

Dario Brandolini 0039 02 7216 4058, dario.brandolini@rasnet.it
Massimiliano Pallotta 0039 02 7216 2065, massimiliano.pallotta@rasnet.it
Raffaele Zenti 0039 02 7216 4097, raffaele.zenti@rasnet.it

Mailing address:
c/o RAS ASSET MANAGEMENT SGR S.p.A.
P.zza Velasca 7/9, 20122 Milano, Italy

Risk Management in an Asset Management Company: a practical case

MILAN, OCTOBER 2000

PROCEEDINGS OF THE EIGHTH INTERNATIONAL CONFERENCE 'FORECASTING FINANCIAL
MARKETS', LONDON, 30 MAY -1 JUNE 2001

Summary

This article considers the differences between the meaning of risk management in a bank vs. that of an asset management company, illustrating the solution, from both a methodological and technological point of view. This, in the context of the challenges found in building up a risk management system in an Italian medium size company.

1. What is “risk” for an Asset Management Company

The expression “risk management” covers many different topics and, even if it is related to market risk only, it is a source of misunderstanding when referred to an institutional investor (fund manager) as opposed to a financial institution (bank).

In our view there are at least four major differences, as outlined below:

1. The institutional investor manages 'third party funds': eventual losses are not its liabilities, but liabilities of the people or the companies that entrusted the funds to it. At the end, the asset management company can (and will) lose the mandate and the clients, not its capital.
2. Therefore, in view of point (1), losses have no immediate impact on the balance sheet of an asset management company: an open mutual fund can be unbundled (a very rare case), a pension fund can destroy the savings of a life (an even rarer event), with no consequences for the management company, if there are no frauds or mismanagement of the funds. Paradoxically market risk does not constitute a major concern for an asset management company (but the clients!) whilst on the other hand operational risk could be very dangerous (think at legal actions for bad internal auditing). In essence, it is evident that, due to poor performances, in the longer term an asset management company can lose in terms of market share and fees.
3. Investment horizon is much longer for an institutional investor than for a financial institution dealing room. Foundations and pension funds need to have a strategic view of risk that must consider several years. Even if we consider only a tactical/operative view of risk, time horizon is relatively long¹, when compared to the typical time horizon of sellside portfolios. This is due to the fact that buy-side portfolios are often rather static. Most of the time, there is an implicit (mutual funds) or explicit (pension funds) agreement between clients and asset management companies on the time span of this investment horizon.
3. Finally is becoming more and more fashionable to consider relative risk (vs. benchmark, e.g. Relative VaR or tracking error) rather than absolute risk (e.g. VaR or volatility) in measuring the market risk exposure of a managed portfolio.

2. A new challenge for risk management

If we agree on previous considerations, the concept of risk management in an asset management company must consider a different set of problems.

First of all, risk has to be considered the other side of reward: the challenge is taking the minimum amount of risk for the maximum return, not to minimize risk in absolute terms. This means that indicators that take into account the expected (or past) return on portfolio net asset value are very important.

Secondly, there are no publicly defined regulatory guidelines (Basle Committee like) for risk measurement and control in an asset management company.

Thirdly, models for estimating future portfolio risk on a medium term-long term horizon are crucial, and the related theoretical problems are not trivial.

This highlights a key question: why should an investment firm implement a risk management system? Without any specific regulatory guidance on risk, many

¹ The chosen tactical time horizon could be associated to periodic formal assessments of portfolio positions, like investment committees and asset allocation meetings, which commonly take place every month.

asset management companies have traditionally spent the greater part of their analytical and technology budgets on expertise and tools to make money rather than measuring and managing risk on an ongoing formal basis. Investment firms have often managed risk in an intuitive manner and risk management systems have been viewed as an expensive investment, which has to constantly demonstrate its utility. On the contrary, we believe that the risk management function, if well implemented, is not only useful to prevent large negative returns, but can enhance the entire investment process. In particular, risk management can enter in the very initial stage of an investment process, helping to evaluate the amount of risk implied in each investment decision. This can help the portfolio manager to estimate the expected return consistent with this amount of risk. Moreover, performance attribution and risk attribution are gradually seen as two sides of the same coin: advanced products of performance attribution can help to understand if realized returns have been in line with expectations and have paid for the risk exposure. Considering the practicality, risk analysis on a strategic basis (i.e. with a time horizon of several months or even years) should answer questions like: is the risk coherent with the portfolio nature? is the risk compatible with predefined goals (e.g. a given target return)? On a tactical basis (i.e. with a time horizon of several weeks, quite often one month) the typical questions are: are we going to erode our risk budget? are we going to erode our relative performance? For example, given realized excess return of 2% since January, and given that we want to beat the benchmark by the end of the year, what is the probability of underperformance with our current portfolio composition?

Finally, a good reason to put at work risk management systems comes from the market itself. In the community of asset management companies, at least in Italy, the quest for reliable risk management techniques has grown recently, mainly in response to the demand of pension funds and banking foundations. More and more active portfolios with controlled risk, vis-à-vis a benchmark, are requested by institutional investors, and a complete analysis of portfolios risk structure is considered a must in the competition among fund managers. In addition, often, institutional investors put limits on the maximum underperformance allowed: risk management can be of help in avoiding such underperformances.

Lastly, as currently witnessed in various markets, direct selling of financial products through Internet, could lead to a wide use of “easy to apply” risk measures and applications, developed specifically for the retail segment.

In summary whilst market forces lead investment firms to implement sound risk management functions and systems, from the theoretical perspectives, we have seen so far little attention to this field of risk management and to its peculiarities.

3. Forecasting returns is crucial

It is clear that risk management cannot avoid the problem of forecasting future returns for each portfolio under management, because this is the crucial point in calculating any risk indicator and evaluating every asset decision. As we mentioned previously, dealing with risk management in an asset management company, medium-long term horizons are the rule. This has some serious implications: i) over medium-long horizons it is not safe to assume that portfolio returns are normal with mean equal to zero; ii) volatilities and correlations change over time². Figure 1 and Table 1, for instance, show some evidence that MSCI

² Of course, conditional normality is compatible with unconditional non-normality: this is typical of GARCH processes with normal innovations. Anyway, this is not our case, because examining our

World monthly log-returns (in Euro terms) over the period March 1971 – May 2000 exhibit skewness and fat tails, so they are not normal. Moreover, they have a mean that is statistically different from zero. Results are basically the same if we consider local returns instead of Euro-denominated returns. We have examined many financial time series of weekly, monthly and quarterly returns, most of the time rejecting both the (null) hypothesis of normality and the zero-mean hypothesis. In Figure 2 we show an example of time varying correlation between MSCI Emerging Markets and MSCI Japan.

Figure 1

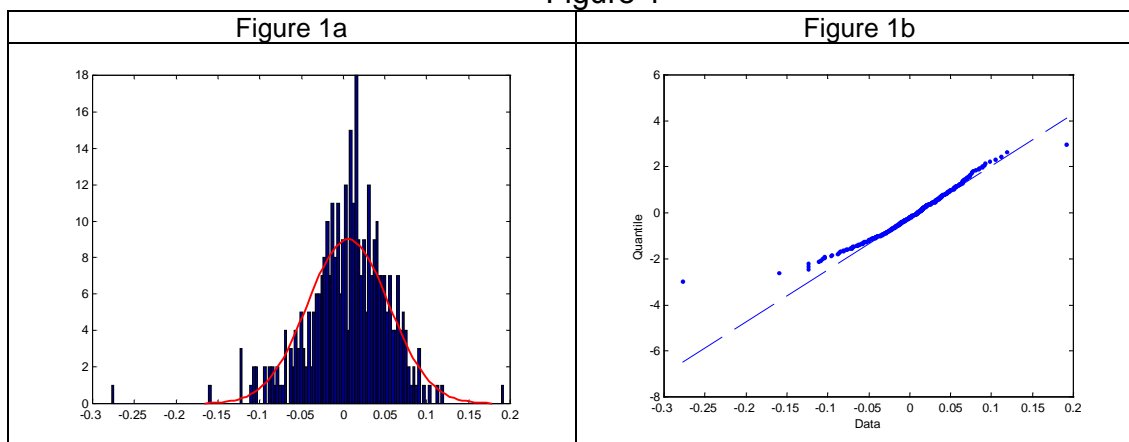


Figure 1a shows the histogram, with superimposed fitted normal density of monthly log-returns of MSCI World in Euro terms (source Datastream International), over the period March 1971 – May 2000. The histogram shows that the pdf is rather asymmetric, with presence of tail events.

Figure 1b shows the Quantile-Quantile plot of the same monthly log-returns, on the horizontal axis, versus the standard normal distribution, on the vertical axis. If the returns were normal, the plot should look roughly linear. Especially the left side of the 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 show that the empirical distribution has a mean that differs significantly from 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. Using the procedure presented by Koedijk, Huisman and Pownall (1998), we estimated also the left tail index, which equals the number of existing moments. Not unexpectedly, the tail index is approximately equal to 4. This means that the first four moments exist, but that higher moments are infinite.

Figure 2

typical portfolio returns we have found that quite often GARCH residuals are strongly non-gaussian.




Figure 2 shows the correlation between MSCI Emerging Markets and MSCI Japan, both in Euro terms (source Datastream International), calculated with a GARCH model over the period January 1990 – May 2000 using daily data. It is apparent that the estimated correlation is time varying.

In summary, the standard hypothesis that asset returns follow a gaussian distribution can be unreliable because empirical distributions may not conform to any known distribution. In a similar way, an analysis based on constant linear correlations and volatilities can be strongly misleading, because they can vary a lot from period to period. Having said this, calculating risk measures for asset managers involves finding a proper way to model future scenarios over medium-long term horizons.

4. Role of risk management in RAS ASSET MANAGEMENT (“RAM”).

Leaving behind the entire classical risk management hypothesis, we have to move into a new territory in which risk management has to demonstrate its utility. In fact, in an asset management company there is a minor commitment to measure and control risk than in a bank submitted to Basle constraints.

Amongst other things, top management expects from the risk management entity help for the fund managers to enhance return without exceeding volatility. It is generally expected that performance attribution and risk management have to go side by side, if this task has to be achieved.

Some of these considerations were clear to us, and some became clear later, when the risk management unit of Gestiras, the mutual fund company of RAS Group, was started in 1997. In a few months the process became more complex with the transformation, in mid 1998, of Gestiras into Ras Asset Management (RAM) and the broadening of the portfolios under management to the life insurance products and the pension funds/institutional investors funds.

Today asset under management are now close to 23 billions euros, distributed in more than 120 portfolios, four of them Luxembourgian, with every kind of financial assets, from money market instruments to bonds, stocks and derivatives, denominated in both euros and foreign exchange.

Although the dimension of RAM can be defined as medium size in the Italian market, it has to deal with a very large set of problems related to the complexity of the structure under management.

As well known, risk management is a data-driven process, so data integration is a critical success factor. In RAM, data storage had been a major obstacle for any kind of development, because basic data were located in several databases, logically and physically away from the front office system.

In fact we have to collect data not only from the back-offices of RAS Group but also from depositary banks in Luxembourg and Italy (respectively for the Luxembourgian mutual funds and the pension funds/institutional clients). This further requires integration of back office with front office data, and other information from several market data-providers.

The “typical” risk tools in RAM were “fund manager oriented” softwares (i.e. Barra AegisTM and CosmosTM) used both by the managers for testing different portfolios, and by the Risk Unit for controlling, on a regular basis, the risk structure of mutual funds portfolios.

The biggest effort was in fact devoted to feed Barra products, through a centralized system, with the different databases where data were dispersed.

The major limitations of this approach are:

- databases dispersion creates problems in testing quickly and easily new portfolios;
- “black box” parametric risk model (like Barra AegisTM) are difficult to be evaluated in their details;
- tests on risk model prediction capacity are very difficult to be implemented, because of the rather closed nature of these tools;
- the risk model performs at best when expected excess returns of single assets are given by the money managers, and this is not easy to be obtained.

Therefore, it was necessary to start a risk management system project having in mind our limits:

- limited time and money to implement a risk system;
- necessity to diffuse a broad risk culture in every part of the company.

5. The RAM approach

First we started with a selection of available risk systems analyzing their cost/benefit trade-off. In fact, it was immediately clear to us that any risk management project in RAM should have three major key success factors, listed below:

- 1) Model: it can be misleading to calculate risk measures without an appropriate model for generating future financial scenarios.
- 2) Technology: the system should process, aggregate and manipulate data in reasonable time; this should be done easily, from the user’s point of view, and in a good-looking layout.
- 3) Data & Data model: we need a robust data model where all the relevant asset data and characteristics are stored and well organized, in order to enable to drill down and to aggregate along arbitrary user-defined data dimensions when performing risk analysis.

With these requisites in mind, we began screening market software in order to find a ready-to-use solution. We examined both simple spreadsheet oriented tools and complex, firm-wide risk management systems. After a rather complete search we concluded that:

- like the risk models, risk systems were mainly bank-oriented, with little attention to the asset management peculiarities;
- there were no solutions available at a price consistent with the needs of RAM;

- there were no easy-to-implement solutions; they all needed a (big or small) customization in order to catch the particular risk approach of an asset management company;
- the database was crucial for every successful risk system implementation.

Therefore, we decided i) to create an integrated database; ii) to develop an in-house system in order to learn about risk models, testing the best models while creating the basis of a risk culture all over the company; iii) to integrate risk management and asset allocation and iv) to consider, in the future, the option of purchasing a heavy multi-tasking risk software.

6. R.A.M. solution: the model

As indicated previously, calculating risk measures for asset managers involves finding a proper way to model future scenarios when time horizon is relatively long and returns are not necessarily zero-mean normal, with constant variance-covariance matrix. We choose to use an “enhanced” historical simulation approach, largely based on the model proposed by Barone-Adesi, Giannopoulos and Vosper (1999). “Enhanced” means we apply bootstrapping techniques to historical simulated data, without strong, “ad hoc” hypothesis on the distribution of returns. In our opinion this method provides a robust, flexible, intuitive and easy-to-explain framework for buy-side risk analysis.

In order to understand the simple idea behind the model, assume we have daily returns, and the investment horizon is 1 month. For simplicity sake we assume that returns are identically and independently distributed (i.i.d.). This is an unnecessary hypothesis that can be removed, as discussed later. In order to calculate any risk indicator over the selected time horizon we require the probability distribution of monthly returns. Therefore, beginning from daily data, a technique must be used to construct the distribution by generating whole month-long paths of asset returns. To do so one can randomly pick-up a strip of 21 daily returns. Compounded, they produce the first one-month return. Repeating this procedure, say, 5,000 times we are able to estimate the probability distribution of monthly returns, using the information content of daily data.

Let us consider a portfolio together with its benchmark that can invest in a universe of K assets. If Δw is the column vector of length K containing all the current differential weights, and r is the matrix $T \times K$ containing the time series of length T for all the asset returns, then

$$er = r \cdot \Delta w \quad (1)$$

is a column vector of length T that contains the constant-mix time series of excess returns, based on the current weights. Therefore, if the goal is to simulate M scenarios (e.g. 5000) in terms of excess returns over a given 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 :

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

Thus, for each scenario, the excess return ER over H periods is given by:

$$ER = \sum_{i=1}^H er(u_i) \quad (3)$$

where $er(u_i)$ is the element of vector er_{PTF} , which is in position u_i . Note that vector er_{PTF} contains “raw” data (e.g. daily excess returns). Repeating this simulation M times allows us to generate a distribution of possible future scenarios based on historical data. Note how we keep together all the returns that refer to a certain date. By doing this we ensure that we capture any kind of (not necessarily linear) time varying correlation existing among assets and any non-normality in asset price changes. In addition, if risk factors coincide one to one with the securities held in the portfolio, one can decompose risk in many ways (e.g. by country, sector, duration, rating and so on). Note, however, that if “raw” returns are not i.i.d., they are unsuitable for historical simulation and can lead to biased results. For instance, we ignore the eventual presence of autocorrelation and volatility clusters. To avoid this and other problems, it is possible to modify the basic scheme as outlined below.

Behind the basic version of enhanced historical simulation with bootstrapping, there is the economic hypothesis that future returns follow the same distribution of past actual returns. This can be heroic. The empirical distribution of returns can exhibit a 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 negative return being incurred is not the same on each period: during periods of higher volatility we will expect larger than usual negative returns. Problems can arise also with regards to autocorrelation; a rather common finding is that financial returns are correlated with their own lagged values. From the practitioners’ perspective, a high serial correlation means, for instance, that large negative changes in returns are likely to be followed by further large negative changes.

Dealing with medium-long term scenarios, one may also want to modify the unconditional mean and unconditional volatility of portfolio returns, so that future scenarios reflect any specific market view. This can be crucial for pension funds and foundations.

There are several variations of the basic historical simulation-bootstrapping techniques to get proper modeling of asset returns that can cope with all this problems. Using the BAGV approach, we filter historical simulated portfolio excess returns (or total returns) using an ARMA-GARCH model to get i.i.d. residuals. 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 scaled by the ratio of current over past conditional GARCH volatility in order to get i.i.d. observations. Then one can use (2) 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 Appendix contains a short analytical description of the approach we use, but see Pallotta and Zenti (2000) for a detailed description of the basic methodology and some variations of its scheme.

Bootstrapping procedures can cope with options in a very natural way, through “full option revaluation”. If the investment horizon is 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 completely misleading,

obfuscating the real risk of the options. Delta-gamma methods are unstable for large asset price changes, which can easily occur if the investment horizon is long. The historical simulation/bootstrapping model is also well suited if the portfolio contains path-dependent contracts; in this case the whole path taken has to be modeled.

It is clear that this method is similar to a Monte Carlo simulation. However, traditional approaches based on Monte Carlo simulation typically use a set of stochastic differential equations for generating returns over the time horizon. So traditional Monte-Carlo simulation uses arbitrary assumptions about the distribution of returns, which define “a priori” the structure of risk that is supposed to investigate. Instead, we use real-world asset returns movements, reflecting actual time-varying correlations.

From our point of view, our approach has some advantages also on multifactor models. Multifactor model based on factors like, for instance, value, growth, size and momentum, or pure statistical principal component are in our opinion less intuitive. We see difficulties in explaining, say, that some international equity portfolio exhibits growing risk due to high exposure to, say, “value”. Maybe it is more direct to say that high-tech companies give the highest contribute to portfolio risk. There are also considerable challenges in estimating and explaining to non-quantitative users the expected return of a given factor. Moreover, multifactor methods are based on a strongly parametric approach, typically based on the normality assumption.

Figure 3a and Figure 3b show in an intuitive way the typical out-of-sample behaviour of VaR estimates using the historical simulation-bootstrapping approach, compared with standard parametric VaR estimates.

Figure 3a




Figure 3b




Figure 3a 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. We use a rolling window of 18 months of daily data to estimate VaR for the next month, then each day we calculate the actual past monthly return. 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 3b 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 VaR estimated with the historical simulation-bootstrapping method is better than the other VaR estimator. Apart from formal tests of conditional and unconditional accuracy³, the former is more reactive to different risk conditions and hence has more signaling power. In our opinion, a good risk measure should help warn portfolio managers of large price moves.

Using the historical simulation-bootstrap approach we calculate not one but many risk indicators, both relative and absolute, e.g. Relative VaR, tracking error, absolute volatility, shortfall probability, worst case, and so on. This allows us to get a broad view of portfolio market risks. One can debate if it is more appropriate to use tracking error, a more common measure of relative risk, or Relative VaR (ReVaR). Tracking error is usually defined as 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 (as it is often the case), the use of tracking error leads to an underestimation of relative risk. The use of ReVaR calculated using a bootstrap-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

³ See Pallotta, Zenti (2000) for a detailed comparison of several VaR estimators.

negative events. For an asset manager ReVaR is a more robust way to quantify relative market risk. In addition, computing ReVaR using an historical simulation-bootstrapping 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.

Another key question is: VaR or ReVaR? Our idea is that pension funds, mutual funds and other buy-side portfolios could use absolute risk indicators like VaR 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.

Our approach to scenario generation and risk measurement has the advantage that can be applied not only to any kind of mutual or pension funds, but also to funds of funds, an increasingly important market segment.

7. R.A.M. solution: data & technical architecture

Having to deal with a complex data feed structure (legacy systems, data providers, custodian banks) and with different priorities given to our risk management project by each IT department, as a first step we decided to create our own data warehouse.

What we need was a system capable of managing medium volume of data, merging, updating and processing positions for many portfolios and benchmarks. Our solution was Hicare's Lilith® database, an innovative Italian product, that has been engineered to overcome the inefficiency and ineffectiveness showed by standard relational databases when dealing with multi-dimensional and hierarchical data structures. Lilith® stores data in compact hyper-cubic structures in order to minimize disk space occupation (from the technical point the database is at the same time Hierarchical, Cartesian and Relational). Despite being technically advanced, Lilith® is strongly end-user-oriented and, in a sense, can be seen as a natural evolution of spreadsheets. Lilith® receives portfolio positions directly from the back-office using flat text files.⁴

Securities' prices, indices and other market data come from several providers, through Lilith®'s macros that dynamically write requests to the info-provider. Mathworks' Matlab® is the calculation engine. We used Matlab® as a high level programming language, taking advantage of many existing functions to develop routines that implement our proprietary models. We implemented our Matlab® library rather quickly (more time was spent on model testing and database creation). Matlab®'s ability to manipulate big matrices, avoiding many heavy "For" or "While" cycles, facilitates reasonably fast computation (for our needs!), even if one has to deal with hundreds of portfolios. Lilith® and Matlab® communicate through a specific gateway developed by Hicare, so Matlab® is virtually a Lilith® add-in.

Lilith® is not just a database, it is also a data-manager. Hence Lilith® is used for data visualizations and calculations results, given its extensive reporting

⁴ In order to simplify the system integration process and to minimize the intervention of external IT departments we adopted a file transfer system. As well known, the main disadvantage associated with this approach is that, in principle, data can be easily corrupted hence data integrity becomes a crucial issue. Anyway, having a unique front office system, this was not a real problem for us.

capability, using tables, many kind of graphs, in both PC screen and printed version. The final user can navigate through all the reports on PC, examining different portfolios, changing dates, and looking at different risk measures. The program also facilitates web based reporting.

8. Implementation details

We choose to model risk at the single security level, so risk factors coincide one to one with securities.⁵ This means we store in our database the time series of all the securities contained in any portfolio or benchmark. This allows us: i) to decompose risk in many ways; ii) for foreign securities, to use the same exchange rates for each portfolio and its benchmark, avoiding one of the biggest sources of errors in risk measurement, especially on the fixed income side, where FX risk is more relevant.

Presently we cover 18 months of price history, with plans to extend the database's depth. If a price time series is incomplete we automatically fill the gap with the more appropriate index. If, say, a recently issued Italian corporate bond rated A, which matures in 5 years, has only 3 months of data, for the remaining 15 months we use an index of the Euro-block, 3-5 years A-rated corporate bonds. If such an index does not exist, we automatically search for the best proxy.

We use the following risk indicators (however, the open structure of the system allows us to implement any new risk indicator as required):

- on the relative risk side:
 - ReVaR, with investment horizon 1 month and confidence level 95%;
 - Worst case in 1 month;
 - Tracking error (monthly and annualized);
 - Shortfall probability in one month, setting the goal excess return equal to 0%;
- on the absolute risk side:
 - VaR, with investment horizon 1 month and confidence level 95%;
 - Worst case in 1 month;
 - Volatility (monthly and annualized);
 - Shortfall probability in one month, setting different goal returns, depending on the specific nature of the portfolio.

Note that we calculate tracking error as the difference between the median and the 15.87% quantile of estimated monthly excess returns. If they are normal, this measure coincides with the standard deviation of excess return (i.e. the traditional notion of tracking error), otherwise it reflects the specific empirical shape of the left tail. It is a kind of non-parametric tracking error.

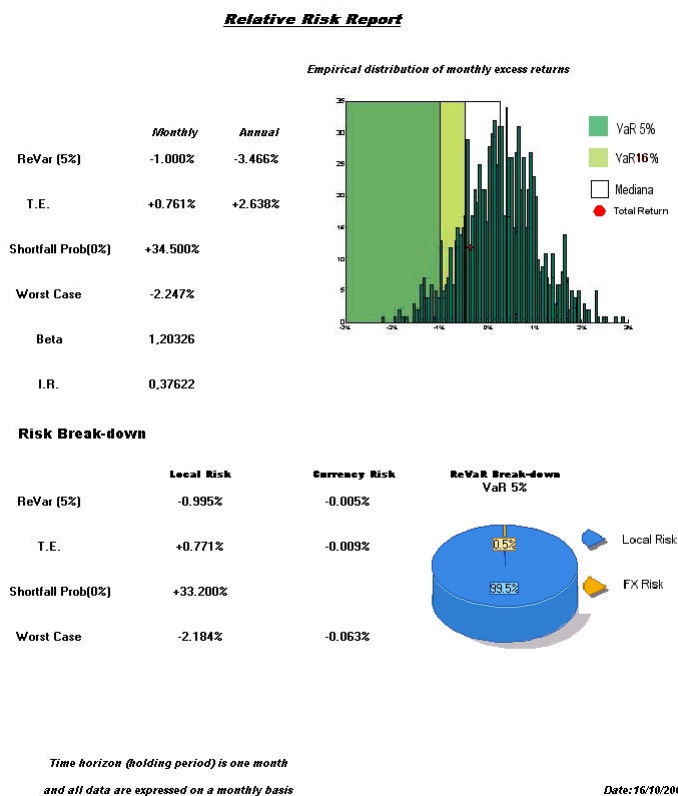
The chosen time horizon is (currently) one month, so our analysis is mainly from a tactical point of view.

Given the rather low portfolio turnover, standard risk assessment is done twice a month on all our portfolios. All the calculations are performed in 6-8 hours, on our current server (with 2 x 500 MHz CPUs and 512 MB of RAM), so in principle, risk

⁵ Actually we manage more than 6000 securities plus hundreds of indices.

assessment could be done more frequently. At any time we can analyze single portfolios “on request”, for instance if suggested by market conditions. Figure 4 shows a sample report on the relative risk (at an aggregate level) of an active mutual fund.

Figure 4



Conclusions

This paper has discussed the development and implementation of a comprehensive risk management system within a medium sized Italian asset management company. The functionality of the system is tailored to the specific needs of an asset management company in term of (1) proper financial modeling and (2) time & resources constraints. Our simulation methodology based on bootstrapping, allows for reliable calculations of a variety of risk measures for a large number of portfolios. The technique considers current market conditions and is based on the empirical multivariate distribution of data; we do not impose any particular probability function.

Our model allows us to generate medium-term to long-term financial scenarios, whilst taking into consideration (if these views exist) heteroskedasticity, serial correlation (if relevant) and specific market views. From the IT perspective, the system is based on a flexible database, Lilith®, and a standard calculation engine, Matlab®, both well-known market available applications. Compared to traditional bank-oriented risk systems, our system is relatively straightforward and cheaper to implement. Total set up and implementation phase has taken about one year time for both the database structure and the risk management tool, total external costs (including the licences for Matlab® and Lilith®) were around 250 thousand euros.

Acknowledgements

We gratefully acknowledge helpful comments, suggestions and full support from Andrea Beltratti. Thanks to Anna Vigliotti for reviewing the paper. The authors are responsible for any errors and inaccuracies.

References

1. Barone-Adesi, G., Giannopoulos, K., Vosper, L. (1999). VaR without correlations for non-linear portfolios. *Journal of Futures Markets*, 19, August.
2. Efron, B., Tibshirani, R. (1998). *An introduction to the bootstrap*, 2nd edn. Chapman & Hall.
3. Koedijk, K. Huisman, R., Pownall, R. (1998). Var-x: fat tails in financial risk management. *Journal of Risk*, Volume 1, n° 1.
4. McNeil, A.J., Frey, R. (1999). Estimation of tail-related risk measures for heteroskedastic financial time series: an extreme value approach. ETHZ, working paper.
5. Pallotta, M., Zenti, R. (2000). Risk analysis for asset managers: historical simulation, the bootstrap approach and value at risk calculation. *Ras Asset Management*, working paper (submitted to Social Science Research Network, available at http://papers.ssrn.com/paper.taf?abstract_id=251669).

Appendix: how we simulate future returns

Our simulation is based on the constant-mix time series $\{r\}$ of returns, or excess returns, we get using the current portfolio composition. The problem is that usually this time series is unsuitable for a naïve application of bootstrapping. If we apply directly the basic bootstrap procedure to $\{r\}$ neither we model conditional volatility (so we “miss” volatility clusters) nor we capture any autocorrelation showed by data. From another perspective, the basic bootstrap procedure is based on the assumption that “raw” excess returns are identically, independently distributed (i.i.d.). Therefore, if “raw” returns are not i.i.d., and note that often $\{r\}$ is unlikely to be i.i.d., they are unsuitable for bootstrapping and can lead to biased results.

To avoid this problem, it is possible to modify the basic scheme. We filter $\{r\}$ using the ARMA(1,1)-GARCH(1,1) model proposed by Barone-Adesi et al. (1999). This means that $r(t)$, the excess return (or, alternatively, total return) at time t obeys to the following equations:

$$[A1] \quad r(t) = a + b \cdot r(t-1) + c \cdot e(t-1) + e(t) \quad \text{ARMA}(1,1) \text{ equation}$$

where $e(t)$ is the residual at time t , which follows its empirical distribution, with zero mean and variance $s(t)$, that follows a GARCH process:

$$[A2] \quad s(t) = \alpha + \beta \cdot s(t-1) + \gamma \cdot e(t-1)^2 \quad \text{GARCH}(1,1) \text{ equation}$$

We estimate the parameters a , b , c , α , β , γ using Pseudo Maximum Likelihood procedures; our estimation approach is rather similar to the approach used by McNeil and Frey (1999). Note that equation [A1] models the mean, and takes into account first order autocorrelation, while equation [A2] models the volatility of returns.

Residuals of the ARMA model are divided by the past conditional GARCH volatility in order to get i.i.d. observations. Then we bootstrap these observations over the desired time horizon. Finally, the paths of i.i.d. observations we get are used as innovations to simulate the ARMA-GARCH process. It is important to notice that: i) bootstrapping is applied to i.i.d. residuals, so results are unbiased; ii) residuals are standardized but not normal.

In general, to model $r(t)$ it is possible to use specifications where the mean equation is any desired time series model, while the variance equation is a GARCH process of any kind and order, e.g. a TARCH(2,2).