the business on 26 November 1999.

Using the same simulated scenarios for the linear portfolios we computed sets of 5000 payoffs and profits for the non-linear portfolio. The profit at VaR horizon  $t$  for a single simulation run is given below

$$\text{Profit}_t = -(\text{Max}[P_t - X, 0], -C) \quad (4)$$

where  $X$  is the strike,  $C$  is the market value of the Call on 26 November 1999 and  $P_t$  is the simulated value of the underlying asset at VaR horizon  $t$  for that simulation run. The simulated profits for the in-the-money Call and for a 10 Day horizon are shown in figure 5. The top histogram displays the HS profits and the other two the profits for the FHS under low and high prevailing volatility regimes prevailing on the last trading date.

**Figure 5 Simulated payoff for the “in-the-money” European Call at 10-Day VaR**



The upper percentile on each of the simulated set of prices less the call premium is our VaR estimate at 99% probability. These values together with the ones for day 1, 5 and 20 are reported in Table 2. Note that the effect of low volatility is relatively stronger at short horizons, but it is still important at longer horizons. In summary, neglecting to adjust historical simulation to current market conditions through FHS may cause substantial errors in the computation of risk.

**Table 2 VaR estimates (at 99%) on a short, “in the money” European Call**

DAY	HS	FHS-07	FHS-30
1	19.2815	5.60252	26.43452
5	48.60474	19.43428	59.78121
10	70.23241	33.21167	80.42265
20	97.42933	59.60564	106.8141

On table 3 we report the equivalent VaR values for the portfolio invested in the short, out-of-the-money Call option, with a strike of \$828.00.

**Table 3 VaR estimates (at 99%) on a short, out-of-the-money, European Call**

DAY	HS	FHS-07	FHS-30
1	0.0	0.0	0.0
5	0.0	0.0	0.0
10	0.0	0.0	5.28878
20	22.55335	0.0	29.50686

The effect of different volatility on VaR estimates is even more important at the longer horizon.

## 6 Conclusions

In principle filtered historical simulation (FHS) can produce risk measures that are consistent with the current state of markets at any arbitrarily large confidence level. These impressive gains allow FHS to dominate historical simulation (HS) easily. In fact historical simulation fails to condition forecasts on the current state of the markets. The difference between the two simulation methods is magnified by the presence of options in the portfolio.

Although the superiority of FHS on HS is obvious, further development may mitigate some of the current limitations of FHS. These limitations of FHS rest mostly on two of the assumptions on which it is based: that correlations across asset returns are not related to the scale of returns and that the scaling process is accurately describing the tails of the distribution of asset returns. Neither of these assumptions is literally true. Therefore our current FHS technique should be considered a first-order approximation to a simulation engine accurately describing portfolio returns.

## Bibliography

- Barone-Adesi G, F Bourgoïn and K Giannopoulos, 1998, “Don’t Look Back”\*, *Risk*, **11**, August, pp100-104.
- Barone-Adesi G, K Giannopoulos and L Vosper, 1999, “VaR Without Correlations for Non-Linear Portfolios”\*, *Journal of Futures Markets*, **19**, August, 583-602.
- Barone-Adesi G, K Giannopoulos and L Vosper, 2000, “Backtesting the Filtered Historical Simulation”\*, unpublished manuscript.
- Basle Committee on Banking Supervision, 1995, “An Internal Models-Based Approach to Market Risk Capital Requirements”, Basle (available at <http://www.bis.org>).
- Boudoukh J, M Richardson and R Whitelaw, 1998, “The Best of Both Worlds”, *Risk*, **11**, May, pp64-67.
- Brooks C and G Persaud, 2000, “The Pitfalls of VaR Estimates”, *Risk*, May, pp63-66.
- Efron and Tibshirani, 1993, *An Introduction to Bootstrapping*, Chapman and Hall.
- Hendricks D, 1996, “Evaluation of Value-at-Risk Models Using Historical Data”, Federal Reserve Bank of New York Economic Policy Review, April, pp39-69.
- Pritsker M, 2000, “The Hidden Risks of Historical Simulation”, Federal Reserve Board, Washington.
- Van den Goorbergh R and P Vlaar, 1999, “Value-at-Risk Analysis of Stock Returns. Historical Simulation, Variance Techniques or Tail Index Estimation?” Mimeo.
- Vlaar P, 2000, “Value at Risk Models for Dutch Bond Portfolios”, *Journal of Banking and Finance*, **24**, pp1131-1154.

(\*) draft copies of these papers could be found on our web site: [www.fh-simulation.com](http://www.fh-simulation.com)