ALEA
Tech Reports

Mixture models for VaR and stress testing

Marco Bee

Tech Report Nr. 12
Giugno 2001

Alea - Centro di ricerca sui rischi finanziari
Dipartimento di informatica e studi aziendali
Università di Trento - Via Inama 5 - 38100 - Trento
http://www.aleaweb.org/

ALEA, Centro di ricerca sui rischi finanziari è un centro di ricerca indipendente costituito presso il Dipartimento di informatica e studi aziendali dell'Università di Trento. Il centro si propone di elaborare conoscenze innovative in materia di risk management, *e di favorirne la diffusione mediante pubblicazioni e programmi di formazione. Il centro è diretto dal prof. Luca Erzegovesi. Nella collana ALEA Tech Reports sono raccolti lavori di compendio e rassegna della letteratura e della prassi operativa su tematiche fondamentali di* risk management.

MIXTURE MODELS
FOR VAR AND STRESS TESTING

Marco Bee

Banca Intesa

Risk Modelling - Risk Management

first draft - March 2001 (*)

ABSTRACT

In this paper we deal with the use of multivariate normal mixture distributions to model asset returns. In particular, by modelling daily asset returns as a mixture of a low-volatility and a high-volatility distribution, we obtain three main results : (i) we can use posterior probabilities to identify hectic observations; (ii) we are able to compute a nonparametric fat-tails Value at Risk by sampling repeatedly from the mixture and computing the quantile of the empirical distribution; (iii) we can use the estimated parameters of the hectic distribution for stress testing purposes. We show how these three items can be addressed using either real data and simulation methods.

(*) I wish to thank Luca Cazzulani, Vladimiro Ceci and Walter Vecchiato for their comments and suggestions. All remaining mistakes are mine own.

1 Introduction

Since the work by Mandelbrot (1963) and Fama (1965), many empirical analyses have found evidence of heavy tails in the marginal distributions of asset returns. The implication is that the usual normal-based Value at Risk (VaR) calculations are likely to be misleading, in the sense that they would produce figures that underestimate actual losses.

On the other hand, stress testing analyses are often based on the volatility and correlation of the "hectic" observations, i.e. the observations corresponding to periods of financial markets' crisis, but it is not clear how to identify these periods. Whereas this last problem is still essentially open, many solutions to the fat-tails problem have been proposed in the literature, including ARCH-type models, Extreme Value Theory and Stochastic Volatility models.

In this paper we focus on the use of mixture distributions, which provide a good way of tackling both problems. The basic idea is that the data generating process of a time series of asset returns is a mixture of two distributions with a similar mean but different covariance matrices: in particular it is natural to expect that most observations are generated by the low-volatility distribution, while the remaining observations are generated by the high-volatility (hectic) distribution. As will be seen in the following sections, this approach can be considered as a tool for reaching three different goals:

(i) identifying correctly the observations generated by each distribution;

(ii) estimating a fat-tails VaR;

(iii) perfoming stress test analyses using the parameters of the high-volatility distribution.

The rest of the paper is organized as follows: in section 2 we introduce the statistical model and its most important properties; in section 3 we explain how it can be used in a risk management framework; in section 4 we apply the methodology to real and simulated data and present the results; section 5 concludes.

2 A Model for Asset Returns

In this section we define the statistical model which will be used to describe the behavior of asset returns and derive some of its properties. In the next section we will show how these properties can be used for our purposes.

Given a vector of p asset returns $\mathbf{Y}' = (\,Y_1 \quad \cdots \quad Y_p\,)$, it is convenient to partition it into a set of core assets $\mathbf{Y}_1' = (\,Y_1 \quad \cdots \quad Y_q\,)$ and a set of peripheral assets $\mathbf{Y}_2' = (\, Y_{q+1} \quad \cdots \quad Y_p \,)$. Let the density of the $(p \times 1)$ random vector $\mathbf{Y}' = (\, \mathbf{Y}_1' \quad \mathbf{Y}_2' \,)$ be

$$f_{\mathbf{Y}}(\mathbf{y}) = \pi_1 f_{\mathbf{Y}^{(1)}}(\mathbf{y}) + \pi_2 f_{\mathbf{Y}^{(2)}}(\mathbf{y}),$$

where $\pi_2 = 1 - \pi_1$ and $f_{\mathbf{Y}^{(i)}}(\mathbf{y})$ is the $N(\boldsymbol{\mu}^{(i)}, \boldsymbol{\Sigma}^{(i)})$ density. In partitioned form, the expected value and covariance matrix of $\mathbf{Y}$ in the two populations are given by

$$E_i(\mathbf{Y}) = \boldsymbol{\mu}^{(i)} = \begin{pmatrix} \boldsymbol{\mu}_1^{(i)} \\ \boldsymbol{\mu}_2^{(i)} \end{pmatrix}, \qquad i = 1, 2;$$
$$\mathrm{cov}_i(\mathbf{Y}) = \boldsymbol{\Sigma}^{(i)} = \begin{pmatrix} \boldsymbol{\Sigma}_{11}^{(i)} & \boldsymbol{\Sigma}_{12}^{(i)} \\ \boldsymbol{\Sigma}_{21}^{(i)} & \boldsymbol{\Sigma}_{22}^{(i)} \end{pmatrix}, \qquad i = 1, 2,$$

where the notation E_i, $\mathrm{cov}_i(\mathbf{Y})$ means that the expectation and covariance matrix are computed with respect to the i-th distribution. The maximum likelihood estimators of the parameters $\pi_1, \boldsymbol{\mu}^{(1)}, \boldsymbol{\mu}^{(2)}, \boldsymbol{\Sigma}^{(1)}, \boldsymbol{\Sigma}^{(2)}$ can be obtained using the EM algorithm [for details see McLachlan and Krishnan (1996)]. As a by-product, the algorithm provides us with the so called posterior probabilites $\tau_i(\mathbf{y}_j)$:

$$\tau_i(\mathbf{y}_j; \hat{\boldsymbol{\theta}}) = \frac{\pi_i f_{\mathbf{Y}^{(i)}}(\mathbf{y}_j; \hat{\boldsymbol{\theta}})}{\pi_1 f_{\mathbf{Y}^{(1)}}(\mathbf{y}_j; \hat{\boldsymbol{\theta}}) + \pi_2 f_{\mathbf{Y}^{(2)}}(\mathbf{y}_j; \hat{\boldsymbol{\theta}})}, \tag{1}$$

where $\boldsymbol{\theta}' = (\, \pi_1 \quad \boldsymbol{\mu}^{(1)'} \quad \boldsymbol{\mu}^{(2)'} \quad \mathrm{vec}(\boldsymbol{\Sigma}^{(1)})' \quad \mathrm{vec}(\boldsymbol{\Sigma}^{(2)})' \,)$. (1) represents the probability, computed on the basis of the estimated parameters, that an observation comes from the i-th population. It is often called *posterior probability* of the j-th observation.

Finally, it is possible to evaluate the conditional expectation of $\mathbf{Y}_2$ given $\mathbf{Y}_1$ in each population:

$$E_i(\mathbf{Y}_2 | \mathbf{Y}_1 = \mathbf{y}_1) = \boldsymbol{\mu}_{2.1}^{(i)} = \boldsymbol{\mu}_2^{(i)} + \boldsymbol{\Sigma}_{21}^{(i)} (\boldsymbol{\Sigma}_{11}^{(i)})^{-1} (\mathbf{y}_1 - \boldsymbol{\mu}_1^{(i)}), \quad i = 1, 2. \tag{2}$$

We expect this conditional expectation, when computed using the (estimated) parameters of the high-volatility distribution, to reflect the reaction of the peripheral assets according to the volatilities and correlations of the hectic periods.

3 Mixture Models in Risk Management

Mixture models can be used for risk management purposes as outlined in section 1.

Consider first the problem of identifying periods of financial markets turmoil. This issue is often important because statistical properties of the data generating process change considerably in turbulent periods; in particular, correlations between asset returns are different in stressful and quiet market conditions, a phenomenon known by the name *correlation breakdown* [see Boyer, Gibson and Loretan (1999) and Loretan and English (2000)]. According to this finding, a naive methodology of performing stress test analyses consists in estimating correlations using only observations from hectic periods, which poses the problem of identifying the observations to be used. Implicitly, this approach assumes that the observations are generated by a mixture of low- and high-volatility distributions and therefore, if we introduce explicitly this assumption, we can use (1) to conclude that $\mathbf{y}_j$ is an hectic observation if and only if $\tau_2(\mathbf{y}_j;\hat{\boldsymbol{\theta}}) > 0.5$.

Once we have estimated the parameters of the mixture, we can easily get a fat-tails VaR: a mixture distribution is indeed able to approximate arbitrarily well any distribution, and therefore, if the distribution of the data is leptokurtic, the non-parametric VaR obtained by Monte Carlo simulation from this distribution will reflect this feature, that is, will be larger (in absolute value) than the normal one. Notice that this is still a result related to "normal" market conditions because it uses all the observations to address the problem of fat tails. On the other hand, if we want to perform stress test analyses (i.e. we decide to focus on hectic observations), we can obtain a (parametric or nonparametric) VaR using the covariance matrix $\boldsymbol{\Sigma}^{(2)}$ of the high-volatility distribution.

As for stress testing, the idea [see Kim and Finger (2000)] consists in shocking the core assets $\mathbf{Y}^{(1)}$ and computing the movement of the peripheral assets $\mathbf{Y}^{(2)}$ as the conditional expectation $E_2(\mathbf{Y}^{(2)}|\mathbf{Y}^{(1)})$.

4 Applications and results

(i) Identifying hectic observations. In order to assess whether posterior probabilities give correct indications about the population membership of an observation, we performed an experiment based on simulated data: we generated 500 observations from the mixture

$$f(\mathbf{y}) = 0.9 f_1(\mathbf{y}) + 0.1 f_2(\mathbf{y}), \tag{3}$$

where $f_1(\mathbf{y}) \sim N_5(\mathbf{0}, \mathbf{\Sigma}^{(1)})$ and $f_2(\mathbf{y}) \sim N_5(\mathbf{0}, \mathbf{\Sigma}^{(2)})$. $\mathbf{\Sigma}^{(1)}$ and $\mathbf{\Sigma}^{(2)}$ are given by

$$\mathbf{\Sigma}^{(1)} = \begin{pmatrix} 0.536 & 0.170 & -0.30 & -0.845 & 0.211 \\ 0.169 & 1.558 & 0.508 & -0.285 & -0.796 \\ -0.296 & 0.508 & 1.208 & -0.136 & -0.008 \\ -0.845 & -0.285 & -0.136 & 2.954 & -0.774 \\ 0.211 & -0.796 & -0.008 & -0.774 & 1.899 \end{pmatrix};$$

$$\mathbf{\Sigma}^{(2)} = \begin{pmatrix} 9 & 0.170 & -0.30 & -0.845 & 0.211 \\ 0.169 & 8 & 0.508 & -0.285 & -0.796 \\ -0.296 & 0.508 & 6 & -0.136 & -0.008 \\ -0.845 & -0.285 & -0.136 & 7 & -0.774 \\ 0.211 & -0.796 & -0.008 & -0.774 & 10 \end{pmatrix}.$$

Notice that $\mathbf{\Sigma}^{(2)}$ differs from $\mathbf{\Sigma}^{(1)}$ only for the variances, which are much larger in the second population.

Unlike the multivariate normal case, in the multivariate normal mixture setup the estimates of the parameters obtained with marginal and full dimensional data are not the same [for details see Bee (1998)]. For this reason we show below the estimates obtained with five-, two- and one-dimensional data. With five-dimensional data we get:

$$\begin{aligned} \hat{\pi}_1 &= 0.916; \\ \hat{\boldsymbol{\mu}}^{(1)'} &= (\,-0.044 \quad -0.108 \quad -0.048 \quad 0.110 \quad 0.071\,); \\ \hat{\boldsymbol{\mu}}^{(2)'} &= (\,-0.377 \quad 0.307 \quad -0.389 \quad 0.566 \quad -0.751\,); \\ \mathrm{diag}(\hat{\mathbf{\Sigma}}^{(1)})' &= (\,0.539 \quad 1.647 \quad 1.182 \quad 2.654 \quad 2.170\,); \\ \mathrm{diag}(\hat{\mathbf{\Sigma}}^{(2)})' &= (\,12.216 \quad 8.482 \quad 7.126 \quad 6.453 \quad 11.071\,). \end{aligned}$$

With bivariate data (the data of the first two components of $\mathbf{Y}$) the results are:

$$\hat{\pi}_1 = 0.920; \quad \hat{\boldsymbol{\mu}}^{(1)'} = (\,-0.026 \quad -0.111\,); \quad \hat{\boldsymbol{\mu}}^{(2)'} = (\,-0.595 \quad 0.360\,);$$
$$\mathrm{diag}(\hat{\mathbf{\Sigma}}^{(1)})' = (\,0.552 \quad 1.631\,); \quad \mathrm{diag}(\hat{\mathbf{\Sigma}}^{(2)})' = (\,12.435 \quad 8.943\,).$$

Finally, with univariate data (the data of the first component of $\mathbf{Y}$) we get

$$\hat{\pi}_1 = 0.866; \quad \hat{\mu}^{(1)} = -0.001; \quad \hat{\mu}^{(2)} = 0.392;$$
$$\hat{\sigma}_1^2 = 0.550; \hat{\sigma}_2^2 = 5.014.$$

Figures 1 to 3 show the posterior probabilities that an observation comes from the high-volatility distribution, computed respectively with five-variate, bivariate and univariate data. Notice that, although at first sight the graph in figure 1 is more convincing, the graphs in figures 2 and 3 are more informative, as a few small (in absolute value) returns are actually generated by the high-volatility distribution. For comparison, figure 4 shows the results of the same Monte Carlo experiment, but with observations generated by a mixture whose second component has a mean vector equal to $\boldsymbol{\mu}^{(2)'} = (\,6 \quad 6 \quad 6 \quad 6 \quad 6\,)$. In this case, where population membership is essentially known, estimating the posterior probabilities with full data gives extremely precise results. Anyway, in all cases the posterior probability indicates that we are essentially certain about population membership of the large (in absolute value) observations, and this is the most important result for our purpose.

Fig. 1 - Posterior probability as function of returns - univariate data

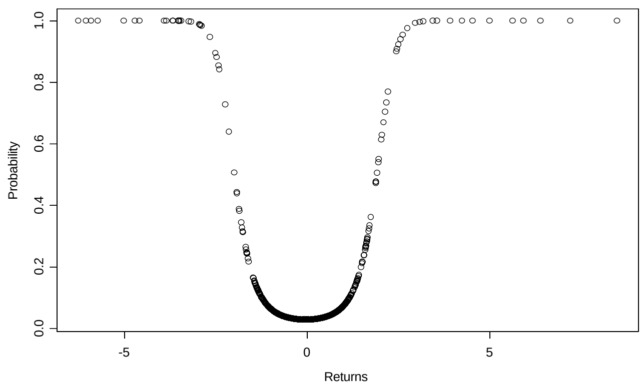

Fig. 2 - Posterior probability as function of returns - bivariate data

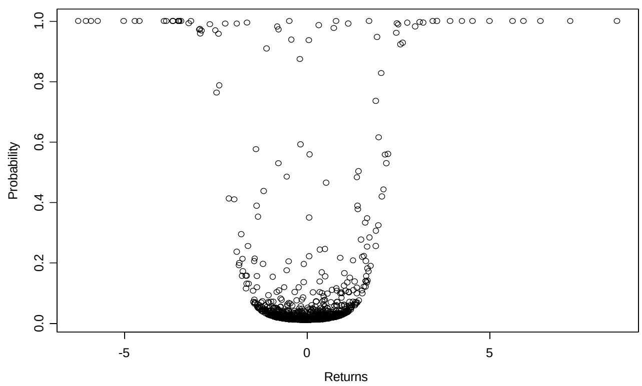

Fig. 3 - Posterior probability as function of returns - multivariate data

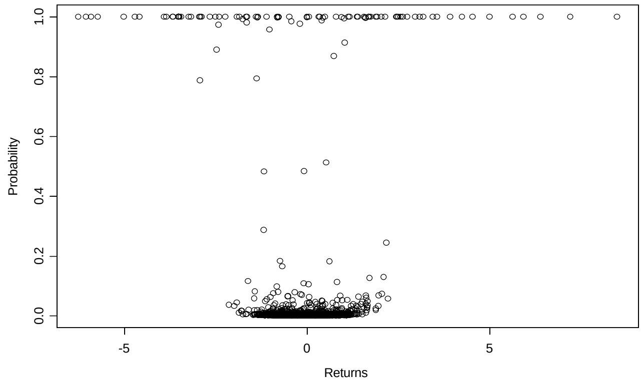

Fig. 4 - Posterior probability as function of returns - different mean vectors

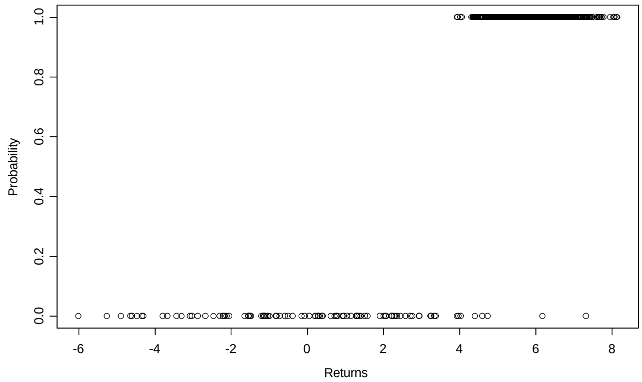

From these results it is not clear whether one should use univariate, bivariate or p-variate data. As for the estimation of parameters, some results have been derived by Bee and Flury (2000). Roughly speaking, they showed that the use of additional marginal distributions is only convenient when they provide relevant informations about the separation of the two populations. In our case the means of the two populations are almost identical, and therefore the contribution of additional marginal distributions is almost negligible. When the means are similar, the results obtained with the two approaches are similar as well.

Kim and Finger's (2000) method consists in estimating separately each bivariate mixture distribution. Therefore, if we are working with p risk factors and we want to estimate all the parameters, we have to estimate $p(p-1)/2$ different mixture distributions.

As to the problem of using the full-dimensional approach or Kim and Finger's approach, from a computational point of view two remarks have to be done. When we use the full-dimensional approach we get all the estimates by running the algorithm once, whereas if we estimate separately all the bivariate mixtures we have to run the algorithm $p(p-1)/2$ times. In addition, in the bivariate approach the estimates of the expected value and variance of a variable will be different depending on which bivariate distribution we consider: for example, if we have just three variables, we need to estimate either $\text{cov}(Y_1, Y_2)$ and $\text{cov}(Y_1, Y_3)$. In both cases we get an estimate of μ_1 and σ_1^2, but these two estimates will not be the same, and in general there is no a priori reason to prefer one of them. In our context, it might be reasonable to choose $\hat{\sigma}_i^2 = \max_{1 \leq i \leq p} \hat{\sigma}_i^2$, $i = 1, \ldots, p$.

On the other hand, if the dimension p is "too large", in the full dimensional case a problem of multicollinearity might arise, and the algorithm would not converge. From numerical experiments, p becomes too large when it is larger than, approximately, a number between 30 and 40.

The bivariate approach used in Kim and Finger (2000) overcomes the dimensionality problem, but is computationally much heavier.

(ii) VaR. As for the computation of a fat-tails VaR, we present either a real data example and some simulation results. We estimate a 13-dimensional normal mixture consisting of equities (Mibtel, Ftse, Dax, Standard & Poor's and Nikkey indices) and bond indices (1-3 years and 7-10 years baskets of government European, American and Japanese bonds); the data consist of daily returns covering the period Jun 2, 1993 - Feb 2, 2001 and were obtained from *Bloomberg*^{TM}. The portfolio is balanced so that the weights of equities and bonds are the same.

Table 1 shows three different VaR measures computed at four different confidence levels for the data at hand: we computed first the VaR based on the "hectic" covariance matrix $\hat{\mathbf{\Sigma}}^{(2)}$, which turns out to be, as expected, the largest one for all confidence levels. In addition, we computed a "normal" VaR based on the estimate of the covariance matrix obtained under the assumption that the observations are multinormally distributed and a non-parametric VaR obtained from the mixture data. As for this last measure, we generated 10000 observations from the mixture whose parameters are the estimates at hand and obtained the measures shown in the fourth column of table 1. This number should still be interpreted as the maximum loss (with 95% probability) in normal market conditions, as it uses all the observations, but it takes care of the fat-tails problem. Notice that at the 95% level the "normal" and "mixture" VaR are essentially identical, whereas, as the level gets larger, the mixture VaR is also larger than the normal VaR.

Table 1

	$\mathbf{VaR_{hec}}$	$\mathbf{VaR_{norm}}$	$\mathbf{VaR_{mix}}$
95%	−1.64	−1.04	−1.05
98%	−2.02	−1.30	−1.45
99%	−2.27	−1.49	−1.75
99.5%	−2.50	−1.65	−2.03

To give an idea of the precision of these VaR measures, we look at the actual losses. Indicating with r_c the proportion of times the return is smaller than c, we get the results shown in table 2.

Table 2

$\mathbf{r}_{-1.64}$	$\mathbf{r}_{-1.04}$	$\mathbf{r}_{-1.05}$
1.34%	5.44%	5.29%
$\mathbf{r}_{-2.02}$	$\mathbf{r}_{-1.30}$	$\mathbf{r}_{-1.45}$
0.75%	2.7%	2%
$\mathbf{r}_{-2.27}$	$\mathbf{r}_{-1.49}$	$\mathbf{r}_{-1.75}$
0.5%	2%	1%
$\mathbf{r}_{-2.50}$	$\mathbf{r}_{-1.65}$	$\mathbf{r}_{-2.03}$
0.25%	1.25%	0.75%

The conclusion is that the normal VaR underestimates the actual ex post loss, in particular for high confidence levels, whereas the mixture VaR shows a much better performance. On the other hand, the VaR obtained with the hectic covariance matrix overestimates the actual loss; this fact is not surprising as it is based only on the distribution of the hectic observations. The three distributions are shown in figure 5.

Fig. 5 - Three VaR measures - 99% level

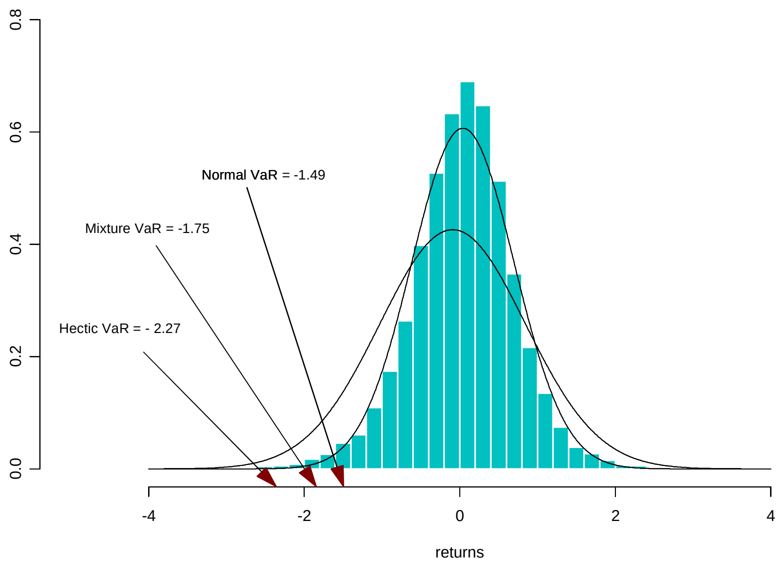

We have also estimated separately each bivariate mixture distribution and therefore each element of $\mathbf{\Sigma}^{(2)}$: in this case the "hectic" VaR turns out to be -1.65 (obviously the normal VaR obtained under the normality assumption remains equal to -0.71).

In a simulation setup equal to the one used for computing posterior probabilities, but with $n = 10000$, the three VaR measures (for a portfolio where each risk factor has a weight equal to 0.2) are shown in table 3 and the ex post actual losses in table 4. These results seem to confirm those obtained in the preceding real data example. For high confidence levels the mixture VaR is much more precise. Similar results are obtained by Zangari (1996), who uses mixture models in a Bayesian framework.

Table 3

	$\mathbf{VaR_{hec}}$	$\mathbf{VaR_{norm}}$	$\mathbf{VaR_{mix}}$
95%	-1.96	-0.85	-0.76
98%	-2.47	-1.05	-1.08
99%	-2.79	-1.21	-1.55
99.5%	-3	-1.35	-2.02

Table 4

$\mathbf{r}_{-1.96}$	$\mathbf{r}_{-0.85}$	$\mathbf{r}_{-0.76}$
0.52%	3.68%	4.76%
$\mathbf{r}_{-2.47}$	$\mathbf{r}_{-1.05}$	$\mathbf{r}_{-1.08}$
0.19%	2.19%	2.04%
$\mathbf{r}_{-2.79}$	$\mathbf{r}_{-1.21}$	$\mathbf{r}_{-1.55}$
0.12%	1.62%	1.08%
$\mathbf{r}_{-3}$	$\mathbf{r}_{-1.35}$	$\mathbf{r}_{-2.02}$
0.08%	1.31%	0.49%

(iii) Stress Testing. We now turn to the use of this model for stress testing purposes. The approach proposed by Kim and Finger (2000) consists in computing the conditional expectation (2) in the hectic population, i.e. using the parameters $\hat{\boldsymbol{\mu}}^{(2)}$ and $\hat{\boldsymbol{\Sigma}}^{(2)}$. Again, we ran a simulation to compare the performance of the five-variate and bivariate approach and to get an idea of how many observations are needed in order to obtain a good estimate of the conditional expectation. Therefore we first generated n observations from the five-dimensional mixture (3). Applying a shock equal to $+3\sigma_1$ to the last value of Y_1, we computed the expected value using the estimated parameters obtained with five- and two-dimensional data. With $n = 10000$, the true conditional expectation, the five-dimensional estimate and the bivariate estimate are respectively equal to

$$\boldsymbol{\mu}^{(2)'} = (\,0.264 \quad -0.462 \quad -1.320 \quad 0.330\,)\,;$$
$$\hat{\boldsymbol{\mu}}_2^{(2)'} = (\,0.237 \quad -0.630 \quad -1.756 \quad 1.041\,)\,;$$
$$\hat{\mu}_2^{(2)} = 0.282.$$

With $n = 5000$, we get

$$\boldsymbol{\mu}^{(2)'} = (\,0.196 \quad -0.344 \quad -0.981 \quad 0.246\,)\,;$$
$$\hat{\boldsymbol{\mu}}_2^{(2)'} = (\,-0.522 \quad -0.603 \quad -0.900 \quad 1.094\,)\,;$$
$$\hat{\mu}_2^{(2)} = -0.519.$$

With $n = 1000$, the results are

$$\boldsymbol{\mu}^{(2)'} = (\,0.266 \quad -0.466 \quad -1.330 \quad 0.332\,)\,;$$
$$\hat{\boldsymbol{\mu}}_2^{(2)'} = (\,1.00 \quad -1.777 \quad -3.132 \quad -0.024\,)\,;$$
$$\hat{\mu}_2^{(2)} = 1.167.$$

With $n = 500$ we get

$$\boldsymbol{\mu}^{(2)'} = (\,0.127 \quad -0.222 \quad -0.634 \quad 0.159\,)\,;$$
$$\hat{\boldsymbol{\mu}}_2^{(2)'} = (\,-0.891 \quad -1.166 \quad -0.087 \quad 0.198\,)\,;$$
$$\hat{\mu}_2^{(2)} = -0.873.$$

Whereas the differences between the bivariate and five-variate approach are negligible, it is clear that the sample size n plays a crucial role. In fact, it is known that, when the two populations are not well separated, the maximum likelihood estimates converge very slowly to the true values of the parameters [for details see McLachlan and Krishnan (1996), pag. 105-108].

Finally, we stress our portfolio. We consider a scenario where the stock markets fall: in particular we choose Standard & Poor's, Mib30 and Dax as core assets. The exogenous shocks on the core assets and the estimated movements of the peripheral assets are shown in table 5 and 6.

Table 5 - Shocks on the core assets

S&P	MIB30	DAX
−10%	−5%	−7%

Table 6 - Reaction of peripheral assets

Asset	Return
Nikkey	-3.93%
FTSE	-6.67%
USD1YR	-2.79%
USD41YR	-2.61%
GER1YR	0.003%
GER4YR	-0.02%
JAP1YR	-2.05%
JAP4YR	-1.92%
UK1YR	-2.02%
UK4YR	-2.2%

As was to be expected, yields fall as a result of people switching from equity to bonds.

Notice that using Kim and Finger's approach it would only be possible to compute the conditional expectations

$$E_2(Y_i|Y_j = y_j), \qquad i = 1, 2, \dots, p, \quad i \neq j.$$

This means that we can't stress more than one asset at a time and see how the peripheral assets react to this joint movement. Thus it would not be possible to perform this last experiment.

5 Conclusions

In this paper we showed how mixture models can be used for risk management purposes. The main applications concern fat-tails VaR calculations and stress testing. As for VaR calculations, it turns out that the non-parametric VaR obtained via Monte Carlo simulation using the covariance matrix of the mixture distribution provides a measure which is more precise than the one given by the normal VaR; a comparison with the actual ex post loss confirms the appropriateness of this VaR measure. The use of the parameters of the hectic distribution allows to perform stress test analyses using correlations and volatilities of the high-volatility periods; these periods are identified by the procedure via the posterior probabilities.

If possible (i.e., if the number of risk factors is not too large), it appears that the multivariate (full dimensional) approach should be used. Only when we deal with a very large number of variables we may be forced to use the bivariate approach, but in this case the estimation of parameters would become very expensive from a computational point of view.

References

Bee, M. (1998), *Misture Normali Multivariate: Inferenza con Dati di Dimensione Diversa*, Ph.D. thesis, University of Trento.

Bee, M., and Flury, B. (2001), "A Problem of Dimensionality in Normal Mixture Analysis", to appear in *The Scandinavian Journal of Statistics.*

Boyer, B., Gibson, M. and Loretan, M. (1999), "Pitfalls in Tests for Changes in Correlations", *International Finance Discussion Papers*, 597.

Fama, E. (1965), "The Behavior of Stock Market Prices", *Journal of Business*, 38, 34-105.

Kim, J. and Finger, C. C. (2000), "A Stress Test to Incorporate Correlation Breakdown", *Riskmetrics Journal*, May 2000, 61-75.

Loretan, M. and English, W. B. (2000), "Evaluating Correlation Breakdowns During Periods of Market Volatility", working paper, Board of Governors of the Federal Reserve System.

Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices", *Journal of Business*, 36, 394-411.

McLachlan, G. J. and Krishnan, T. (1996), *The EM Algorithm and Extensions*, New York, Wiley.

Zangari, P. (1996), "An Improved Methodology for measuring VaR", *RiskMetrics Monitor*, Second Quarter 1996, 7-25.

Collana ALEA Tech Reports

Nr.1 F. Sguera, *Valutazione e copertura delle opzioni binarie e a barriera*, Marzo 1999.

Nr.2 A. Beber, *Introduzione all'analisi tecnica*, Marzo 1999.

Nr.3 A. Beber, *Il dibattito su dignità ed efficacia dell'analisi tecnica nell'economia finanziaria*, Marzo 1999.

Nr. 4 L. Erzegovesi, *Capire la volatilità con il modello binomiale*, Luglio 1999.

Nr. 5 G. Degasperi, *La dinamica delle crisi finanziarie: i modelli di Minsky e Kindleberger*, Agosto 1999

Nr. 6 L. Erzegovesi, *Rischio e incertezza in finanza: classificazione e logiche di gestione*, Settembre 1999

Nr. 7 G. Degasperi, L. Erzegovesi, *I mercati finanziari come sistemi complessi: il modello di Vaga*, Settembre 1999.

Nr.8 A.Beber e L.Erzegovesi, *Distribuzioni di probabilità implicite nei prezzi delle opzioni*, dicembre 1999.

Nr.9 M.Filagrana, *Le obbligazioni strutturate nel mercato italiano: principali tipologie e problematiche di valutazione e di rischio*, marzo 2000.

Nr.10 A.Beber, *Determinants of the implied volatility function on the Italian Stock Market*, marzo 2001.

Nr.11 F.Bazzana, *I modelli interni per la valutazione del rischio di mercato secondo l'approccio del Value at Risk*, giugno 2001.

Nr.12 M.Bee, *Mixture models for VaR and stress testing*, giugno 2001.

I Tech Reports possono essere scaricati gratuitamente dal sito di ALEA: http://www.aleaweb.org.