MASSIMILIANO PALLOTTA, RAFFAELE ZENTI

RISK ANALYSIS FOR ASSET MANAGERS: HISTORICAL SIMULATION, THE BOOTSTRAP APPROACH AND VALUE AT RISK CALCULATION

NOVEMBER 2000

massimiliano.pallotta@rasnet.it
raffaele.zenti@rasnet.it

EXECUTIVE SUMMARY

From the risk management's perspective, one of the main differences between asset management companies and banks concerns the investment horizon: typically, asset managers have longer investment horizons. We compare different ways to deal with medium/long horizons, when the aim is to calculate absolute or relative VaR using an historical simulation approach and its variations, like bootstrapping procedures. We use several indices to test the accuracy of the different methods analyzed. We find these methodologies:

- can provide satisfactory assessments of tactical risk;
- can inform portfolio managers of changes in market risk;
- are also promising for strategic risk analysis.

RISK ANALYSIS FOR ASSET MANAGERS: HISTORICAL SIMULATION, THE BOOTSTRAP APPROACH AND VALUE AT RISK CALCULATION

1. INTRODUCTION

On the one hand in the community of asset management companies the quest for reliable risk management techniques has grown in recent years mainly in response to the demand of sophisticated investors such as pension funds and foundations. Falloon (1999) reports on the growing interest of pension funds for risk management techniques and tools. On the other hand, in this decade Value at Risk (VaR) has became a very popular measure of market risk. VaR is the loss on the portfolio that will not be exceeded with a specified probability over a specified time horizon. VaR is an extremely powerful risk measure, because looks at downside risk, i.e. it is well suited for asymmetrical distributions, and because in principle it can be calculated assuming any kind of distributions of portfolio returns. VaR is widely used for controlling traders, for determining capital requirements and for disclosure to external subjects, both investors and regulators. Unfortunately, the majority of the models regarded as standard practice were born for banking use (sellside portfolios). Among the assumptions commonly made in the banking/trading environment we find i) short time horizons (1-10 days); ii) normally distributed returns; iii) zero expected return; iv) constant variance-covariance matrix; v) absence of a benchmark. In contrast, for an investment (buyside) portfolio is reasonable to assume that: i) horizon is relatively long, at least 1 month, often several months or even years; ii) over such horizons it is not safe to assume that portfolio returns are normal; iii) the expected change in the portfolio value is not zero; iv) volatilities and correlations vary over time and are not necessarily linear; v) usually, there is a benchmark.

Since VaR is an interesting and common risk measure, it does therefore make sense to adapt it to the asset management environment. Adapting VaR measures for asset managers (rather than traders) involves finding a proper way to model future scenarios, preserving the multivariate properties of asset returns, when time horizon is relatively long.

The simplest solution, parametric models based on the hypothesis that asset returns are normally distributed, can be unreliable because empirical distributions may not conform to any known distribution. In addition, this kind of model requires estimating the correlations among market risk factors. The elements of a correlation matrix increase with the square of the number of risk factors in the portfolio: for large portfolios the correlation matrix may not be positive definite. In order to avoid this problem and deal with smaller correlation matrices it is common practice to map the effective portfolio in a simpler portfolio that preserves some properties of the original one. This mapping process can oversimplify the portfolio structure. For instance, this is the case when a stock is mapped onto a single country stock index using the CAPM.

The RiskMetrics Group's proprietary LongRun™ methodology (1999) provides an integrated methodology, for generating market rate scenarios over long horizons using two forecasting methodologies: one based on current market information, the other based on econometric models. The forecasts based on current market prices make intensive use of spot, futures, forwards and options price data and apply some derivatives theory to extract information from price data, while the forecasts based on economic fundamentals rely on historical time series of financial and economic data and the econometric modeling of time series. Problems arise if the number of portfolio risk factors is high. In this case, if one chooses to use the first method it is hard to find derivatives for each risk factor. Additionally, the use of multivariate econometric models is often restricted to a few series at a time. Therefore, in order to generate scenarios for a wide number of assets, one needs different methodologies.

Historical simulation provides a flexible and intuitive framework for risk analysis, but its basic version uses only one path of returns ("the realized world") and therefore produces risk indicators with high variance. When the goal is to model returns on a horizon longer than data frequency, Monte Carlo simulation or bootstrapping techniques can be seen as sensible choices. Usually, the approach based on Monte Carlo simulation uses a set of stochastic differential equations for generating returns over the time horizon. Note that Monte-Carlo simulation uses arbitrary distributional assumptions, imposing the structure of risk that it is supposed to investigate. Instead, the bootstrapping approach can be seen as a variation of the historical simulation approach, where one resamples from the empirical distribution of portfolio returns. So it does mix Monte Carlo and historical simulation. This method guarantees that the multivariate properties of original data are preserved and is flexible enough to incorporate an update of both mean and volatility.

The aim of this paper is to discuss and compare different ways to deal with a relatively long horizon, when the aim is to calculate VaR using the historical simulation/bootstrapping methodology and its variations.

The plan of this paper is as follows. Historical simulation and VaR, both absolute and relative (versus benchmark), are introduced in the following section. In section 2 we point out some key issues concerning VaR computation when time horizon is relatively long, such as in the asset management field. In section 3 and section 4 we discuss, respectively, a simulation framework to deal with long time horizons, and some variations of the basic historical scheme that enable to extend the capabilities of the approach beyond the original limits. In section 5 we show how to introduce options in the bootstrapping scheme. In section 6 we compare several approaches an asset management company can use to estimate VaR. Conclusions are then drawn in the final section.

2. HISTORICAL SIMULATION AND VAR

The formal definition of VaR, expressed in percentage terms, is given implicitly by the following equation:

$$\Pr\left[r_{PTF} \leq VaR\right] = \alpha \tag{1}$$

where r_{PTF} is the portfolio return over a given time horizon H, with probability α. Equation (1) states that VaR is the portfolio return during the time period H that is exceeded with probability α. If calculated in absolute terms, VaR measures portfolio total risk, while in relative terms (i.e. using excess returns versus benchmark, instead of total returns) it measures active risk and it is named ReVaR. For a classic reference on VaR see Jorion (1996).

From the above discussion it is obvious that VaR is the α-th quantile of r_{PTF}, denoted by $Q_\alpha(r_{PTF})$. In order to calculate quantiles one needs the probability density (pdf) function of future returns, which characterizes the future scenario. One of the major problems in implementing VaR analysis is the specification of the pdf. For computational convenience asset returns are often assumed to be conditionally normal. Normality greatly simplifies VaR calculation because implies that all percentiles are multipliers of the standard deviation. Unfortunately, return distributions quite commonly exhibit fatter tails than the normal distribution. There is thus a need for alternative methodologies that adequately describe the tails, especially the left tail. Historical simulation is a nonparametric way to estimate VaR that allows for fat tails. The idea behind this approach is quite simple: using the time series of all portfolio risk factors (e.g. all the assets currently held in the portfolio) and their current weights, one creates the simulated time series of portfolio returns. These returns do not represent an actual portfolio but rather reconstruct the history of a hypothetical portfolio using the current positions. In order to implement an historical simulation one needs: (i) for each risk factor, a time-series of actual returns, and (ii) risk factors' weights.

Let us assume there are K risk factors, for the sake of simplicity corresponding to K linear instruments. Their percentage weights are the elements of vector ω:

$$\omega = \begin{bmatrix} \omega_1 \\ \omega_2 \\ \cdot \\ \omega_k \end{bmatrix} \tag{2}$$

and their time series, in terms of log-returns, are in the matrix r, which has K columns and N rows, corresponding to the number of observations in the database. For the moment we assume these returns have a frequency corresponding to the time horizon of VaR analysis. This means that with a time horizon of 1 week, data frequency is weekly. Defining the oldest returns as $r_{i,-N}$, i = 1,2,3,...K, the matrix containing the risk factors time series is:

$$r = \begin{bmatrix} r_{1,-N} & r_{2,-N} & \cdot & \cdot & r_{K,-N} \\ r_{1,-N+1} & r_{2,-N+1} & \cdot & \cdot & r_{K,-N+1} \\ r_{1,-N+2} & r_{2,-N+2} & \cdot & \cdot & r_{K,-N+2} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}. \tag{3}$$

The "constant mix" time series of portfolio returns $\{r_{PTF}\}$, for brevity r_{PTF}, is given by the following product:

$$r_{PTF} = r \cdot \omega \tag{4}$$

and VaR is $Q_\alpha(r_{PTF})$. Of course, if asset returns are normally distributed, VAR estimates obtained with the historical simulation method and the parametric-normal method are the same.

Basically, instead of making assumptions about the exact shape of the pdf of returns, the only assumption is that future returns follow the same distribution of actual past returns. Therefore an historical simulation approach accurately reflects empirical skewness, kurtosis and any non-linear time-varying variance-covariance matrix. Note indeed that each element of vector r_{PTF} is a linear combination of the actual realized returns of the K risk factors currently held in the portfolio. This allows us to model the pdf of portfolio returns in a way that accurately reflects the historical multivariate pdf of risk factors. Barone-Adesi, Giannopoulos and Vosper (1999) and Hull and White (1998) show that historical simulation can incorporate volatility updating. We will consider the volatility updating method in section 4.

We can easily extend the historical simulation approach switching from the distribution of total returns to the distribution of relative returns, or excess returns. The "constant mix" time series of portfolio excess returns er_{PTF} is given by:

$$er_{PTF} = r \cdot \Delta\omega = r \cdot (\omega_{PTF} - \omega_{BMK}) \quad (5)$$

where ω_{PTF} and ω_{BMK} are column vectors with K rows containing respectively the portfolio weights and the benchmark weights of the K risk factors. The Relative Value at Risk (ReVaR) is defined as the α -th quantile of er_{PTF} .

For pension funds, mutual funds and other buyside portfolios absolute risk indicators like VaR can be used to measure risk before deciding on a strategic asset allocation policy. VaR can also be used on an ordinary basis when portfolio management is on a total return basis. ReVaR, instead, is a risk management tool more appropriate whenever a benchmark exists, and can give insight on portfolio risk, on an ordinary or tactical basis. However it is evident that historical simulation allows us to calculate any desired risk indicator, e.g. shortfall probability, worst case, mean shortfall return/excess return and so on.

One can debate if it is more appropriate to use tracking error, a more common measure of relative risk, or ReVaR. Tracking error is the standard deviation of excess returns. It is therefore appropriate only when one believes returns are normally distributed: in this case mean and standard deviation are sufficient to fully characterize excess returns. Otherwise the use of tracking error can be misleading. For example, if the portfolio excess returns pdf is negatively skewed and shows leptokurtosis, the use of tracking error leads to an underestimation of relative risk. The use of ReVaR calculated using an historical simulation approach helps to cope with this problem. It is a quantile-based risk measure, so in principle it takes into account not only the first two moments, but also higher moments; it is a downside risk indicator, focusing on negative events. In addition, it is calculated using an historical simulation approach, therefore it uses a more realistic description of the nature of the tail of the distribution. For an asset manager ReVaR is a more robust way to quantify relative risk. In addition, computing VaR using an historical simulation approach can be appealing from the practitioner point of view because, among the various pros, the method is rather intuitive and relatively simple to implement.

3. REVAR AND VAR MEASURES FOR ASSET MANAGERS: A MULTI-STEP HISTORICAL SIMULATION FRAMEWORK BASED ON BOOTSTRAPPING

As we mentioned before, for an investment portfolio it is reasonable to assume that horizon is relatively long, probably in the range 1 month-1 year or even more. For example, foundations and pension funds need to have a strategic view of risk that must consider multiple years. Even if we consider only an ordinary view of risk, i.e. a view linked to tactical/operative portfolio management, time horizon is relatively long, if compared to the typical time horizon of sellside portfolios. This is due to the fact that buyside portfolios are often rather static. Active asset managers adding value through security picking or strategic/tactical asset allocation decisions tend to have a low turnover, if compared to sellside portfolios. The same is true for passive or semi-passive portfolios: rebalancing frequency and turnover have to be low, in order to reduce management costs. If a portfolio has a high turnover, time horizon ideally corresponds to the period of time it would take to unwind positions under adverse market conditions. On this point, the Basle Committee on Banking Supervision requires banks to set a 10-day time holding period. It is therefore reasonable to think that asset managers need to use a longer holding period, because of the likely difficulties of unwinding large positions without moving the market. Probably it is reasonable to think the more static is portfolio management style, the longer is time horizon. We add that the chosen time horizon could be associated to periodic portfolios formal assessments, like investment committees and asset allocation meetings, which commonly take place every month.

Over medium/long horizons it is not safe to assume that portfolio returns are normal with mean equal to zero. Figure 1 to Figure 6 and Table 1 and 2 show some evidence that monthly log-returns over the period March 1971 – May 2000 exhibit time varying volatility, skewness and fat tails (they are not normal) and have a mean which is statistically different from zero. This is true both for a broad stock index like MSCI World (Figure 1 – Figure 3, Table 1) and a volatile stock index like Nasdaq Composite (Figure 4 – Figure 6, Table 2). Both indexes are expressed in

Euro terms, but results are basically the same if we consider local returns.

[insert Figures 1-3]

[insert Table 1]

[insert Figures 4-6]

[insert Table 2]

Similar results are rather common, when examining monthly returns of equity and fixed income indices. See for example Zenti (2000).

The above discussion suggests that adapting VaR measures for asset managers (rather than traders) involves finding a proper way to model future scenarios when time horizon is relatively long and returns are not necessarily zero-mean normal. If returns were i.i.d. normal one could simply use the well-known "square root rule". In the historical simulation framework, where returns do not follow any predetermined distribution and are not necessarily i.i.d., there are basically two ways to generate future scenarios, depending on the timescale of interest and the timescale for one's model or data.

The first way requires a database of returns with a frequency corresponding to the time horizon. If one is, say, interested in a horizon of one month and has one month returns, then she or he can use directly this data to generate a distribution of future scenarios and thus to estimate risk measures. The first problem with this approach is that it can be difficult to get the time series of the length needed for a historical simulation for all the risk factors. A bigger problem is that, even if the length of the time series is reputed satisfactory, one has the problem that historical data may correspond to completely different economic circumstances than those that currently apply.

The second way uses a database of returns with a frequency higher than the time horizon. Suppose data is for a short timescale, let us say we have daily data, and the horizon is one month. Then a technique must be used to build up a one-month distribution by generating whole month-long paths of asset returns starting from daily data. This is by far more time consuming but important because it uses the information content of daily data about the relative or absolute risk that a portfolio might exhibit. It is also essential if the portfolio contains path-dependent contracts when the whole path taken has to be modeled. To do so one can apply bootstrapping techniques using actual asset returns movements taken from historical data. For an introduction to the bootstrap see Efron and Tibshirani (1998). Suppose we have three years of daily log-returns for K risk factors and our time horizon is one month, approximately 21 days. Applying (4) for total risk analysis or (5) for relative risk analysis we get the "constant mix" time series of portfolio returns. Let us say we use (5) and we get er_{PTF}. It is a vector with, say, 750 rows. Then we assign a position index to each excess return; that is, we assign 750 natural numbers, one for each row. Now we draw a uniformly distributed number from 1 to 750; assume it is 118. We take the 118th element of er_{PTF}. Then we draw another uniformly distributed number from 1 to 750; it is 515. We take the 515th element of er_{PTF}. We continue this process until we have collected 21 daily log-excess returns; their sum is one realization of the path of the portfolio excess returns over the required time horizon. Repeating this simulation 5000 times (for instance) allows us to generate an accurate distribution of all future excess returns.

Note that we have described the procedure in the univariate case, i.e. we apply the bootstrapping procedure to portfolio returns or excess returns. One can easily extend this approach to the multivariate case, i.e. bootstrapping several asset returns at the same time. Let reconsider the example, assuming we have, say, ten assets. If the first uniformly distributed number, drawn from 1 to 750, is 118, we take the 118th elements of all the ten series. If the second uniformly distributed number from 1 to 750 is 515, then we take the 515th elements of all the ten series, and so on. It is clear that we keep together all the observations occurred on the same day. This multivariate bootstrapping techniques is commonly known as parallel bootstrapping. If the aim is to calculate portfolio risk measures and one uses the same uniform random variates, it is easy to verify that parallel bootstrapping leads to the same results one can get bootstrapping directly the portfolio returns as defined by (4) or (5). Therefore for the moment we will concentrate on the univariate method; we will need parallel bootstrapping later, in section 5, when dealing with options.

Applying 5000 times the univariate procedure to the daily log-returns of the equity index MSCI World (in Euro terms), from 8/9/1971 to 8/5/2000, we can estimate the return distribution with one-month horizon. In Table 3 we compare the first four moments of the estimated distribution to the first four moments of the actual distribution of monthly return. Although the procedure is simple, it seems to work rather well, because the estimated moments are not significantly different from the actual moments.

[insert Tables 3]

Formally, if the goal is to simulate M scenarios in terms of excess returns[1] over a time horizon of H periods (e.g. 21 business days ~ one month) we need to generate, for each scenario, a vector of H (pseudo) integer random numbers, uniformly distributed in the range between 1 and N:

$$u_i \sim U[\text{int}(1, N)], \quad i = 1, 2, 3, ..., H. \tag{6}$$

where N is the length of the time series. Thus, for each scenario, the excess return ER over H periods is given by:

$$ER = \sum_{i=1}^{H} er(u_i) \tag{7}$$

where $er(u_i)$ is the element of vector er_{PTF} that is in position u_i. Note that vector er_{PTF} contains "raw" data (e.g. daily excess returns). Repeating this simulation M times allow us to generate a distribution of future excess returns keeping together all the returns that happen on a

[1] The procedure is the same for absolute returns.

certain date. By doing this we ensure that we capture any correlation that there may be among assets and any non-normality in asset price changes, as Jorion (1998, pp. 237-239) points out. This method of bootstrapping is very simple to implement. A naïve application of this method does not capture any autocorrelation in the data, but then neither does a Monte Carlo simulation in its basic form. This procedure is based on the assumption that "raw" excess returns are identically, independently distributed. If "raw" returns are not i.i.d., they are unsuitable for bootstrapping and can lead to biased results, because, for example, we ignore the eventual presence of autocorrelation and volatility clusters. To avoid this problem, it is possible to modify the basic scheme. In the following section we will examine some variations of the basic scheme.

4. SOME VARIATIONS OF THE BASIC BOOTSTRAPPING SCHEME

Behind the basic version of historical simulation, with or without bootstrapping, there is the economic hypothesis that future returns follow the same distribution of actual past returns. This can be heroic. Consider Italian or Spanish bonds: in the period 1994-1998 they exhibited high total returns due to the convergence of interest rates in the EMU-block, clearly an exceptional economic and political event. Performing historical simulations using raw data from Italian or Spanish bond markets going back, for instance, five years in the database, means that one includes the "convergence period" and implicitly assumes those exceptional circumstances could repeat again. In particular, the empirical distributions of those bond returns show a rather high mean, the so-called "phantom drift", that not necessarily reflects one's expectations.

Similar problems can arise with regards to volatility. As well known, most financial series exhibit volatility clusters, i.e. large changes in returns are likely to be followed by further large changes. In risk analysis, volatility clusters imply that the probability of a specific loss being incurred is not the same on each period: during periods of higher volatility we will expect larger than usual losses.

Problems can arise also with regard to autocorrelation; a rather common finding is that financial returns are correlated with their own lagged values. From the practitioner point of view, a high serial correlation means, for instance, that large negative changes in returns are likely to be followed by further large negative changes.

There are several techniques to get proper modeling of asset returns that can help to create reasonable future scenarios over medium-long term investment horizons. This scenarios can be used both for tactical and strategic risk analysis.

The following procedure is a variation of the basic bootstrapping method described in the preceding section. It takes into account the possibility that time series exhibit serial correlation. It's a simple non-parametric method that can be used if one is concerned with the possibility that data exhibits autocorrelation.

Let us go back to the example of the preceding section and suppose we have three years of daily log-returns for K risk factors. Our time horizon is one month, e.g. 21 days and we are interested in a relative risk analysis, so we apply (5) to get the "constant mix" time series of portfolio excess returns, er_{PTF}, which is a vector with 750 rows. As before, we assign a position index to each excess return, that is, we assign 750 integer numbers, one for each row. Now we draw a uniformly distributed number from 1 to 750; assume it is 118. We take the 118th element of er_{PTF}. In order to keep the actual autocorrelation we take also the 119th element of er_{PTF}, the 120th element, and so on, till we have collected 21 consecutive daily log-excess returns in a row. That is, we take the whole path of 21 excess returns, starting from the 119th element of er_{PTF}. From the simulated path, it is easy to calculate the realization of portfolio excess return over the required time horizon. Repeating this procedure many times allows us to get an estimate of the future excess returns' distribution. The main limit is that time horizon has to be significantly shorter than raw time series length, in order to have a sample of acceptable size.

From a more formal point of view, if the goal is to simulate M scenarios in terms of excess returns over a time horizon of H periods, for each scenario we generate a (pseudo) integer random number, uniformly distributed in the range between 1 and $N-H+1$:

$$u \sim U[\text{int}(1, N-H+1)]. \tag{8}$$

The excess returns ER over the required time horizon *H* is then given by:

$$ER = \sum_{i=u}^{u+H-1} er(i) \tag{9}$$

where $er(u_i)$ is the element of vector er_{PTF} that is in position u_i. By executing this procedure *M* times, we estimate the distribution of future excess returns taking into account autocorrelation. Of course, one can use this approach; we call it non-parametric autocorrelation adjustment, also with absolute returns. Parallel bootstrapping techniques can be used as well.

When the goal is to incorporate volatility updating in future scenarios, there are several options. Hull and White (1998) show how to take into account volatility clusters into the basic historical simulation method (without bootstrapping), by scaling observations by the ratio of current over past conditional volatility. Barone-Adesi et al. (1999) propose filtering historical simulated portfolio returns (or excess returns) by ARMA-GARCH processes to get i.i.d. residuals. Basically, the ARMA equation for the conditional mean allows modeling any serial correlation, while the GARCH conditional variance equation copes with volatility clusters. Residuals of the ARMA models are divided by the past conditional GARCH volatility in order to get i.i.d. observations. Then one can use (6) and (7) to bootstrap these observations over the desired time horizon. Finally, the paths of i.i.d. observations is used as innovations to simulate the ARMA-GARCH process followed by the historical simulated portfolio returns (or excess returns). The proposed method is actually a flexible parametric-non parametric approach: it is an ARMA-GARCH model where residuals do not follow any predefined distribution. Bootstrapping is applied to i.i.d. residuals, so results are unbiased. McNeil and Frey (1999) propose a similar bootstrapping approach, where the residuals of the ARMA-GARCH model follow an EVT distribution. Of course, instead of using an ARMA model, one could incorporate mean updating in future scenarios using different models, ranging from simple Exponential Weighted Moving Average techniques to structural models. In general, to model the returns r (or excess returns) it is possible to use specifications of the form:

$$r_t = \varphi(x) + \varepsilon_t, \quad \varepsilon_t \sim (0, \sigma_t) \tag{10}$$
$$\sigma_t^2 = \alpha + \beta\varepsilon_{t-1}^2 + \gamma\sigma_{t-1}^2$$
$$z_t = \varepsilon_t / \sigma_t$$

where $\phi(\cdot)$ is some functional form and x is a vector of explanatory variables (observed at time t or lagged), ε_t is the disturbance term with zero mean and standard deviation σ_t, which follows the GARCH(1,1) equation. Bootstrapping is applied to the series of standardized residual, $\{z\}$. They are i.i.d. but, in general, not normal. The simulated standardized residuals are then recursively used in the three equations of the model (10) to generate future returns.

Especially when dealing with long term scenarios, one could also modify the unconditional mean and unconditional volatility of portfolio returns, so that future scenarios reflect a specific market view. This can be done in a straightforward way using the following simple linear transformation:

$$\tilde{r}_{PTF,i} = \left[r_{PTF,i} - \overline{m}\right] \cdot \frac{\sigma_{PTF}}{\overline{\sigma}_{PTF}} + \mu \tag{11}$$

where σ_{PTF} is the expected unconditional volatility, $\overline{\sigma}_{PTF}$ is the historical unconditional volatility, $\overline{m}$ is the historical unconditional mean, μ is the expected unconditional mean; $r_{PTF,i}$ is the (simulated) portfolio return at time *i*, and $\tilde{r}_{PTF,i}$ is the adjusted (simulated) portfolio return at time *i*. Using (11) it is possible, for example, to examine how a pension fund behaves under different market conditions, testing the validity of a given asset allocation.

A final crucial point using historical simulation concerns the use of raw empirical frequency distributions, without fitting a smooth density estimate. This problem may be especially significant as we are dealing with the left tail of the distribution, where the number of observations in the historical sample period may be few. VaR and ReVaR estimates based on small tail samples have a relatively high variance. Figure 7A and Figure 7B clearly show this phenomenon for the daily log-returns of MSCI World in Euro terms over the period September 1971 – May 2000. The shorter the sample, the greater is this phenomenon.

[insert figures 7a – 7b]

In order to lower the impact of this problem, one could use the ARMA-GARCH-EVT approach proposed by McNeil et al., mentioned before. As we have seen, this parametric approach is based on fitting an EVT function to the empirical tails of the distribution. Alternatively, limiting ourselves to non-parametric models, it is possible to estimate the pdf of portfolio returns (or excess returns) using a smooth density estimator. To do so one can use a kernel estimator, that can be thought of as a way of smoothing a histogram constructed with the sample data. See Silverman (1986) for a detailed description. A suitable alternative to the kernel estimator is given by the RTDM distribution of Ramberg et al. (1979). It is a probability distribution based on the first four moments, and can cope with a wide variety of curve shapes.

Smooth density estimators can be used, in principle:

1. before boostrapping (i.e. one draw the returns form the smoothed pdf);
2. after bootstrapping (i.e. returns are drown from raw data, one get the returns over the desired time horizon, then estimate a smooth density function);
3. before and after bootstrapping.

The use of smoothing density estimator before bootstrapping forces us to modify the bootstrap procedure as follows. If we have a smooth pdf estimate and a cumulative function F, assuming we want to simulate M excess returns over a time horizon H, our first step is to generate H pseudo random numbers uniformly distributed in the range between 0 and 1:

$$u_i \sim U[0,1], \quad i = 1, 2, 3, \ldots, H \,. \tag{12}$$

Then, for each u_i, we apply the inverse transformation:

$$er(u_i) = F^{-1}(u_i) \tag{13}$$

where F^{-1} is the inverse of the cumulative function. In each simulation run, the excess return ER over the period *H* is given by:

$$ER = \sum_{i=1}^{H} er(u_i) \tag{14}$$

and the procedure must be repeated *M* times. The main drawback of this method is that the time structure of data (e.g. serial correlation) is lost.

5. OPTIONS AND MEDIUM-LONG TERM INVESTMENT HORIZONS

If the investment horizon is medium to long, options must be treated carefully. The effect of time can be sensible and must be taken into account. The asymmetric, nonlinear nature of options is particularly evident. In this situation the use of methods based on Taylor series expansion can be misleading, obfuscating the real risk related to options. Delta-gamma methods are unstable for large asset price changes, which can occur if the investment horizon is long. Basic historical simulation, as well, is not appropriate, unless the chosen time horizon coincides with data frequency. Otherwise it can lead to wrong results: with daily data one could calculate daily option prices, then daily option returns, find the desired quantile (daily VaR) and then scale it with the square root of time horizon. For instance, assume that we have our daily VaR for a call option; is -23%. If time horizon is one month (e.g. 21 business days), our monthly VaR is

$$-23\% \cdot \sqrt{21} = -105.4\% \,.$$

This is clearly impossible because this implies we lose more than the premium.

Parallel bootstrapping procedures, instead, can cope with options in a very natural way. In order to get the distribution of option returns over the chosen time horizon we need the current option price and the option price at horizon. The latter depends on the value of the relevant variables at horizon, essentially the underlying, the volatility and the risk-free rate for plain vanilla options. The basic idea is that the underlying price paths (including the final value) are obtained from the corresponding return paths. The bootstrap procedure is also used to get the path for each relevant market variable, e.g. volatility, interest rates. Then, using an option pricing model applied to each path and other relevant option parameters (e.g. strike price, time to maturity) one gets the distribution of option prices at horizon. Together with the current option price, they allow us to calculate the probability distribution of option returns over the required time horizon.

From a more formal point of view, let us assume we are at time t, dealing with an European option, with price given by $f(S_t, K, T, r_t, \sigma_t)$, where S_t is the underlying asset price at current time, K is the strike price, T is the time to maturity, r_t is the risk-free interest rate, and σ_t is the volatility. If the goal is to calculate a risk measure over the time horizon *H*, we run *M* simulations and generate *M* option total returns. The generic option total return has the form:

$$f(S_{t+H}, K, T-H, r_{t+H}, \sigma_{t+H}) / f(S_t, K, T, r_t, \sigma_t) - 1 \tag{15}$$

where $T-H$ is the time to maturity at horizon, and S_{t+H}, r_{t+H}, σ_{t+H} are calculated using the corresponding bootstrapped returns. The estimate σ_{t+H} can be obtained alternatively using a volatility forecasting model, therefore following the approach proposed by Barone-Adesi, Giannopoulos and Vosper (for the sake of brevity BAGV), or directly from implied volatilities, using bootstrapping techniques.

At this point the distribution of the option total returns over the required time horizon has to be combined with the rest of the portfolio. It is worth to stress that in order to reflects the actual historical correlations the uniform variates used in the parallel bootstrapping procedure must be the same for all the risk factors involved. This means that if the first path of the option underlying returns is obtained using a vector of uniform variates, then the first path of all the variables bootstrapped is obtained using the same vector.

For example, let us consider a portfolio made of an option contract, with weight ω_{option}. The rest of the portfolio (we call it "ex-option portfolio") has weight ω_{stocks} and is invested in several stocks (i.e. linear instruments). Our goal is to calculate the portfolio's VaR, with time horizon *H*. In order to get the j-th ($j = 1,2,3,...,M$) of M simulations, we can use the following procedure:

- generate a vector of *H* (pseudo) integer random numbers, uniformly distributed in the range between 1 and *N*; where *N* is the length of the time series available;
- use this vector to get a strip of *H*:
    1. ex-option portfolio returns;
    2. underlying returns (and consequently the levels);
    3. interest rates
    4. volatilities
- using the strips of underlying levels, interest rates and volatilities, put their terminal value into (15), together with the current option price, to get the option's total return in the j-th simulation run, $TR_{option,j}$;
- the sum of the H ex-option portfolio returns is the ex-option portfolio's return in the j-th simulation run, $TR_{ex-option,j}$;
- finally the overall portfolio's total return is given by:

$$TR_j = \omega_{option} \cdot TR_{option,j} + \omega_{ex-option} \cdot TR_{ex-option,j};$$

- repeat this till you have completed *M* simulation runs; at this point you have the (non-smooth) estimate of the portfolio's return distribution.

The parallel bootstrapping procedure can be easily extended to deal with all the option contracts in the portfolio. The results obtained with this technique are clearly based on the full portfolio revaluation and reflect the assets' actual multivariate distribution. Figure 8 is an example of application of the procedure described above. It shows the probability distribution of total returns of a European call option on MIB 30, over a monthly investment horizon. The option is 10% out of the money and will expire in 6 months. Applying the BAGV model to MIB 30 (the underlying) and Libor ITL 3 month (the assumed risk-free rate), using daily data from 13/06/97 to 14/06/00 (source Datastream International) we obtain the option probability distribution. The pricing model is Black & Scholes.

[Insert Figure 8]

Problems arise if one uses a smoothing density estimator before parallel bootstrapping. In this case the time dimension of data is lost and parallel-bootstrapping techniques can not work. Therefore, in order to take into account the correlation one could use multivariate simulation techniques such as the algorithm proposed by Stein (19).

6. COMPARISON OF VALUE AT RISK ESTIMATES AND ACCURACY TESTS

We test the different methodologies we have seen so far. The goal is to see how they behave when put at work. We assume the tactical point of view. We compute VaR[2] with confidence level equal to 95% and 1 month horizon for several total return indices: MSCI USA, MSCI Far East, MSCI EMF, MSCI Italy, MSCI Consumer Goods, MSCI Capital Equipment, J.P. Morgan US Govt. Bond, J.P. Morgan Italy Govt. Bond, MSCI World, J.P. Morgan Global Govt. Bond. We use daily data (source Datastream International) from 1/1/1990 to 5/5/2000.

[2] Although we are more interested in relative risk indicators, like ReVaR, in our test we use VaR because i) we can work on indices, without creating arbitrary portfolios; ii) indices allow us to backtest results in a simpler way; iii) constant-mix portfolio excess returns and total returns are both linear combinations of individual securities returns, so there are no reasons to think they should exhibit different statistical properties.

We test the methodologies listed in Table 4. Basically we test different bootstrapping methods: basic or "plain vanilla", with autocorrelation adjustment, with ARMA-GARCH filtering (BAGV model), with the Hull and White volatility scaling using a GARCH model (for the sake of brevity, HW model)[3]. If technically possible, we also test the use of kernel estimates of the distribution of returns before and/or after the bootstrap. Following Butler and Schachter (1996), we use a standard Gaussian kernel estimator. Note that the use of a Gaussian kernel does not imply any hypothesis of normally distributed returns, as data smoothing can be done with any continuous distribution. The benchmark methods are the standard parametric normal methods, based on a variance-covariance approach, and the basic historical simulation. With these two methods, in order to switch from daily-based metrics to monthly-based metrics, we use the square root rule.

[insert Table 4]

The test is as follows. We use a rolling window of 18 months of daily data to estimate VaR for the next month. Then each day we test, out of sample, whether the actual past monthly return is smaller then the corresponding estimated VaR. Then we repeat the procedure, rolling out the sample window one day. We judge the different methodologies purely on the basis of their risk-forecasting capabilities. First of all, VaR should accurately predict the true amount of risk, on average. For instance, if confidence level is 95% we expect to see that actual returns are smaller than VaR only 5% of the times. This is a test of unconditional accuracy. In order to test the hypothesis that VaR forecasts are unconditionally correct, we use the binomial test, discussed by Kupiec (1995). As pointed out by Lopez (1999), this test is currently the standard chosen by regulators. The second issue we are concerned with is conditional accuracy of VaR estimates, i.e. they should be correct at each point in time and should not be serial correlated. That is, if at any given time the model is wrong, this should not affect the next forecast. To test this hypothesis one could use the approach proposed by Christofferesen (1998). Note, however, that we

[3] In practice, whenever we have to use a GARCH(1,1) model we adopt an IGARCH(1,1) model. This means that we estimate volatilities with a standard Exponential Moving Average Approach. Having to crunch a lot of data, this shortcut reduces our computational effort, apparently without major practical drawbacks.

calculate VaR estimates using overlapping samples, so, for two subsequent risk measures, the information set is almost equal. Therefore, VaR estimates tend to exhibit serial correlation and there are not the conditions to use this kind of test.

Table 5 shows the fraction of times actual returns are smaller than VaR estimates with confidence level equal to 95%. We denote this fraction with α^*. When α^* is in the range [3.9%, 6.2%], the corresponding VaR estimator is unconditionally correct at a 1% significance level. Examining the results we notice that many risk indicators are not unconditionally accurate. Basically, this is due to the relatively long time horizon we use. As concerns the test of unconditional accuracy, basic historical simulation (M1) and BAGV (M4) are the two methodologies with the highest number of successes in passing the test.

[insert Table 5]

It is clear that the gaussian VaR (M2) performs in a poor way, if compared with VaR measures based on historical-simulation and bootstrapping. In order to sort out in a rough way the "best" VaR estimates, we calculated the mean absolute deviation of α^* from 5%, across all the indices. Results, reported in Table 6, show that basic historical simulation (M1) and basic historical simulation using kernel density estimator of the pdf (M10) are the best performers, followed by the BAGV model (M4) and the BAGV model enhanced by the kernel estimator (M11). Unfortunately, as pointed out in section 5, if our portfolio contains options both M1 and M11 are not appropriate. Conversely, the BAGV model is well suited to cope with options.

[insert Table 6]

A good risk indicator should also signal changes in market risk, helping asset management companies to predict, to some extent, extreme moves in asset returns, reacting properly. This is particularly true in the asset management field, where risk indicators should help portfolio managers to calibrate their market exposures in order to outperform the benchmark and the competitors. For instance, if ReVaR goes up, a portfolio manager is warned of the danger that her or his information ratio is going to deteriorate (other things being equal), so a sensible decision could be to reduce relative risk. The most responsive risk estimates are those based on the BAGV model, as shown by Figure 9a and 9b.

[Insert Figure 9a - 9b]

This fact is also apparent when we examine MSCI EMF (see Figure 10). Although none of the risk estimates is unconditionally correct, the BAGV model seems to respond well to changes in market risk.

[Insert Figure 10]

7. CONCLUSIONS

We examine and test several methods of calculating risk measures in the asset management field, when the horizon is relatively long. They are based on historical simulation and its variations, like bootstrapping. The historical simulation framework takes into account current market conditions and is based on the empirical multivariate distribution of the data without imposing any particular probability function, or compressing the tails, or altering the skewness of portfolio returns. In particular, simulation methodologies based on bootstrapping allow for a reliable calculation of several risk measures for a large number of portfolios of any dimension. Although basic historical simulation (enhanced by the use of kernel estimator, too) works well for linear portfolios over medium term horizons, the most promising model seems to be BAGV, especially if the portfolio contains derivatives. It allows us to generate medium-term to long-term financial scenarios, taking into account heteroskedasticity, serial correlation (if relevant), and specific market views (if these views exist). This simulation methodology provides fast evaluation of relative and absolute risk measures, taking into account current market conditions. It does not rely on the knowledge of the correlation matrix of security returns. According to our tests, this methodology provides satisfactory assessments of portfolio tactical risk and can inform in advance portfolio managers of changes in market risk. The bootstrapping approach is also promising for creating financial scenarios for strategic risk analysis.

Acknowledgements

We gratefully acknowledge helpful comments, suggestions and full support from Dario Brandolini and Andrea Beltratti. The authors are responsible for any errors and inaccuracies.

References

1. Barone-Adesi, G., Giannopoulos, K., Vosper, L. (1999). VaR Without Correlations for Portfolios of Derivative Securities. *Journal of Futures Markets*, August. Available at: http://www.lu.unisi.ch/istfin/papers/index.htm.
2. Christoffersen. P.F. (1998). Evaluating interval forecast. *International Economic Review*, 39.
3. Butler, J.S., Schachter, B. (1996). Improving value-at-risk estimates by combining kernel estimation. *Conference on Bank Structure and Competition*. Federal Reserve Bank of Chicago. Proceedings.
4. Efron, B. and Tibshirani, R. (1998), *An Introduction to the Bootstrap*, 2nd edn. Chapman & Hall.
5. Falloon, W. (1999). Growin'up. *Risk*, February. pp. 26-31.
6. Hull, J., and White, A. (1998). Incorporating volatility updating for value-at-risk. *Journal of Risk*, Vol. 1, No. 1.
7. Kupiec, P. (1995). Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives*, 3.
8. Lopez, J.A. (1999). Regulatory evaluation of value at risk models. *Journal of Risk*, Vol. 1, No. 2.
9. McNeil, A.J., Frey, R. (1999). Estimation of tail-related risk measures for heteroskedastic financial time series: an extreme value approach. ETHZ, working paper.
10. Ramberg, J.S., Tadikamalla, P.R., Dudewicz, E.J., Mykytka, E.F. (1979). A probability distribution and its uses in fitting data. *Technometrics*, Vol. 21., No. 2. May.

11. RiskMetrics Group. (1999). *LongRun Technical Document*. 1st edn. Available at: http://www.riskmetrics.com.

12. Silverman, B.W. (1986). *Density Estimation for Statistics and Data Analysis*, Chapman & Hall.

13. Stein, M. (1987). Large sample properties of simulations using latin hypercube sampling. *Technometrics*, Vol. 29, No. 2. May.

14. Zenti, R.. (2000). Some statistical facts about monthly returns: a short note. Ras Asset Management SGR, working paper. Available on request.

Figure 1

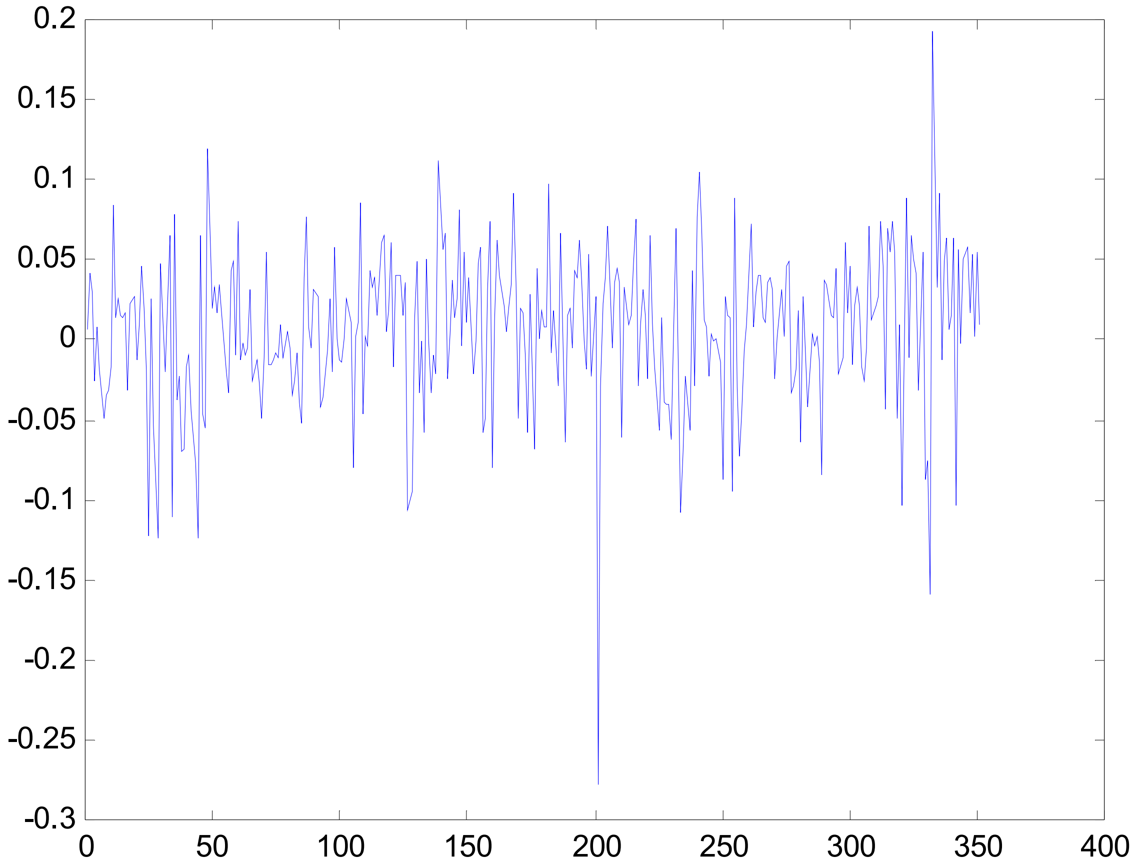

Historical plot of monthly log-returns of MSCI World in Euro terms (source Datastream International), over the period March 1971 – May 2000. The plot shows that the mean is rather stable while variance is non-stationary.

Figure 2

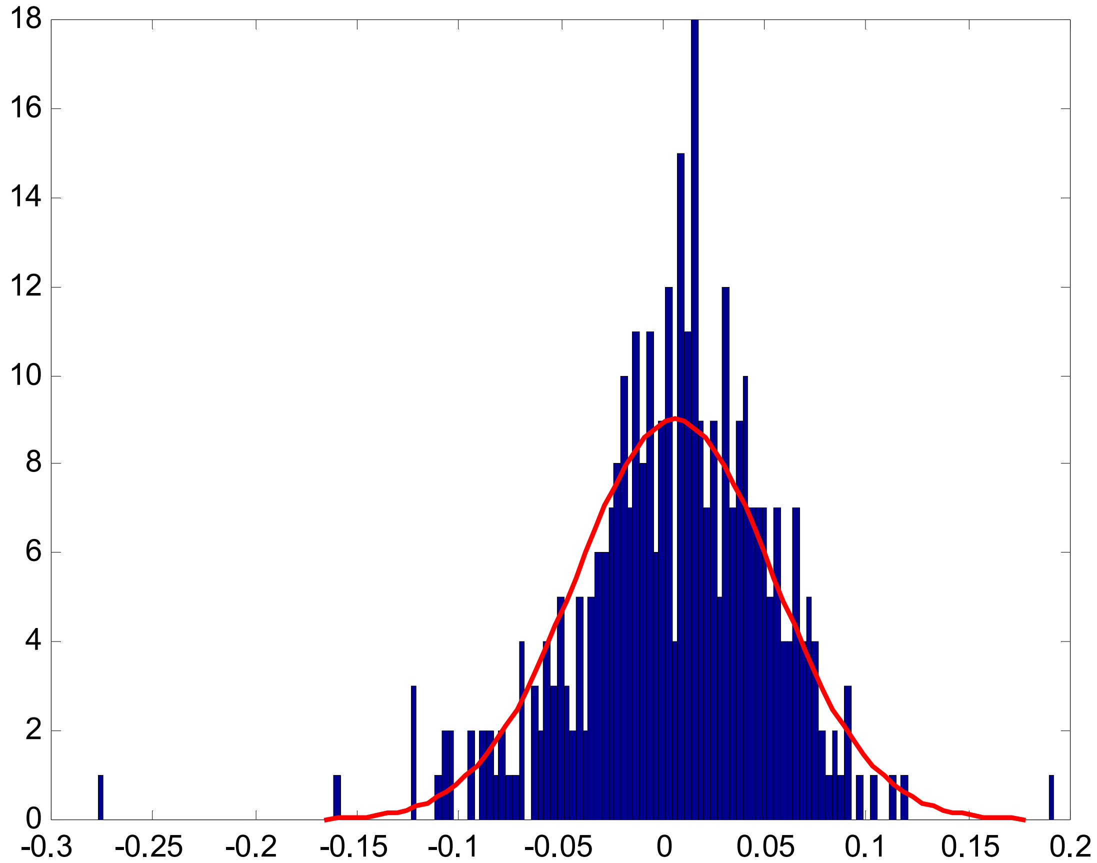

Histogram, with superimposed fitted normal density of monthly log-returns of MSCI World in Euro terms (source Datastream International), over the period March 1971 – May 2000. The histogram shows that the pdf is rather asymmetric with presence of tail events.

Figure 3

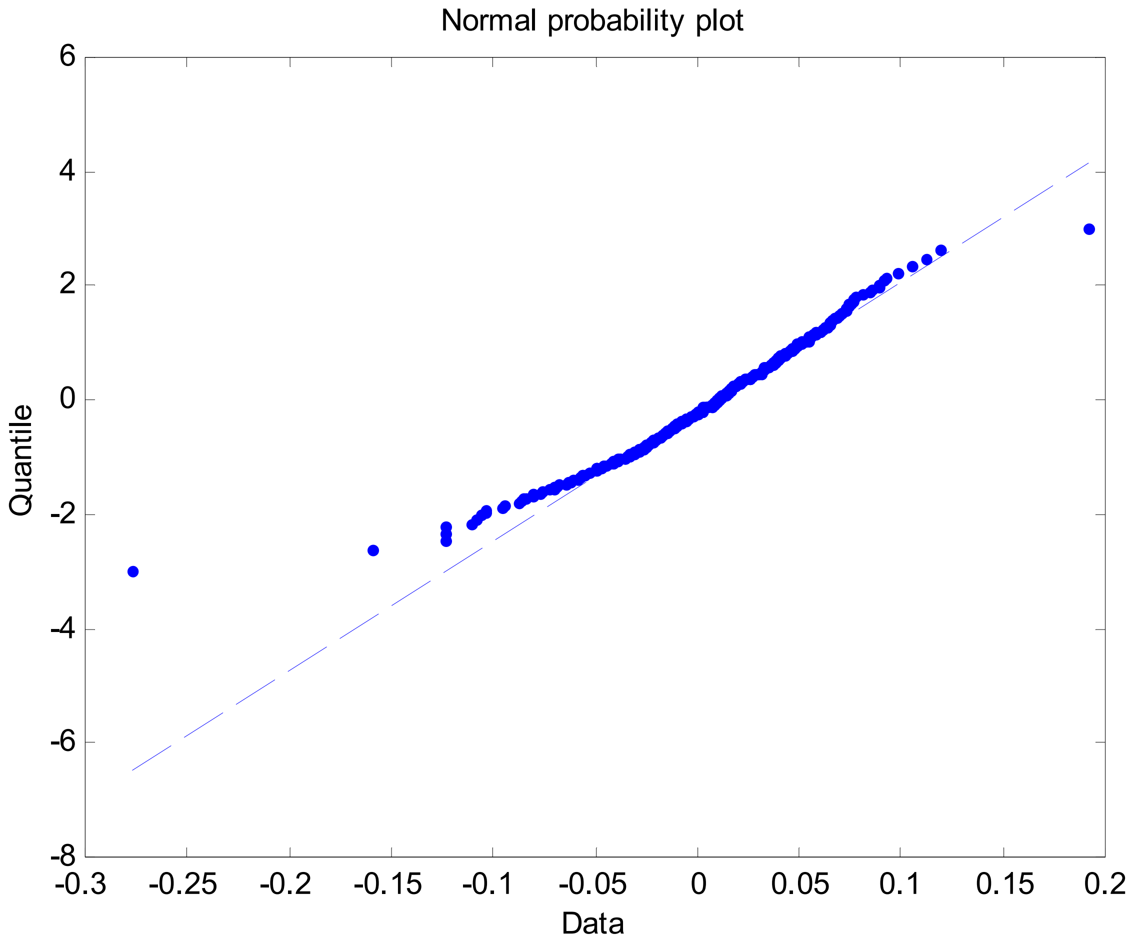

Quantile-Quantile plot of monthly log-returns of MSCI World in Euro terms (source Datastream International), over the period March 1971 – May 2000, on the horizontal axis, versus the standard normal distribution, on the vertical axis. If the returns were normal, the plot should look roughly linear. The left side of plot curves up, indicating the presence of a heavy left tail in the index distribution.

Table 1

	Left limit	Point estimate	Right limit
Mean	0.0007	0.0060	0.0110
Standard deviation	0.0434	0.0487	0.0550
Skewness	-1.5853	-0.7817	0.0759
Kurtosis	3.0468	6.4546	9.9425
	Test	P-value	
Jarque-Bera	210.2830	0.0000	

Descriptive statistics of monthly log-returns of MSCI World in Euro terms (source Datastream International), over the period March 1971 – May 2000. The left limit and the right limit are the extremes of a 95% confidence interval calculated using the Efron's percentile method, with 1000 resamples. The statistics shows that the empirical distribution has a mean which differs significantly form zero, is negatively skewed, with a kurtosis significantly higher than 3. Jarque-Bera is a test statistic for testing whether the series is normally distributed. Under the null hypothesis of a normal distribution, the Jarque-Bera statistic is distributed as a χ^2 with 2 degrees of freedom. The reported P-value is the probability that a Jarque-Bera statistic exceeds (in absolute value) the observed value under the null: a small probability value leads to the rejection of the null hypothesis of a normal distribution.

Figure 4

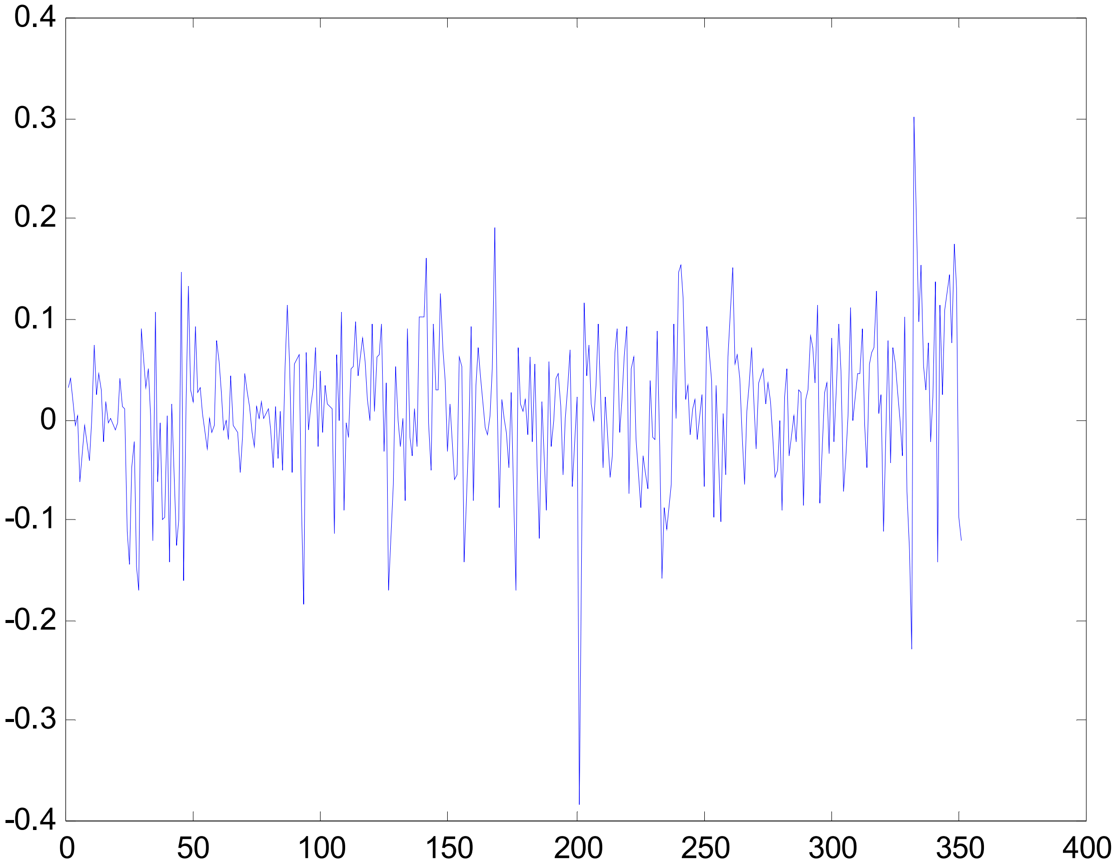

Historical plot of monthly log-returns of Nasdaq Composite in Euro terms (source Datastream International), over the period March 1971 – May 2000. The plot shows that the mean is rather stable while variance is non-stationary.

Figure 5

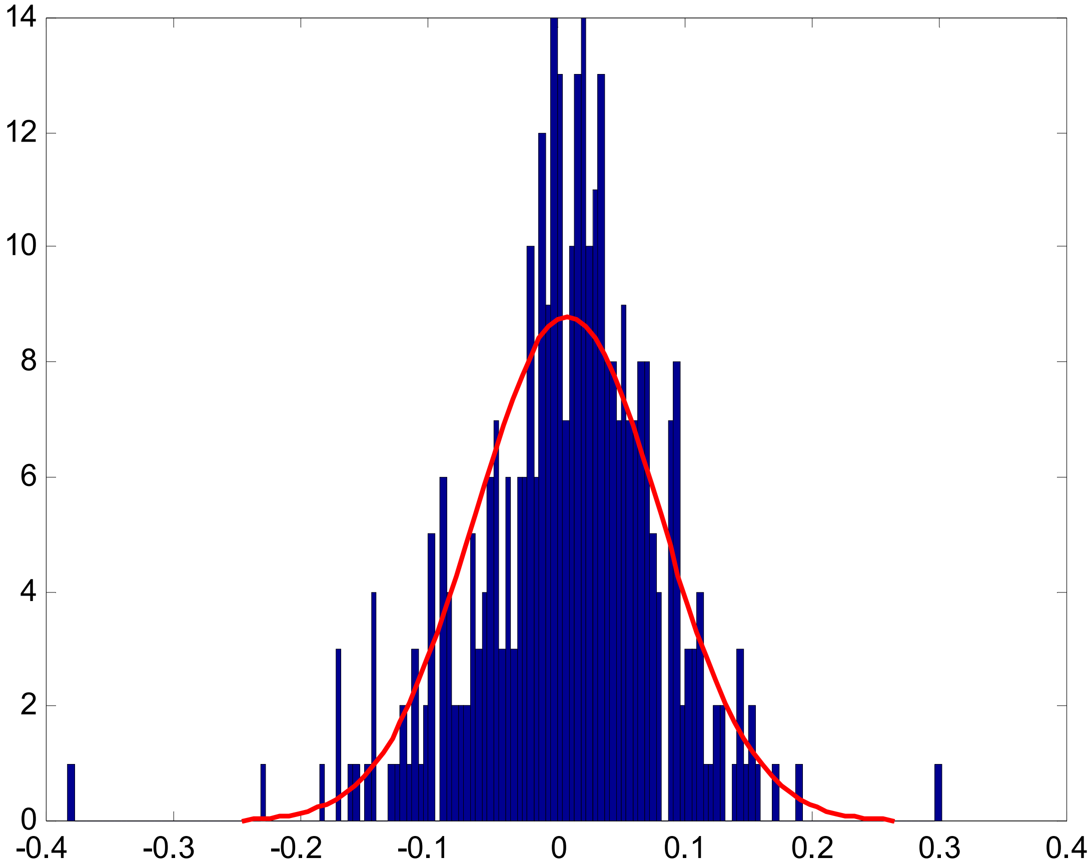

Histogram, with superimposed fitted normal density of monthly log-returns of Nasdaq Composite in Euro terms (source Datastream International), over the period March 1971 – May 2000. The histogram shows the existence of tail events.

Figure 6

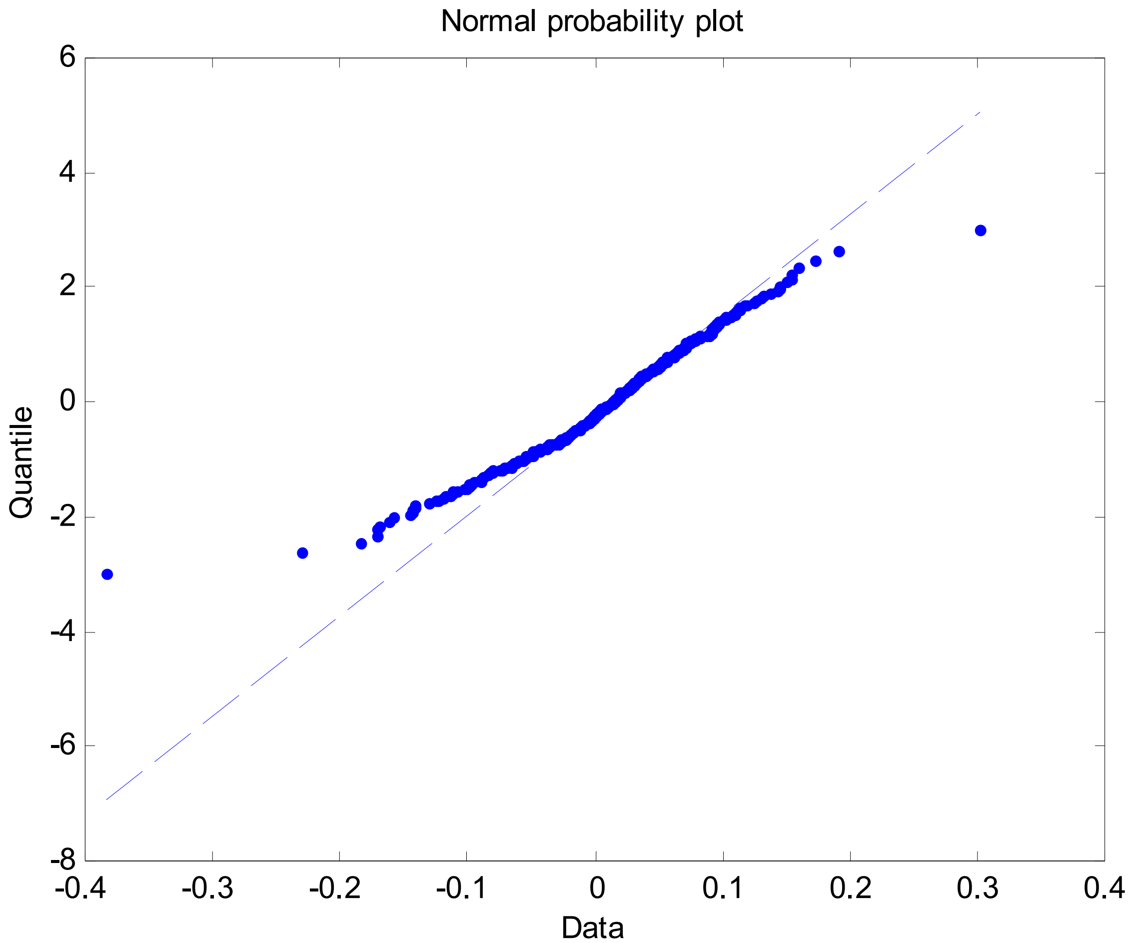

Quantile-Quantile plot of monthly log-returns of Nasdaq Composite in Euro terms (source Datastream International), over the period March 1971 – May 2000, on the horizontal axis, versus the standard normal distribution, on the vertical axis. If the returns were normal, the plot should look roughly linear. The left side of plot curves up and the right side of the plot curves down, indicating the both the tails of the index distribution are heavy.

Table 2

	Left limit	Point estimate	Right limit
Mean	0.0013	0.0088	0.0164
Standard deviation	0.0640	0.0727	0.0812
Skewness	-1.2608	-0.5616	0.1888
Kurtosis	3.1587	5.7800	8.5026
	Test	P-value	
Jarque-Bera	131.4829	0.0000	

Descriptive statistics of monthly log-returns of Nasdaq Composite in Euro terms (source Datastream International), over the period March 1971 – May 2000. The left limit and the right limit are the extremes of a 95% confidence interval calculated using the Efron's percentile method, with 1000 resamples. The statistics shows that the empirical distribution has a mean which differs significantly form zero, is negatively skewed, with a kurtosis significantly higher than 3. Jarque-Bera is a test statistic for testing whether the series is normally distributed. Under the null hypothesis of a normal distribution, the Jarque-Bera statistic is distributed as a χ^2 with 2 degrees of freedom. The reported P-value is the probability that a Jarque-Bera statistic exceeds (in absolute value) the observed value under the null: a small probability value leads to the rejection of the null hypothesis of a normal distribution.

Table 3

	Estimated monthly returns	Actual monthly returns
Mean	0.006	0.006 *[0.002, 0.010]*
Standard deviation	0.048	0.049 *[0.044, 0.054]*
Skewness	-0.694	-0.779 *[-1.464, -0.095]*
Kurtosis	5.680	6.394 *[3.280, 9.007]*

Table 3 compares the first four moments of the estimated monthly return distribution of the equity index MSCI World (in Euro terms, source Datastream), using a basic bootstrapping approach, and the first four moments of the actual monthly return. Both daily and monthly returns are calculated over the period 8/9/1971 to 8/5/2000. The estimated moments are not significantly different (with a 90% confidence level) from the actual moments. In brackets we show the left limit and the right limit of the 90% confidence interval, which is calculated using the Efron's percentile method, with 1000 resamples.

Figure 7a

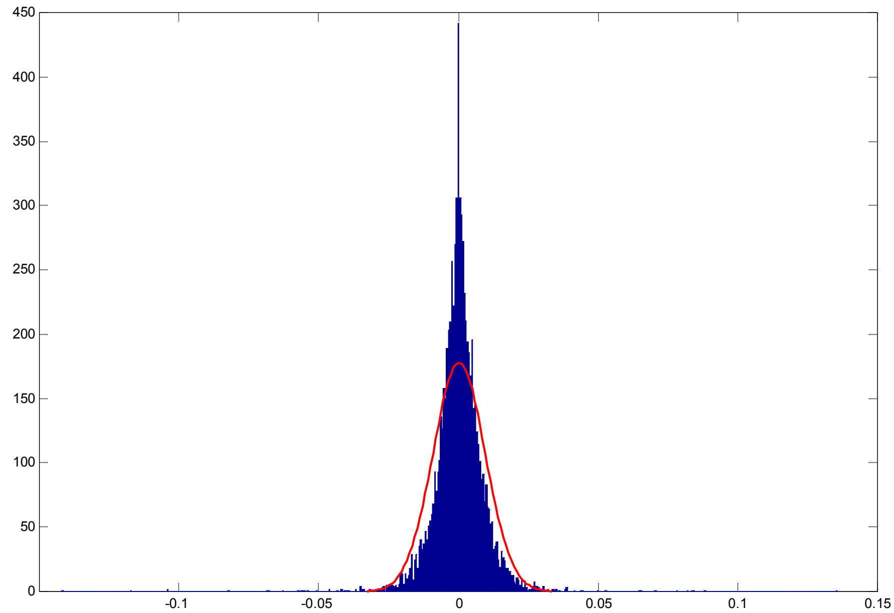

Figure 7b

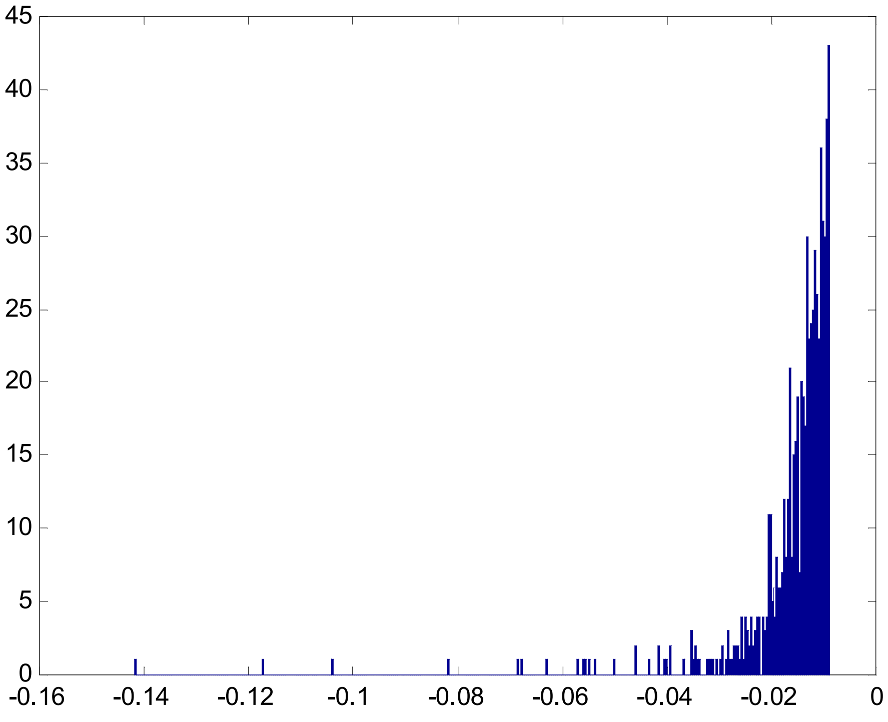

Figure 7a and Figure 7b show, respectively, the histogram with superimposed fitted normal density of daily log-returns of MSCI World in Euro terms (source Datastream International) over the period September 1971 – May 2000, and the histogram of the worst 10% of the same data, i.e. a zoom on the left tail of the distribution. Figure 7b clearly shows how the number of observations tends to decline moving away from the center of the distribution. This means that tail observations are usually a small sample and that the estimators based on that sample (i.e. using basic historical simulation) have a relatively high variance.

Figure 8

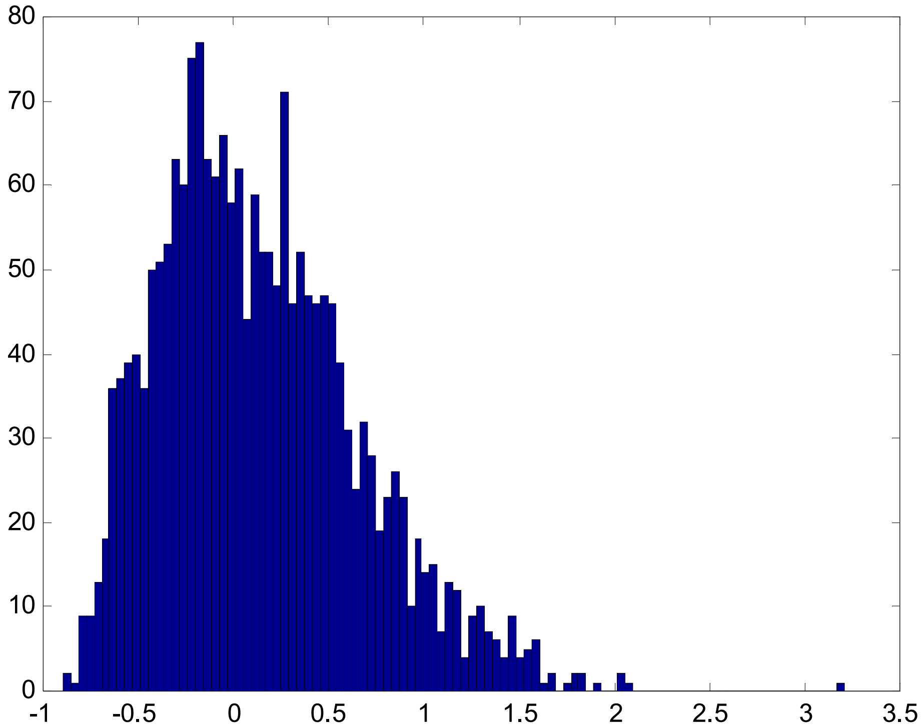

Probability distribution of the total returns of a European call option, over a monthly investment horizon. The underlying is MIB 30 (in Euro terms), the option is 10% out of the money, and time to maturity is 6 months. The probability distribution is obtained applying the BAGV bootstrapping technique to MIB 30 (underlying) and Libor ITL 3 month. Data range is 13/06/97 - 14/06/00. The pricing model is Black & Scholes. The distribution is clearly asymmetric and is bounded to the left by -1. This reflects the fact that the worst case is represented by the loss of the entire premium.

Table 4

Name	Description of VaR estimation methodology	Kernel pdf estimate of daily data	Calculation of monthly risk measures	Basic bootstrapping	Bootstrapping with autocorrelation adjustment	BAGV (ARMA-GARCH filter)	Use of kernel on monthly (estimated) data	HW volatility scaling (GARCH filter)
M1	Basic historical simulation	No	Square root rule applied to daily risk measures	-	-	-	No	No
M2	Parametric normal	No	Square root rule applied to daily risk measures	-	-	-	No	No
M3	Basic bootstrapping	No	Bootstrapping	Yes	No	No	No	No
M4	Bootstrapping using the BAGV model	No	Bootstrapping	No	No	Yes	No	No
M5	Bootstrapping with non-parametric autocorrelation adjustment	No	Bootstrapping	No	Yes	No	No	No
M6	Bootstrapping with non-parametric autocorrelation adjustment and HW volatility scaling	No	Bootstrapping	No	Yes	No	No	Yes
M7	Bootstrapping of daily data, previously smoothed using a kernel density estimator to fit the pdf	Yes	Bootstrapping	Yes	No	No	No	No
M8	Bootstrapping of daily data previously (i) adjusted with "constant" volatility scaling; (ii) smoothed using a kernel estimator	Yes	Bootstrapping	Yes	No	No	No	Yes
M9	Basic historical simulation, using kernel density estimator to fit the pdf	Yes	Square root rule applied to daily risk measures	-	-	-	Yes	No
M10	Basic bootstrapping with kernel density estimate of the pdf of monthly returns	No	Bootstrapping	Yes	No	No	Yes	No

M11	Bootstrapping using the BAGV model, with kernel density estimate of the pdf of monthly returns	No	Bootstrapping	No	No	Yes	Yes	No
M12	Bootstrapping with non-parametric autocorrelation adjustment, with kernel density estimate of the pdf of monthly returns	No	Bootstrapping	No	Yes	No	Yes	No
M13	Bootstrapping with non-parametric autocorrelation adjustment and HW volatility scaling, with kernel density estimate of the pdf of monthly returns	No	Bootstrapping	No	Yes	No	Yes	Yes
M14	Bootstrapping of daily data previously smoothed using a kernel estimator, with kernel density estimate of the pdf of monthly returns	Yes	Bootstrapping	Yes	No	No	Yes	No
M15	Bootstrapping of daily data previously (i) adjusted with "constant" volatility scaling; (ii) smoothed using a kernel estimator. Kernel density estimate of the pdf of monthly returns	Yes	Bootstrapping	Yes	No	No	Yes	Yes

Table 4 describes all the methodologies we analyzed and tested.

Table 5

	MSCI WORLD	MSCI USA	MSCI EUROPE	MSCI JAPAN	MSCI ITALY	MSCI EMF	NASDAQ COMP.	JPM GLOBAL GOVT.	JPM US GOVT.BOND	JPM GERMAN GOVT.BOND	JPM ITALY GOVT.BOND
M1	6.1%*	3.6%	6.7%	5.7%*	5.7%*	10.7%	6.2%*	2.3%	1.9%	4.1%*	5.7%*
M2	8.7%	6.0%*	9.1%	5.2%*	6.8%	14.7%	9.7%	6.1%*	4.7%*	10.5%	13.5%
M3	8.7%	6.2%*	8.8%	5.3%*	6.7%	14.4%	9.5%	6.0%*	4.7%*	10.3%	13.6%
M4	8.1%	5.8%*	6.0%*	5.3%*	7.7%	11.1%	7.8%	3.9%*	4.1%*	3.6%	5.3%*
M5	6.0%*	5.5%*	9.1%	6.4%	6.3%	8.4%	8.0%	6.4%	6.9%	8.3%	9.7%
M6	8.2%	8.3%	10.1%	6.6%	7.5%	10.3%	9.1%	8.5%	8.9%	9.2%	10.3%
M7	7.3%	4.8%*	7.5%	3.9%*	5.5%*	13.0%	8.5%	5.3%*	3.3%	8.5%	11.8%
M8	9.9%	7.3%	9.9%	6.6%	7.9%	15.5%	10.2%	7.3%	5.3%*	10.7%	13.0%
M9	6.1%*	3.4%	6.6%	5.7%*	5.4%*	11.0%	6.3%	2.2%	1.8%	4.0%*	5.7%*
M10	8.7%	6.0%*	8.9%	5.3%*	6.6%	14.4%	9.4%	6.0%*	4.7%*	10.2%	13.6%
M11	8.1%	5.7%*	6.6%	5.7%*	7.6%	11.0%	7.8%	3.8%	4.1%	3.5%	5.3%*
M12	6.1%*	5.5%*	9.0%	6.3%	6.3%	8.3%	8.0%	6.4%	6.9%	8.4%	9.6%
M13	8.1%	8.2%	10.0%	6.6%	7.5%	10.0%	9.0%	8.5%	8.8%	8.8%	10.3%
M14	7.3%	4.6%*	7.5%	3.9%*	5.4%*	12.9%	8.4%	5.2%*	3.3%	8.5%	11.8%
M15	9.8%	7.3%	9.9%	6.5%	7.9%	15.4%	10.0%	7.2%	5.1%*	10.6%	12.9%

Table 5 shows the fraction of times actual returns are smaller than one-month VaR estimates, with confidence level equal to 95%. We test unconditional accuracy using the binomial test discussed by Kupiec (1995) and Lopez (1999). The test is based on the following likelihood ratio statistic:

$$LR(\alpha) = 2\left[\ln\left(\alpha^{*n}\left(1-\alpha^*\right)^{N-n}\right) - \ln\left(\alpha^n\left(1-\alpha\right)^{N-n}\right)\right]$$

where $\alpha = 5\%$, N is the sample size (2319 daily observations), n is the number of times portfolio returns exceed VaR estimates, and $\alpha^* = n/N$. This statistic is asymptotically distributed as a $\chi^2(1)$. In our specific case, when α^* is in the range [3.9%, 6.2%], the corresponding VaR estimator is unconditionally correct at a 1% significance level and has an asterisk as superscript.

Table 6

M1	*M2*	*M3*	*M4*	*M5*	*M6*	*M7*	*M8*	*M9*	*M10*	*M11*	*M12*	*M13*	*M14*	*M15*
1.82%	3.69%	3.62%	1.89%	2.37%	3.83%	2.76%	4.42%	1.85%	3.59%	1.97%	2.33%	3.70%	2.75%	4.34%

Table 6 contains the mean absolute deviation of α^* from $\alpha = 5\%$, across all the indices examined. This is a rough but intuitive way to pick-up the best estimators: they are those with the smallest mean absolute deviations.

Figure 9a

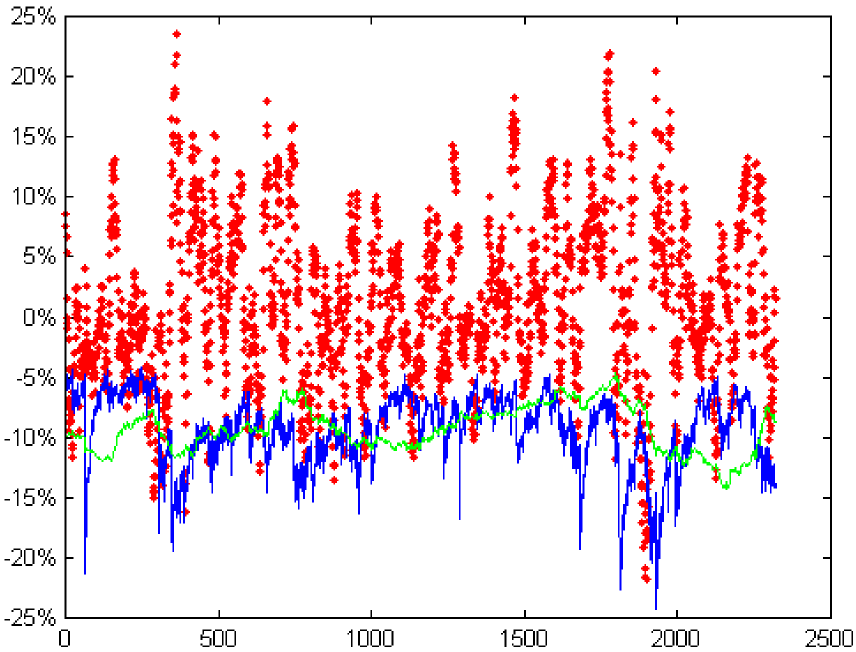

Figure 9b

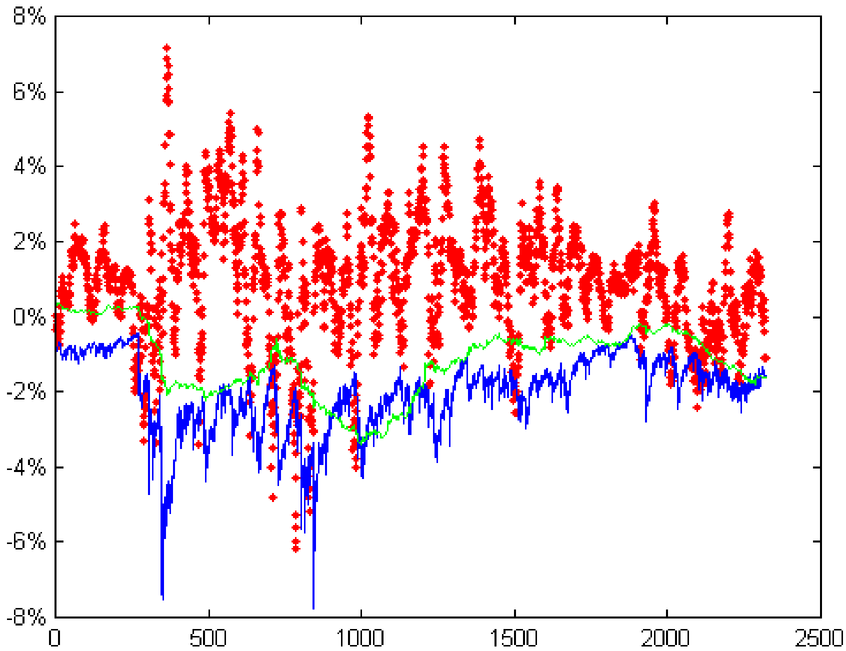

Figure 9a shows out-of-sample VaR estimates, with time horizon equal to 1 month and confidence level 95%, for MSCI ITALY. We used daily log-returns from 1/1/1990 to 5/5/2000 and a rolling window of 18 months of daily data to estimate VaR for the next month. Each day we calculate the actual past monthly return, too. Then we repeat the procedure, rolling out the sample window one day. VaR estimates are plotted against the corresponding actual monthly returns of MSCI ITALY, the red dots. The blue line corresponds to VaR out-of-sample estimates with the historical simulation-bootstrapping method, while the green line corresponds to standard normal-VaR, calculated with equally weighted volatility estimates.

In a similar way, Figure 9b shows the same kind of VaR estimates for the J.P. Morgan Italy Government Bond Total Return Index.

In both cases it is rather evident that the VaR estimator calculated by bootstrapping is better than the parametric VaR estimator, because it is more reactive to different market risk conditions and hence has more signaling power. This can help warn portfolio managers of large price moves.

Figure 10

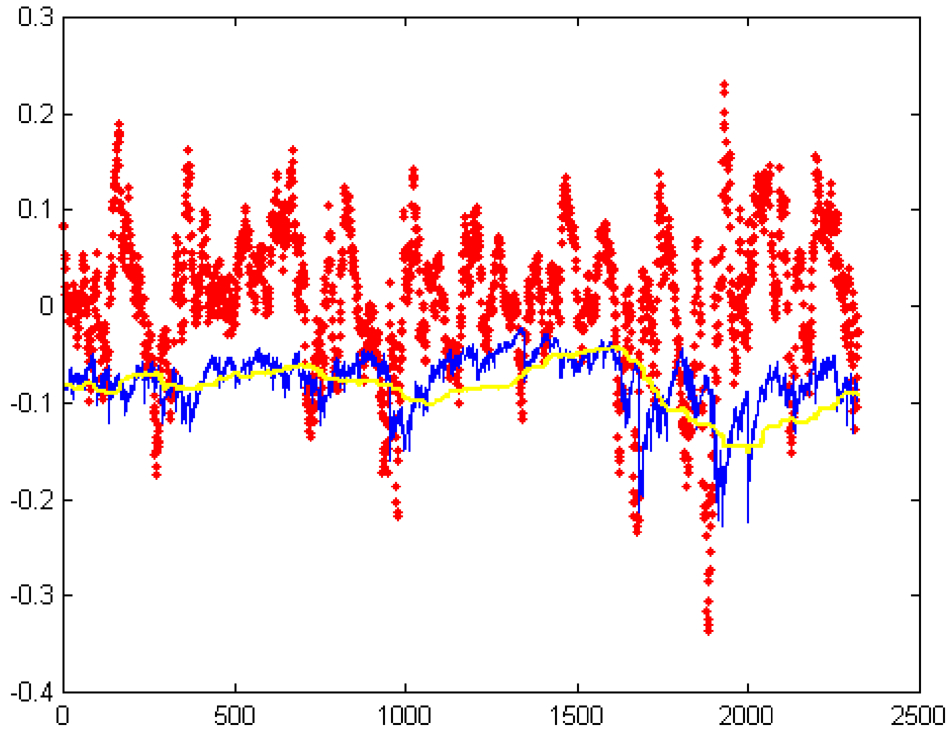

Figure 10 shows out-of-sample VaR estimates, with time horizon equal to 1 month and confidence level 95%, for MSCI EMF. We used daily log-returns from 1/1/1990 to 5/5/2000 and a rolling window of 18 months of daily data to estimate VaR for the next month, Each day we calculate the actual past monthly return, too. Then we repeat the procedure, rolling out the sample window one day. VaR estimates are plotted against the corresponding actual monthly returns of MSCI EMF, the red dots. The blue line corresponds to VaR out-of-sample estimates with the bootstrapping method, while the yellow line corresponds to basic historical simulation-VaR. Even if both estimators fail to pass the test of unconditional accuracy, it is apparent that the VaR estimator calculated by bootstrapping is more responsive than the simple historical simulation VaR estimator. It reacts to different market risk conditions and shows more signaling power.