Dipartimento di Economia e Finanza

Cattedra in Econometric Theory

Black Litterman for non-normal markets: focus on view's confidence level and performance

Relatore

Prof. Giuseppe Ragusa

Candidato

Mattia Menchi

664641

Correlatore

Prof. Nicola Borri

Anno accademico 2015/2016

Contents

Acknowledgements

I would first like to thank my thesis coordinator Prof. Giuseppe Ragusa for the expertise he provided, leading me through this project with ambition and humbleness. Moreover, my gratitude goes to the co-advisor Prof. Nicola Borri for his support.

I would like to thank the teams I have worked with from Ladbrokes and Fasanara Capital Ltd. They gave me a technical support with practical suggestions to improve the standards of this work.

I want to express my huge gratitude to my family who has motivated me until the last word of this thesis. Special thanks go to my friends and people who have been close to me.

Abstract:

Financial industry asks for tail risk protection, as the assumption of normally distributed markets does not hold anymore. The recipe presented by Meucci (2006) to generalise Black -Litterman Model without assuming any underlying distributions is here analysed on its ability to consider further moments of a multi-asset portfolio. Above all, we are questioning if it is convenient to input views. The asset allocation is computed from the Mean-CVaR efficient frontier. The performance of the resulting asset allocation aims at efficiency in cutting tail risk while it is tested on non-normally distributed returns with tail distributions modelled employing Extreme Value Theory. We then formulate a synthetic indicator to evaluate the impact of the view on portfolio allocation according to its confidence level.

Introduction and brief literary review

In the last 70 years, different theories have offered solutions to the problem of portfolio optimisation. The first pioneer in this discipline was Harry Markowitz in 1952 with his seminal paper "Portfolio Selection," published by the "*Journal of Finance*". This was the beginning of Modern Portfolio Theory: risk-averse investors can construct portfolios to maximise expected return for a given level of market risk, implying that risk is an inherent part of higher reward. This path leads to the construction of an efficient frontier of portfolios, which are expected to achieve the highest return (**mean**) for a given level of risk (**variance**). Mean variance model brought to life the revolutionary idea of diversification.

This basic portfolio model has been developed through several paths in the following years. The first alternative to minimum variance came out substituting the risk measure: Mean-Absolute Deviation model (MAD) by Konno and Yamazaki (1992) substitutes the mean squared error with the mean absolute deviation. Even this measure suffers from the "symmetric curse" giving the same weight to positive and negative variations around the mean. Along the same concept in 2000, Uriasev S. and Rockafellar, R.T. found an algorithm for efficient portfolio in a Mean-CVaR framework. For the first time the risk has been measured as a loss, or more precisely, the average loss in the "negative" tail of the worst percentile (usually 5%) of returns. As we will see later in this paper, this approach for efficient allocation does not rely on a specific distribution, allowing for better considerations on tail risk whatever would be the shape taken.

In 1992 Black, F. and Litterman, R. (1992) published on the Financial Analyst Journal an article called "Global Portfolio Optimization" providing an outstanding model to import active management into portfolio optimisation process. This fine technique allows combining investor's views and opinion on assets with CAPM and Mean-variance model. The pure strength was the tool to create a posterior distribution able to blend and to capture information from both market distribution and manager's opinion.

The underlying assumption of normal distributions involved in the model supports the Black-Litterman (BL) innovation. Growing needs from financial industry and challenges for the academic world have brought attempts to generalise BL model to every sort of distribution. We will discuss broadly the result obtained by Meucci who in 2006 proposes a solution by replacing Bayesian statistics of traditional BL model with a Copula Opinion Pooling approach to implement the view without assuming any specific distribution for market returns.

The following chapter set the ground for Black Litterman Model for non-normally distributed returns, describing in detail the previous step for efficient asset allocation.

In the first chapter, we are presenting the basic concepts for modern portfolio theory. We explain and compare Markowitz and Black Litterman models giving a deep perspective on the assumption in which they rely on. The aim of the work actually concerns the possibility to extend Black Litterman model beyond the normality assumption. Moving from the basic tool in portfolio optimisation we analyse the financial risk management measure of VaR and CVaR, introducing the advantages for portfolios' efficient frontiers, and therefore, optimal allocation in the Mean-CVaR framework.

The second chapter focuses on non-normality of asset's returns. First feature from the third and fourth moment are presented. As a second step, we describe the methodology suggested by Meucci (2006) to generalise Black Litterman Model and therefore on how to insert views in asset allocation process employing a Copula Opinion Pooling approach. The final step to complete the non-Gaussian framework needs to structure a Monte Carlo simulation method consistent with our purposes: Extreme Value theory provides a way to model tails of the distribution and a practical way to assess risk brought by excess kurtosis.

The third chapter puts the model at work. The main question to answer is whether it is convenient to insert a view. The performances are tested for scenarios where the view is correct or wrong. In addition we formalises a synthetic indicator to express how intensively the view influences the portfolio allocation. The analysis is concluded by considerations on the ability of this model to face non-normal markets.

CHAPTER 1: Portfolio theory

1.1 Mean-Variance Model

Markowitz approach moves from a set of background assumption:

The investors want to maximize the returns for a given level of risk: this statement is implied by the fact that most of the investors are risk adverse.
Your portfolio includes all of your assets and liabilities: comprehensive approach implying the hypothesis of complete markets.
A good portfolio is not simply a collection of individually good investments: the portfolio maximises the utility function of a risk adverse investors and takes into account the relationships among assets returns in the portfolio.

The first practical and quantitative step is to estimate a set of parameters:

1.1.1 Return

The return of an asset over a single period is computed by

$$R_t = \frac{P_t}{P_{t-1}} - 1$$

For every asset, we need to estimate the average return over a one-period of time. Usually simulations or historical data are the sample for which the mean return is computed. For the upcoming analysis, data are collected on a daily basis and the arithmetic average is computed.
For a portfolio with N assets, its return is the weighted average of the return of the assets composing the portfolio, while the weights are exactly the percentage of portfolio's value invested in the respective asset. Organising the weights in a vector $w = (w_1 , w_2 , w_3 , \dots , w_n)$ the return on a portfolio for a given period is given by

$$R_p = w' R$$

where R is the vector of assets returns $R = (R_1 ,R_2 ,R_3 ,\dots,R_N)$.
In Portfolio management, rates of return are modelled as random variables for which a certain probability distributions function is assumed: the choice of the

random variable depends on the assumptions, namely the most famous and common models rely on the assumption of normally distributed returns; the purpose of this paper is to find a reasonable alternative to this binding perspective. Expected return is therefore the firs moment of the portfolio.

$$E(R_p) = w' \, E(R)$$

1.1.2 Covariance Matrix

Risk is the other fundamental parameter to be estimated. Risk is the uncertainty about future outcomes. An alternative, and a more intuitive, definition would be the probability of certain negative scenarios. The most common measure of risk is the variance of returns namely the average of squared deviation from the expected return.

In a portfolio framework with N assets, we have to deal with N random variables generating returns. This problem can be modelled using a multivariate random variable (n-variate random variable) where the expected value is a 1xN vector of expected returns, and the variance a symmetric NxN matrix. It is actually fundamental to take into account how assets' returns covariate, enabling to appreciate the diversification effect.

The computation of the covariance is performed through the usual formula:

$$cov(R_i \, , R_j) = \sigma_i \, \sigma_j \, \rho_{ij}$$

where ρ is the Pearson correlation coefficient

The variance covariance matrix would be composed as follow

$$\sum = \begin{bmatrix} cov(R_1 \, , R_1) & \cdots & cov(R_1 \, , R_N) \\ \vdots & \ddots & \vdots \\ cov(R_N \, , R_1) & \cdots & cov(R_N \, , R_N) \end{bmatrix}$$

According to Markowitz and to the mean variance framework, is now possible to assess the risk of a portfolio as the variance (a scalar) around its expected return given a pre-determined vector of weights to be applied to its assets

$$\sigma^2 = \sum_{i=1}^{n} w_i \; \sigma_i^2 + \sum_{i=1}^{n} \sum_{j<i} w_i \, w_j cov(R_i \, , R_j \,) = w' \, \Sigma \, w$$

1.1.3 Portfolio optimisation

The aim of Modern Portfolio Theory is to select the optimum portfolio. So far, we have described how to achieve the first and the second (centralised) moment of the portfolio. Mean and standard deviation are the only sufficient parameters to describe a Normal distribution. Markowitz optimisation computes the efficient allocation according to these two parameters: it is actually easy to guess an underlying Normality assumption for market returns

The optimisation is performed subject to the weights of every asset in the portfolio and hence performed to maximise the utility function of a hypothetical investor. According to the theory developed by Markowitz, most of the investors are risk averse. The choice of an utility function showing the risk aversion feature can be performed among a wide set of functions and the selection of a specific one is led by subjective reasons on a case by case basis.

This problem can be brilliantly solved by noting that every consideration on risk aversion ends up with a minimisation of the risk, namely the variance of our portfolio computed as above, keeping the expected return at a certain level.

$$\min_{w \in \mathbb{R}^n} \frac{1}{2} w' \Sigma w$$

$$\text{s.t.} \quad w' E(R) = \mu$$

$$w'\mathbf{1} = 1$$

This minimisation problem is solved using numerical methods. Speed consistency and stability of estimated parameters are crucial points that we will discuss later about weaknesses of this model.

Using Lagrange multipliers to solve the optimisation,

$$\min_{w \in \mathbb{R}^n} \frac{1}{2} w' \Sigma w - \lambda w' E(R)$$

we see that the multiplier λ is the parameter that calibrates the risk aversion. In this way, we can give a generalised and objective measure to the risk aversion.

The solution for the optimal weights of the portfolio are defined as follow

$$w^* = (\lambda \Sigma)^{-1} E(R)$$

The solution is a vector giving weights to the N assets included in the portfolio. In other words, it is the optimal portfolio.

Remembering that we have just minimised the variance keeping the expected return at a determined level, it is possible to perform the very same optimisation for every level of expected return. The output would be an optimal portfolio, the one that minimises the risk, for each target expected return.

Plotting every single combination of risk-return, we obtain a hyperbola, where its upper edge represents the efficient frontier.

1.2 Black-Litterman Model

In this paragraph, we describe the original version of Black Litterman starting to appreciate the innovativeness of the model. A pointwise comparison with the famed Markowitz model will follow.

First, Black-Litterman model is an equilibrium model: the starting point is the market portfolio, where the weights are indicated by the full diversification and according to the exposure to the unique risk factor, namely market risk. It is a way to simplify the procedure concerning parameter's estimation, keeping in mind all the weaknesses that such a procedure carries with itself.

Anyway, historical data or simulation can be exploited to estimate the parameters, and this is actually, what we are going to see in the implementation of our modified model.

Before starting to present the model, it is useful to state the aim of BL: estimate the distribution (mean and variance) of the returns taking into account that the investor has expressed views about assets' performance. From the passive market portfolio from equilibrium model, we have switched into the world of active management.

1.2.1 The model

First question arising is: "What is meant by view?"

A view is an opinion or a statement about the market. This might appear quite a generic definition, actually Black –Litterman framework exclusively considers views on expected returns. In other words, these views are linear and arranged in a matrix.

There are two kinds of views: relative and absolute. A relative view concerns the comparison with another asset, namely if an asset will outperform or underperform another one; in this case the weights in the respective row of the matrix will sum up to zero.

An absolute view set up the comparison between an asset and the entire portfolio, again in terms of over/underperformance; in this case, the sum of the weight in the row will sum up to one.

In addition, the set of views do not have to cover every asset, in fact they can even conflict with each other.

Formalizing this theory using an example, we assume K views on a portfolio of N assets. The K views are usually organised into a pick matrix, called $P_{K \times N}$: it is easy to realize that every row describes a single view. Having in mind the structure of linear system described with matrices, we now need a vector containing information about expectation on the view. This vector is called *a* with dimension $K \times 1$. Finally yet importantly, we have to introduce a confidence level for the views: the matrix Ω will be a diagonal matrix containing the percentage of confidence of the views. The matrix Ω is diagonal by construction because the views are assumed to be uncorrelated with each other.

Modelling assets' portfolio with returns $R = (R_1 ,R_2, \dots ,R_N)$ we set the following multivariate normal distribution

$$R \sim N (\mu, \Sigma)$$

Where μ is the vector containing the expected return of every asset while $\sum$ is the variance covariance matrix, both parameters estimated like in Markowitz model. Hence, we need to define the expectations again as a multivariate normal distribution.

$$\mu \sim \mathrm{N}\ (\theta, \Sigma_\mu)$$

It is interesting to notice that Σ_μ is equal to Σ scaled by a parameter τ representing the uncertainty around the estimation of μ.

Coming back to the active management problem, we organise the views according to a normal distribution

$$\mathrm{P}\mu \sim \mathrm{N}\ (a, \Omega)$$

From this statement follows that

$$a = \mathrm{P}\mu + \mathrm{H}$$

where H is a normal distribution $\sim$ N $(0, \Omega)$. Expectations on the views can therefore be modelled as a random variable A conditioned on the realisation of μ

$$\mathrm{A}|\mu \sim \mathrm{N}\ (\mathrm{P}\mu, \Omega)$$

Determining the confidence level matrix Ω is not a straightforward task. A long debate has tried to disentangle a procedure to estimate objectively the confidence level. Meucci (2006) has provided an easy and intuitive solution

$$\Omega = \frac{1}{c}\ \mathrm{P}\ \Sigma\ \mathrm{P}$$

Where the c is a scale, parameter showing the level of confidence on views considering them jointly. The scalar c does not share the same nature of the terms in the principal diagonal of Ω: c is not a percentage, its value varies in the interval $(0, \infty)$ where it is 0 for the null confidence and ∞ meaning the certainty.

Concluding the treatise about the views, we provide an example. Suppose to have 5 assets and 2 views: the first is an absolute view in which the asset 1 will have a return of 1% with confidence level η_{22}. The second is a relative view according to which asset 3 will outperform asset 5 by 5% with confidence η_{11}.We therefore formalise the problem with the following matrices

$$P=\begin{pmatrix}1 & 0 & 0 & 0 & 0\\ 0 & 0 & 1 & 0 & -1\end{pmatrix} \quad \Omega=\begin{pmatrix}\eta_{11} & 0\\ 0 & \eta_{22}\end{pmatrix} \quad A=\begin{pmatrix}5\%\\ 1\%\end{pmatrix}$$

1.2.2 Posterior distribution

The aim of Black-Litterman model is to find a posterior distribution of assets' return, namely the distribution of the views conditioned on the realization of estimated return μ. We can now introduce the basic concepts of Bayesian statistics using the previous notation for distributions.

Considering A and μ two events, in probability theory we have that

$$P(\mu|A)=\frac{P(A\cap\mu)}{P(A)} \quad \text{and} \quad P(A|\mu)=\frac{P(\mu\cap A)}{P(\mu)}$$

From which follows Bayes theorem

$$P(\mu|A)=\frac{P(A|\mu)\times P(\mu)}{P(A)}$$

Or equivalently using probability density functions

$$f_{(\mu|A)}(\mu)=\frac{f_{(A|\mu)}(a)\quad\times f_{(\mu)}(\mu)}{f_{(A)}(a)}$$

We know that μ are the parameters, more precisely a multivariate normal distribution as stated before. This distribution is the prior distribution in Bayesian statistics, while the result of the latest equation, f(μ|A) (μ), is the posterior distribution and $f_{(A|\mu)}(a)$ is the likelihood function.

Practically speaking, the set of estimated parameters μ, are the implied returns from CAPM, defined as a prior equilibrium distributions, or directly estimated from simulation or historical data. The information processing that follows is improperly defined as a Bayesian statistics, while it would be more precise to define it as a conditional probability. Actually, the distribution μ is then conditioned on the "observed data", namely the views. As a result, the posterior distribution is the conditional probability that is obtained after that the relevant new information is taken into account. In other words we are updating information.

$$\textit{Posterior distribution} \propto \textit{prior distribution} \times \textit{likelihood function}$$

Straight back to our Black Litterman portfolio, we obtain the posterior distribution of μ using Bayes' formula

$$\mu|A \sim N(\theta^{post}, \Sigma_{\mu}^{post})$$

where

$$\theta^{post} = [(\Sigma_{\mu})^{-1} + P'\Omega P]^{-1}[(\Sigma_{\mu})^{-1}\,\theta + P'\Omega^{-1}V]$$

$$\Sigma_{\mu}^{post} = [(\Sigma_{\mu})^{-1} + P'\Omega^{-1}P]^{-1}$$

So far so good, we have obtained a posterior distribution of the estimated parameter, so the following step will be to find a posterior distribution of assets' returns R.

$$R|A \sim N(\mu_{BL}, \Sigma_{BL})$$

where

$$\mu_{BL} = [(\Sigma_{\mu})^{-1} + P'\Omega P]^{-1}[(\Sigma_{\mu})^{-1}\,\theta + P'\Omega^{-1}V]$$

$$\Sigma_{BL} = \Sigma + \Sigma_{\mu}^{post} = (1+\tau)\,\Sigma - \tau^2\,\Sigma P'\,(\tau P\,\Sigma\,P' + \Omega)^{-1}\,P\,\Sigma$$

There are several interpretations of the obtained result.

Remembering that we started by estimating a market equilibrium μ in which views apply, a paper by Fabozzi, Focardi and Kolm (2008)[1] suggests to interpret return distribution as a linear combination of market equilibrium θ and the estimated distribution μ weighted by the confidence levels concerning both views and estimated parameters. Given the result, the confidence level for our estimates of market equilibrium is $(\Sigma_\mu)^{-1}$ while the confidence level of our views corresponds to $P'\Omega P$. Describing further the linear combination formalised by Fabozzi, Focardi and Kolm (2006) we can compute the weights given to θ and μ

$$\omega_\theta = [(\tau\Sigma)^{-1} + P'\Omega^{-1}P]^{-1}(\tau\Sigma)^{-1}$$

$$\omega_A = [(\tau\Sigma)^{-1} + P'\Omega^{-1}P]^{-1}(P'\Omega^{-1}P)$$

where $\omega_\theta + \omega_A = I$.

Stress-testing this result when $(\Sigma_\mu)^{-1} \rightarrow 0$ we see that $\omega_A \rightarrow 1$ so the Black Litterman return will get close to the value stated by the views. On the other hand when $(P'\Omega^{-1}P) \rightarrow 0$ follows that $\omega_\theta \rightarrow 1$.

1.3 Brief comparison: Black Litterman vs Markowitz

It is easy to realise how Black-Litterman has introduced active management in the framework of Portfolio Theory and asset allocation problems. The chance to implement views in portfolio optimisation starting from a market equilibrium model, passive by definition, has rewritten the rules in asset management industry.

From a pragmatic perspective, Black-Litterman is able to smooth the principal weaknesses of mean-variance approach. Mean variance optimisation is indisputably difficult to be performed mainly for two reasons: number of parameters to be estimated and sensitiveness of the results.

For every single asset, Markowitz theory needs to estimate expected return, standard deviation and correlation with all the other assets. More precisely for a

[1] Fabozzi, F., Foccardi, S., Kolm, A., 2008, "Incorporating trading strategies in the Black-Litterman framework", The Journal of Trading

portfolio of N assets, we have N expected returns, N variances and $(N^2-N)/2$ covariances. This is extremely time consuming and it embeds a high amount of discretion in selecting the sample in which to perform the estimates. Moreover, the same accuracy is needed for every asset, also for the ones not taking parts into the active investment decision.

Once we have the optimum portfolio or the entire efficient frontier we realise that our result is extremely sensitive to the estimated parameter used as inputs. A common problem with these portfolios is that they end up being unreasonably concentrated in few assets and extremely dependent on the selected sample. Portfolio managers can then use resampling technique to normalise the estimated parameters and to give robustness to their strategies: they compute an average of optima for a certain combination of mean-variance. Another solution to the excessive concentration of portfolios is to set a minimum diversification: constraining the weight to desirable ranges is after all a palliative care. Regarding strong dependence of estimated parameters to the sample can be partially overcome by shrinking the covariance matrix, giving more weight to the principal diagonal composed by variances[2].

Black-Litterman, on the other hand, faces brilliantly this problems estimating two sets of expected returns, one for the passive investment and one for the active side.

The passive side concerns the "equilibrium" expected returns, namely what the market has decided a level of return and the investor is price taker. These returns are obtained using reverse engineering from a benchmark structured as the Capitalization weighted portfolio and they are the appropriate result for a passive investor, so a neutral vantage point for the active manager.

As already mentioned, the active side concerns the views. What about assets not covered by any view? The weight in the portfolio is maintained close to the one revealed by the benchmark. For the other class of asset the weights are adjusted consistently with the views and respecting the trade-off between risk and return.

[2] Ledoit, O., Wolf, M., 2003, "Improved Estimation of the Covariance Matrix of Stock Returns With an Application to Portfolio Selection", Journal of Empirical Finance

Actually, after having obtained the posterior distribution you can opt for several portfolio optimisation solutions, and traditionally mean variance optimisation is preferred. The Black-Litterman portfolios are more stable and less sensitive to input parameters having less degree of freedom in the overall estimation, but above all the asset allocation is better calibrated to fit manager's investment decision, compared to any heuristic approach concerning constrained weights.

Mean variance optimisation of posterior distribution is usually preferred for its applicability to normal distributions. When we analyse assets with non-normally distributed returns, portfolio optimisation can be performed using other risk measure to substitute the variance. The most common alternative is the Mean-CVaR approach. Conditional VaR or Expected shortfall describes a loss function over a certain percentile of returns. The advantage of using CVaR is that the loss function is concave, so it is possible to perform a numerical optimisation. We leave for the following chapters a deep analysis of Mean-CVaR optimisation.

1.4 Financial risk management measures

In financial risk management, it is crucial to select the convenient risk measure to calibrate the models. Talking about risk, the standard deviation of unexpected outcomes is the first risk measure to think of. Anyway, standard deviation is relevant under normality of returns. Moreover, it describes the variation around the expected value for both sides, positive and negative variations with the same weight. As everybody can figure out, a person not familiar with financial measures would be much more concerned about the likelihood of losses than about the deviations from the mean: this is why measures like VaR and CVaR (or expected shortfall ES) have a good practical advantage.

1.4.1 VaR

JPMorgan firstly pioneered VaR, and it captures the idea of maximum loss for a determined confidence interval in a given amount of time. VaR measure is conventionally reported as a positive number X, equal to the amount of the loss. There are two ways to compute VaR: parametric, namely assuming a random variable generating the outcomes, or non-parametric, using only the sorted series of outcomes.

The formula for parametric VaR for a confidence level equal to α of a random X is defined by

$$\mathrm{VaR}_{\alpha}(X) = \inf \{ F_X(x) > \beta \}$$

where F(x) is the cdf of X.

Equivalently, exploiting inverse cdf and quantiles we can rephrase easily the previous equation

$$\mathrm{VaR}_{\alpha}(X) = F^{-1}(\beta)$$

For a non-parametric approach it is sufficient to sort the outcomes of the random variable, namely the returns and to take the element in the $(1-\alpha)^{th}$ percentile.

The problem arises in portfolio optimisation realising the weakness of this risk measure. Actually, VaR is not a coherent risk measure. We are now showing the requirements for a function in order to be a coherent risk measure.

Given a function π of a random variable X, if π shows these properties

- Monotonicity of the function: if an asset performs worse than another in every likely state does, than its risk measure must be higher.
- Translation invariance: $\pi(X+c) = \pi(X)+c$ or in other words if a c amount of risk free asset is added to a portfolio, the risk measure should decrease by c.
- Homogeneity: if $c>0$ then $\pi(Xc)= c\pi(X)$ or in other words the risk measure is scaled by the size of the portfolio, other things being equal.
- Subadditivity: $\pi(X + Y) < \pi(X) + \pi(Y)$ or in other words merging two portfolios, the resulting risk measure should be less than the sum of the single risk measures.

Unfortunately, VaR is not a coherent risk measure because of the lack of subadditivity property. Actually, VaR does not take into account the benefits from diversification. Moreover performing a portfolio optimisation Mean-VaR we could face a non-convex region, importing difficulties in the numerical minimisation.

1.4.2 CVaR

Considering the technical weaknesses of Value at Risk, a conceptually close measure was developed by Rockafellar and Uryasev [3] called Conditional Value at Risk or Expected Shortfall. It is called conditional because it investigates the loss function in the "negative" tail, right after the threshold indicated by the VaR.

Value at Risk states the likely maximum loss with a certain confidence, but it gives no information about what happens right after the VaR loss. Conditional VaR is the expected loss conditional that the loss exceeds VaR level. It answers the question, "if things go bad, how bad can they go?".

$$\mathrm{CVaR}_\alpha(X) = E[X|X > \mathrm{VaR}_\alpha(X)]$$

CVaR is a coherent risk measure and a convex function. A way to compute it non-parametrically consists on taking the average of outcomes worse than VaR threshold.

Knowing the distribution generating the outcomes, Rockafellar and Uryasev provide a formula to compute this measure

$$CVaR_\alpha(X) = \int_{-\infty}^{+\infty} y f_x^\alpha(y)\, dy$$

Where

$$f_x^\alpha(y) = \frac{1}{f_x^\alpha(y) - \alpha} \quad \text{if } y \geq \mathrm{VaR}_\alpha(X) \text{ otherwise is equal to 0.}$$

The parametric approach meets a specific formula different for every kind of distribution. The challenge is to adapt this risk measure in a portfolio optimisation framework.

[3] Rockafellar, R.T., S. Uryasev, 2000 "Optimization of conditional value-at-risk, Journal of Risk and
Rockafellar, R.T., S. Uryasev. 2002, "Conditional Value-at-Risk for General Loss Distributions", Journal of Banking & Finance

1.5 Portfolio optimisation Mean-CVaR

Portfolio optimisation using Mean-CVaR combination is performed in order to achieve a corresponding efficient frontier. As we have seen for the theory by Markowitz, we can reach an efficient frontier by minimising the variance keeping the expected return at a certain minimum level, or equivalently maximising the expected return for a given level of variance.

According to the pioneer work by Rockafellar and Uryasev (2000), the optimisation was considered as minimizing CVaR for a minimum expected return. According to a theorem formalised by Pavlo Krokhmal, Jonas Palmquist, and Stanislav Uryasev (2001)[4] there are three equivalent ways to formulate the problem:

$$\min_w CVaR(w) - \mu R(w) \qquad \text{where } w \in W \text{ and } \mu > 0$$

$$\min_w CVaR(w) \quad \text{s.t. } R(w) \geq r^* \qquad \text{where } w \in W \text{ and } \mu > 0$$

$$\max_w R(w) \quad \text{s.t. CVaR} \leq L \qquad \text{where } w \in W \text{ and } \mu > 0$$

Given that CVaR (w) is a convex function, R(W) is concave and W is a concave set, the three versions to optimise end up with the same frontier allowing short-selling and offsetting the long short positions.

The algorithm computes the optimal weights for which the portfolio belongs to the efficient frontier, namely the best expected return-CVaR combination, using the information provided by the posterior.

The choice of this methodology to compute the optimum portfolio is led by the assumption of non-normal markets. As we have already explained in the previous chapter, most of the features not covered by Markowitz optimization concern the probability of events in the tails of the distribution: fat-tails are unanimously

[4] Pavlo Krokhmal, Jonas Palmquist, and Stanislav Uryasev, 2001, "Portfolio optimization with conditional value-at-risk objective and constraints"

considered a problem for portfolio managers and risk specialists. CVaR measure describes the risk nestled on the tail of the distribution, giving weight to extreme events that normal model would not even take into consideration. Large success of the CVaR approach comes after having realized that extreme negative market conditions can deplete portfolio performances over the worst expectations simply because no defence was taken against these events. In addition more and more investment manager has matched the preferences of their investors by cutting tail risk and offering downside protection as a mandate.

CHAPTER 2: Living in a skewed and leptokurtic world

Modern Portfolio Theory has defined the principal paradigm behind the one best way in constructing efficient asset allocation. Unfortunately, sometimes results are not as good as expected. Portfolio models cannot predict and react to every situation especially during extreme events. The point is that financial markets are excessively complicated, and the best model to represent the world is the world itself. The aim is towards the best approximation.

The friendliest way to fit the distribution of returns is the Gaussian bell: a heavy assumption for a less painful procedure. What could be more painful than a complicated model are the losses the Gaussian distribution would have never considered. Thinking about deviations from Normality is straightforward to address the discussion to asymmetry and fat-tails in distribution of returns.

Skewness

Asymmetric distributions have skewness different from zero: in other words, the mass of the outcome of the distribution does not take place symmetrically around the mean. When the left tail is longer (left-skewed) and the outcomes are concentrated on the right of the “bell”, than the distribution has a negative skew. On the other hand, a positively skewed distribution sees its right tail longer and the mass concentrated on the left.

Fixed income securities are the perfect example for skewed distribution of returns. While the probability of default is relatively small and the probability of receiving the coupons and the face value at maturity is relatively high, the distribution is truncated for more positive returns, so it is possible to assess a negative skewness.
The most common measure of risk, namely the standard deviation, is a measure of dispersion symmetrically around the mean and it would not consider the skewness. Whether to consider desirable positive or negative skewness is empirically sensible to the entity of the asymmetry.
According to classical financial theory, investors always prefer positively skewed returns. After all, positively skewed tails head toward positive infinity while

negatively skewed tails point toward the negative infinite; moreover, for positive skewness the mean falls above the median. Recent developments in behavioural finance have observed empirically a preference for deeply negatively skewed returns. The key to interpret this phenomenon relies on prospect theory: “A change from impossibility to possibility or from possibility to certainty has a bigger impact than a comparable change in the middle of the scale.”

The preference for positively skewed pay-offs is justified through the “longshots bias”, explaining the popularity of lotteries, because of the overweight events that are possible although the low probability. The preference for negatively skewed pay-off relies on the fact that positive outcomes have high probability to occur but not the certainty and for this reason are underweighted. Utility from skewness depends mainly on the weights addressed to the distribution of losses, where many investors are concerned to what would be their fate during severely distressed situation.

After these considerations, the choice of a meaningful risk measure becomes the main driver for efficient asset allocation. Conditional VaR is a risk measure that gives enough relevance to what happens to the portfolio on the left tail of the distribution, answering to the question how much do I expect to lose in the worst X% of cases?

We can compute the skewness for a multi-asset portfolio employing an approach similar to the one used to estimate the portfolio variance. In this case, the covariance matrix would be substituted by the co-skewness matrix. This matrix has dimension N x N^2 and it is here defined:

$$M_3 = \mathrm{E}\,[(\mathrm{R}-\mu) \times (\mathrm{R}-\mu)' \otimes (\mathrm{R}-\mu)'$$

From which follows that the measure for multivariate skewness is computed as:

$$S_p = \mathrm{w}' M_3 (\mathrm{w} \otimes \mathrm{w})$$

Kurtosis

Another deviation from normality concerns the concentration of outcomes clustered on the tails of the distribution.

What makes portfolio managers worried is that Normality assumption underestimates the probability of catastrophic events. The problem with kurtosis is better addressed in terms of "excess Kurtosis": Gaussian distribution has a kurtosis equal to 3, and then a distribution showing a larger value is deemed as "leptokurtic", and "platykurtic" for lower values. A leptokurtic distribution is more peaked than a Gaussian "bell". It follows that small changes are less frequent while extreme events are far more likely to happen than the ones contemplated by Normality assumption because the tails are considerably fatter and longer.

The motivation to dread leptokurtosis relies on assessing the likelihood of potential large losses to our portfolio, so excess kurtosis is undoubtedly perceived as a negative feature for the investor. Again, Conditional VaR is a useful tool to investigate what happens in the tails of the distribution.

Multivariate kurtosis is again measured thanks to a multidimensional matrix called co-kurtosis matrix of dimension N x N^3:

$$M_4 = \mathrm{E}\left[(\mathrm{R}-\mu) \times (\mathrm{R}-\mu)' \otimes (\mathrm{R}-\mu)' \otimes (\mathrm{R}-\mu)'\right.$$

From which follows that the measure for multivariate skewness is computed as:

$$S_p = \mathrm{w}' M_4 (\mathrm{w} \otimes \mathrm{w} \otimes \mathrm{w})$$

In the next paragraphs, we will analyse if the third and the fourth portfolio moments behaves in the process to implement the view. First consideration concerns the posterior and its moments for assets univariate distributions: are they going to change from the same measures found in the prior? Second analysis will cover the multivariate frame, namely how skewness and kurtosis differ for allocations depending on the confidence level.

2.1 Implementing views in a non-normal world

Active portfolio management seeks a method to include manager's view in asset allocation procedures, trying to keep aside naïve and heuristic biases.

Information processing algorithms allows filtering the information from the market jointly with the opinions stated by the manager, creating a posterior distribution of returns. Dealing with non-normally distributed returns (prior) we should look after the features of the posterior distribution, namely how the distribution has been reshaped.

As we have seen in the first chapter, the common model to process the information is the Black-Litterman model. This model is able to blend smoothly the views of the investor with the prior distribution from the market, assuming that every distribution involved is Gaussian. Given its popularity, financial industry and academic world has tried to extend the model to non-normal markets.

Meucci (2006)[5] provided an innovative method to insert views and to infer a posterior distribution starting from non-normally distributed returns using a Copula-Opinion Pooling (COP) approach.

Meucci has applied COP to Black–Litterman's views in two stages in 2006. First stage relies on a specific quasi-normality assumption: the market is assumed to be skew-t distributed just because linear combinations of this kind of distributions are again skew-t distributions creating a situation almost analogous to the normal case. This property is one of the main reason to make Gaussian distribution easy to handle and that allows the functionality of the Black-Litterman model in its original formulation. The second stage starts from an application of COP approach to reach a generalisation of the model for any kind of distribution: constructing a market prior through a very large number of Monte-Carlo simulations, COP can easily compute a posterior distribution.

2.2 Copula Opinion Pooling: theory behind the algorithm

The algorithm aims to compute a posterior distribution. First set of information is provided by the market in a N-dimensional vector: a set of returns from a bunch of

[5] Meucci, A. 2006 "Beyond Black-Litterman: Views on Non-normal Markets"

N assets or a collection of risk factors are the usual input to describe the stochastic variables determining the market randomness. At this point, it is useful to represent this random variable M in terms of their probability density function (f) or cumulative density function (F).

As we have already exposed broadly, the crucial part of this model are the views. As in the BL model, the views are a set of $K \leq N$ statements of any linear combination of the market vectors.

In this model, we observe a slight difference from the usual BL world: we have the option to express views directly on the market realizations M instead of the only distribution's parameters. The above linear combinations give shape to a $K \times N$ dimensional "pick" matrix P.

As a result, the random vector describing the views is defined by

$$\mathbf{V} = \mathbf{P} \times \mathbf{M}$$

We represent each view in terms of a probability density function: the model assumes that the views are expressed through a uniform distribution to better respect the feature of opinion on market realization and to take into account the non-normal market structure. The process to achieve a posterior consistent with non-normal market needs to stay clear from add information assumed exogenously: according to maximum entropy framework, uniform distribution is actually the random variable that maximises the entropy under no constraints.

Market random variable could have any shape and uniformly distributed views needs to end up in a blending posterior: the views are then reshaped by the market distribution M.

The heart of the procedure lies in the opinion pooling procedure: in this specific case, a "bottom up" approach has been selected. Therefore, the first step is about computing the posterior marginal distribution of each view separately: we can actually replicate the view by a weighted average of the market-structured prior pdf and the view's pdf itself.

The weight given to the market or to the view depends on the confidence the investor assesses for the opinion: actually, we can assimilate this weight to the confidence level of the views in the traditional BL framework.

The joint posterior distribution of the views comes as the next step using the dependence structure introduced by the market prior (see Monte Carlo simulation using copula).

We fit a copula to the marginal distribution of V

$$C = (F_{V_1}\ (V_1)\ ,\ F_{V_2}\ (V_2),\ \dots\ ,\ F_{V_k}\ (V_k))\ '$$

Then it possible to constrain on the market structure to achieve the joint posterior distribution as follows:

$$J =\ (\ F_{S_1}^{-1}\ (C_1),\ F_{S_2}^{-1}\ (C_2),\ \dots\ ,\ F_{S_k}^{-1}\ (C_k)\ )\ `$$

where $F_{S_n}^{-1}$ is the quantile function of the cumulative distribution function.

Finally, joint posterior distribution is computed by blending the views in a suitable set of market coordinates. Before that, we need to notice that the market vector can be expressed in a set of view-adjusted coordinates:

$$M \Leftrightarrow \begin{pmatrix} V \equiv PM \\ W \equiv P^{\perp}M \end{pmatrix}$$

Where W is intuitively a matrix with (N-K) rows to describe the entries unaltered by the views. By reverting the now defined matrix M to the market coordinates, we end up with a posterior distribution.

$$M^* \Leftrightarrow \begin{pmatrix} P \\ P^{\perp} \end{pmatrix}^{-1} \begin{pmatrix} J \\ W \end{pmatrix}$$

2.3 Algorithm

The algorithm to compute the posterior is borrowed from Meucci (2006)[6] and we now propose the list of five steps to insert views without assuming a specific distribution, not only traditional Gaussian, for our market random variable.

- ✓ rotate the market prior into the views' coordinates;
- ✓ compute the views prior cdf's and the market-implied prior copula;
- ✓ compute the marginal posterior cdf of each view;
- ✓ compute the joint posterior distribution of the views;
- ✓ compute the joint posterior realizations of the market distribution;

The posterior needs to show some consistent features and characteristics. First intuitive requirement concerns the fact that by imputing a confidence level equal to zero, the posterior is not different from the initial market structure. On the opposite side, we notice that a confidence level equal to certainty 100% will create a posterior where the vector related to the asset interested by the view shows a uniform distribution, namely the distribution of the view itself.

Given our care to describe market randomness according the first four moments, mean, standard deviation, skewness and kurtosis, it is important to check how the process leading to the posterior has affected these parameters. In practice, we see a mean consistently modified according to the view while the higher moments are barely affected. Asset uncovered by any view are left completely unaltered.

[6] Meucci, A., 2006 "Beyond Black-Litterman in practice: a five-step recipe to input views on non-normal markets"

2.4 Monte Carlo Simulation adjusted for EVT[7]

So far, our model has provided the weights to our efficient portfolio consistent with the views. Now we give the priority to assess the efficiency in information processing and therefore how the resulting asset allocation performs.

We are going to evaluate portfolio performances upon simulated scenarios reflecting the characteristics of non-normal markets. More specifically the joint simulation needs to replicate the abnormalities detected from the input market data. Adjustment for Extreme Value Theory (EVT) helps us modelling tails of the joint multivariate distribution of assets in our portfolio.

Recent researches applied EVT for extreme variations to financial markets to handle credit crises, currency shocks and other events considered impossible by traditional parametric models.

The tail behaviour of financial series has been discussed in Longin (1996), McNeil (1999), Jondeau and Rockinger (1999), Neftci (2000), McNeil and Frey (2000). This Monte Carlo simulation employs t-copula and EVT to better replicate tail-risk.

The process is articulated into several steps. First step we estimate an asymmetric GARCH model to data[8]. Autoregressive component deals with autocorrelation on data:

$$R_t = c + R_{t-1}\Theta + \varepsilon_t$$

Asymmetric GARCH component cope with heteroscedasticity and asymmetry:

$$\sigma_t^2 = \varsigma + a\sigma_{t-1}^2 + \phi\varepsilon_{t-1}^2 + \varphi[\varepsilon_{t-1} < 0]\varepsilon_{t-1}^2 .$$

At this point, we are interested in the filtered residuals related to every single asset. Hence, residuals are standardized and modelled according to a Student-t distribution to take into account leptokurtic asset's returns. In case the assets have

[7] The model has been adapted for our purpose starting from the example presented in https://uk.mathworks.com/help/econ/examples/using-extreme-value-theory-and-copulas-to-evaluate-market-risk.html

[8] Glosten, L. R., Jagannathan R., Runkle, D. E., 1993, "On the Relation between Expected Value and the Volatility of the Nominal Excess Return on Stocks." The Journal of Finance, Vol. 48

an approximately normal kurtosis, normal residuals can also be an effective solution.

Residuals are the base series for the EVT estimation of tails CDF modelled as Generalised Pareto Distribution (GPD). CDF for every asset is then constructed using a Gaussian Kernel to smooth the estimation and clean it from patterns due to sampling limits. Anyway, Gaussian Kernel does not provide the best available fit for distribution's tail.

Our main concern is about the lower tail, which is defined as the worst 5% percentile of returns: the issue is to model this area in an effective way, adjusting for non-normality distributions. We chose that percentile to be consistent with the Mean-CVaR framework we are using to compute efficient allocations. Actually, we minimise CVaR 95% for a target return, modelling therefore the tail risk on the extreme 5%.

Residuals belonging to this area can be modelled effectively through a Generalised Pareto distribution. Maximum likelihood approach allows us to estimate the parameters (tail index and scale parameter) for a Generalised Pareto distribution of exceedances.

The CDF of a GPD requires the estimation of the parameters in the following equation

$$F(y) = 1-(1 + \frac{\varsigma y}{\beta})^{-1/\varsigma}$$

Where ς is the tail index parameter larger than -0.5, y are the exceedances (positive by construction) and $\beta>0$ is the scale parameter.

We provide therefore a picture about the estimates of the lower tail to give an intuitive view of the approach used to model the most concerning region of the distribution.

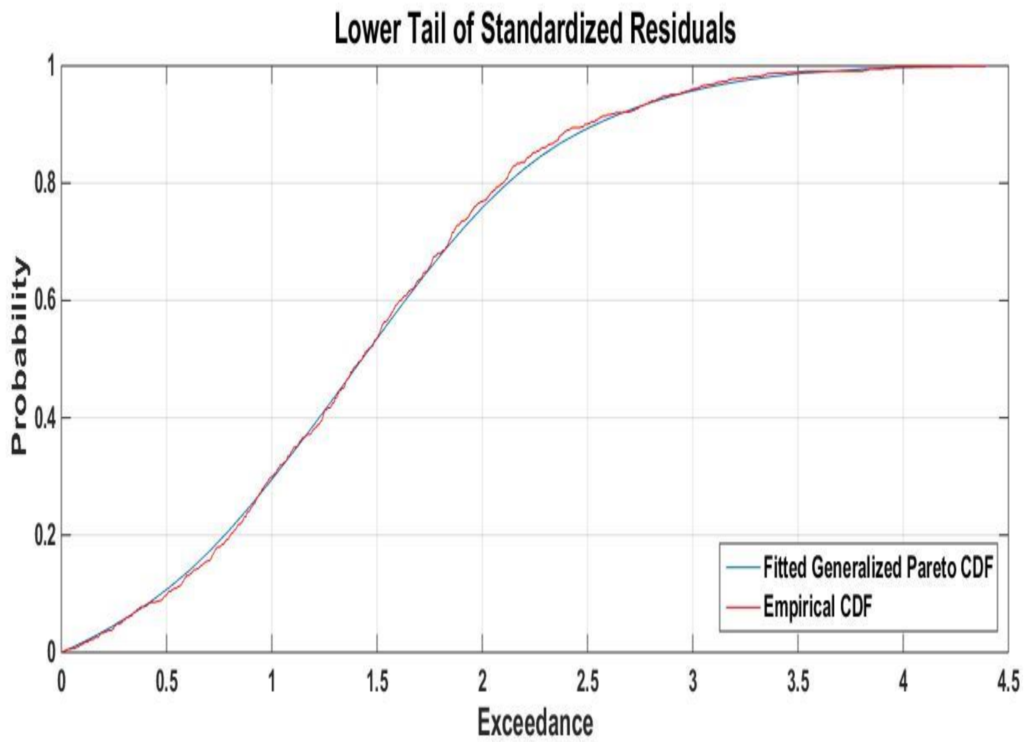

Figure 1 GPD fitting for the CDF of upper tail exceedances of residuals

As a result, it is possible to combine and interpolate the three sections of residual distribution:

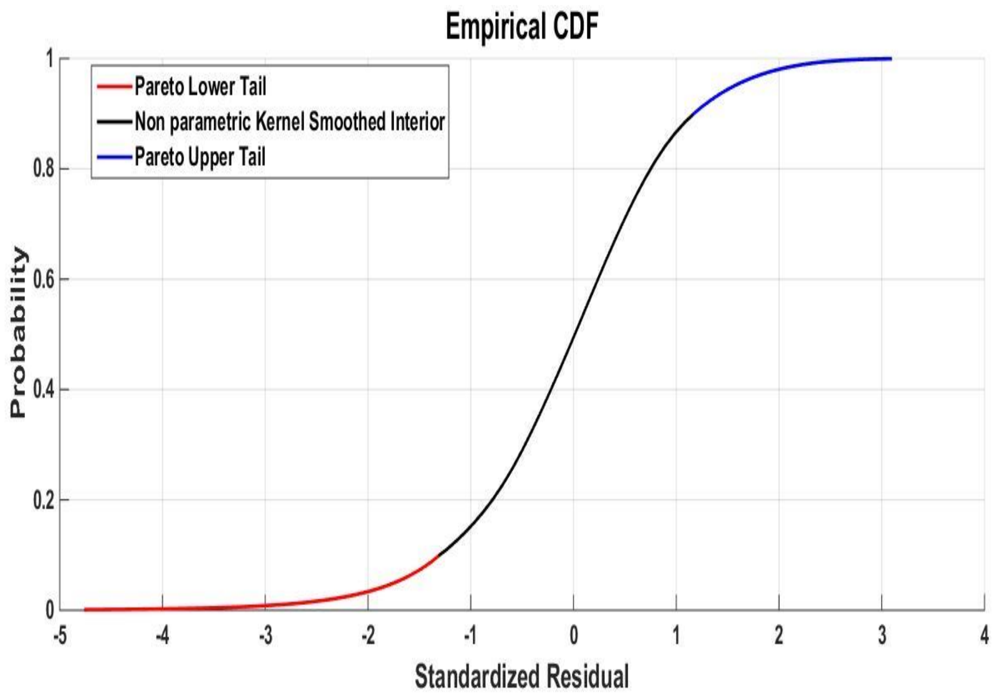

Figure 2 Empirical CDF of residuals

As a final step, we transform the standardized residuals to uniform variables using the semi-parametric empirical CDF just derived, and then we fit the t-copula to the transformed data (canonical maximum likelihood). Student-t copula imports correlation between the simulated residuals of each asset. Other elliptical copulas are available but it is largely recognised that t-copula is efficient for the low number of parameters to be estimated while, at the same time, it fits better than a Gaussian copula in modelling kurtosis.

Copula fit is here performed through an approximate MLE approach for large samples. To summarise the estimation process, the fitting of the copula first constructs the log-likelihood function; it differentiates the function with respect to the linear correlation matrix, which is one of the parameters to be estimated: these iterations are nestled in the maximization process to converge into the estimation of degrees of freedom of the Student-t distribution. In this way you will not reach a convergence of the estimated correlation matrix to the classic conditional MLE parameter but for large samples the approximation is efficient in that takes much less computations to be addressed[9].

Employing parameters from Student-t, it is now possible to simulate asset returns: first, dependent uniform residuals are simulated by inverting the semi-parametric CDF. Output of the simulation is not dependent on time but dependent at any point in time, and most important consistent with the econometric model estimated on the first step. These simulated residuals are filtered in a way to re-establish the autocorrelation and the heteroschedasticity "cleaned" by the ARIMA process. They represent the standard noise process embedding the rank correlation induced by the copula to create returns.

[9] Durrleman V., Nikeghbali, A., Roncalli, T., 2000, "Copulas Approximation and New Families", Groupe de Recherche Operationnelle Credit Lyonnais Paris

CHAPTER 3: Model at work

The best model to explain the world is the world itself.

Reliable models are a proficient way to handle efficiently the complexity of the world. Most famous models are nothing else than a good approximation but as we have seen the effects of extreme crashes and market shocks, investors have started to lower down their confidence in well-established models after having suffered losses. Skewed and leptokurtic market cannot be ignored anymore.

Asset management industry needs to cope with market anomalies. Recent investment philosophies have developed successful strategies (and marketing) minimising losses and cutting tail risk.

In this application of the model, we propose a very practical perspective reflecting the real needs of asset management industry. The main question to be answered is whether it is profitable to input views in the asset allocation process as a general matter. Related to that, we should know how bad it is going to perform in case the view is definitely wrong. Of course, this version of Black-Litterman under non-normality assumption assesses an efficient allocation conditioned on the view for the following unit period, or, in other words, the model is calibrated for one period. As a result, it seems to be reasonable to wonder what happens if the view turns out to be correct in few periods. As a bottom line, we will try to assess empirically an efficient "rule" to calibrate the level of confidence for the view.

Talking about the views to be inserted we will give way to relative views for our practical purpose's sake. Market anomalies occur especially in distressed situations and do not follow closely fundamental analysis for asset pricing; moreover beta strategies are biased because of increased cross correlation among assets during volatility spirals; stock picking is too risky and unstable to even be considered. Cross assets (or cross sector) expectations of over/under performance seem to be the only way portfolio managers can express views. Relative views can then fit better the needs of practitioners.

Our metrics to evaluate and to rank performances is a modified version of Sharpe ratio. Just like the traditional indicator, we compute the amount of return for every unit of risk: what changes is the measure for risk, namely it is not the standard deviation of the returns but the Conditional VaR to be consistent with the asset allocation approach. The Conditional VaR is computed with 95% of confidence level.

Five hypothetical assets compose our portfolio. Returns are simulated and resampled.

In the next paragraphs, we are optimising the allocation conditioned on the view in a 5 assets portfolio and, after having obtained the optimal weights, we are going to simulate risk and returns in different horizons splitting the results on whether the view is correct or not.

In this first section of analysis, we are going to compare performances of portfolios embedding or not a view: these two efficient allocations would then be assessed for their performances on different scenarios, namely where the view is correct and where it is not. The simulation will be run according to the Monte Carlo method adjusted for EVT. Occurrence (or not) of the view is obtained by setting the distribution in order to have on average the realisation (or not) of the view for every single period in one case or for distribution of cumulated returns until a certain horizon in a second case.

Data present daily returns for the five assets and the maximum horizon to cumulate are 22 periods, one effective month.

3.1 To put or not to put the view

Before starting the empirical analysis we want to state clearly that we are going to present statistical results based on views likely to be achieved by the asset in portfolio for certain time terms and consistent with the input data.

Another threshold for reasonability of the view concerns counter-productivity: we are going to exclude a view which would decrease sharply the expected return of the portfolio or enlarge the risk measure or a combination of both, even though the view reveals to be correct on market realizations. A portfolio manager would actually figure out how the view would alter the asset allocation and portfolio performance and it would not make any sense to input a view that underperforms the neutral portfolio even when the view is correct.

The core purpose of this analysis is to measure how bad could it be if the view is not correct or, in other word, if it would be worth to insert it in the allocation process.

The first view states an over-performance of the fourth asset over the first one. We first compute the optimal allocation in case no view is introduced. As a second step, we optimise the posterior to obtain the optimum portfolio conditioned on the view. In the first part of the simulation we manipulate scenarios in order to have a distribution for every single period that on average makes the view occurring. Before implementing the views, we present in the following table the four moments associated with the assets' returns.

	Asset1	Asset2	Asset3	Asset4	Asset5
Mean	0.09%	0.06%	0.03%	0.04%	0.09%
Variance	0.029%	0.024%	0.020%	0.009%	0.022%
Skewness	-0.2913	-0.32173	-0.13961	-0.26036	0.379434
Kurtosis	7.51315	6.027926	5.304139	6.035059	11.77309

TAB 1 Univariate moments of prior distribution

We now implement a relative view claiming that Asset4 will over-perform Asset5 by 0.07% in one period with a confidence level equal to 50%.

We display *una tantum* the four moments related to the posterior for few simple considerations:

	Asset1	Asset2	Asset3	Asset4	Asset5
Mean	0.086%	0.048%	0.026%	0.056%	0.090%
Variance	0.029%	0.017%	0.020%	0.009%	0.022%
Skewness	-0.291	-0.383	-0.140	-0.307	0.379
Kurtosis	7.513	6.955	5.304	5.163	11.773

TAB 2 Univariate moments of posterior distribution

Only average returns have been "touched" by the posterior in that the view concerns the first moment of the distribution.

We now display in a graph the adjusted Sharpe ratio produced by the portfolio allocation conditioned on the view. The efficiency of the return is computed on the simulation of cumulative returns previously described: the maximum horizon to cumulate is 22 periods. The first picture compares these metrics for portfolio with and without the view, when the view is correct for the distribution of every single period. The second picture expresses the same comparison but for scenarios where the view does not occur.

From now on, we define a wrong view calibrating the simulation in order to have the opposite outcome claimed by the view. In other words, let us assume a view saying that asset1 over-performs asset2 by 1%, the performance of the portfolio on wrong view would be tested on scenarios where asset2 outperforms asset1 by 1%

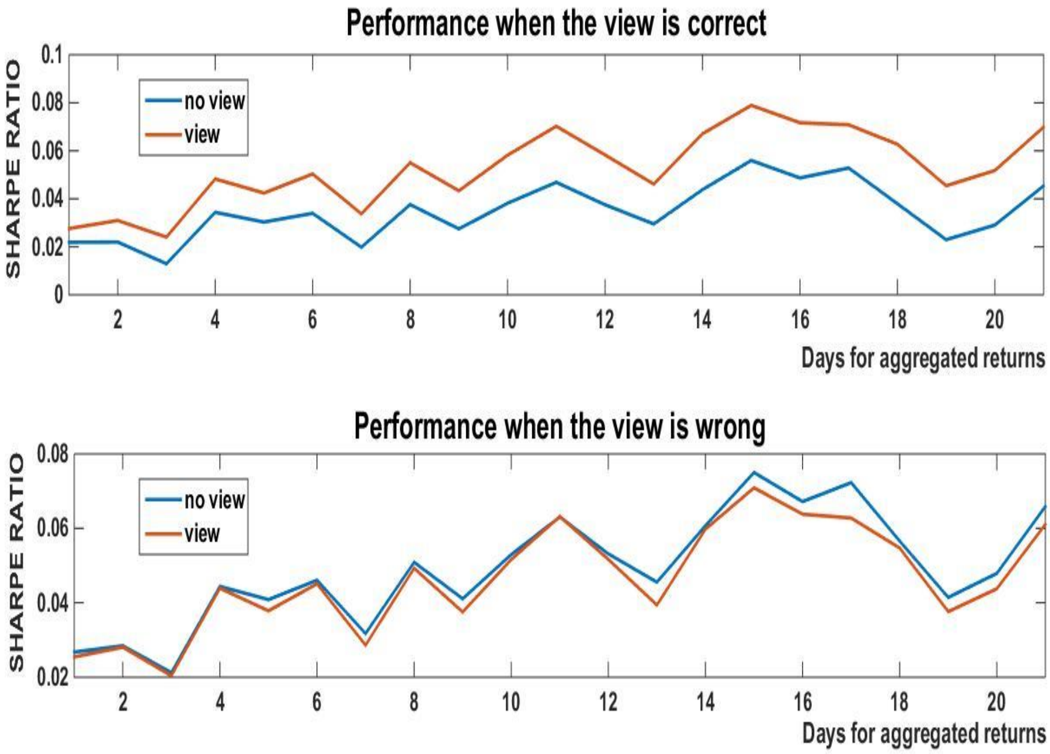

Figure 3 Comparing performances: portfolio with the view (view) vs. neutral allocation (no view). The view is correct (upper plot) or wrong (lower plot) on average in every single period. Shortselling is not allowed.

Portfolio with the view performs far better than the neutral one when the view is correct while it doe slightly worse when the opinion is found out to be a mistake.

It is reasonable to find out that the spread between the two portfolios is larger when the view is actually correct. However, the positive evidence is that when the manager's opinion is wrong, the conditioned portfolio performance lay not far at all from the performance recorded by the neutral portfolio.

The valuable feature of a CVaR portfolio optimization is the effective cut on tail risk. Once the efficient allocation conditioned on the view has been computed the weights for the portfolios allow the manager to keep the risk indicator CVaR almost at the same level, while the expected return undermines a shift downward that decreases the adjusted-Sharpe Ratio.

So far, we did not allow portfolio weights to be negative. Long/short strategies are widely used in asset management industry especially by sophisticated asset managers, which at the same time would need to keep risk below a certain threshold. In this second case study, we repeat the very same situation as before, but letting weights moving inside the range [-1; 1] for every single asset.

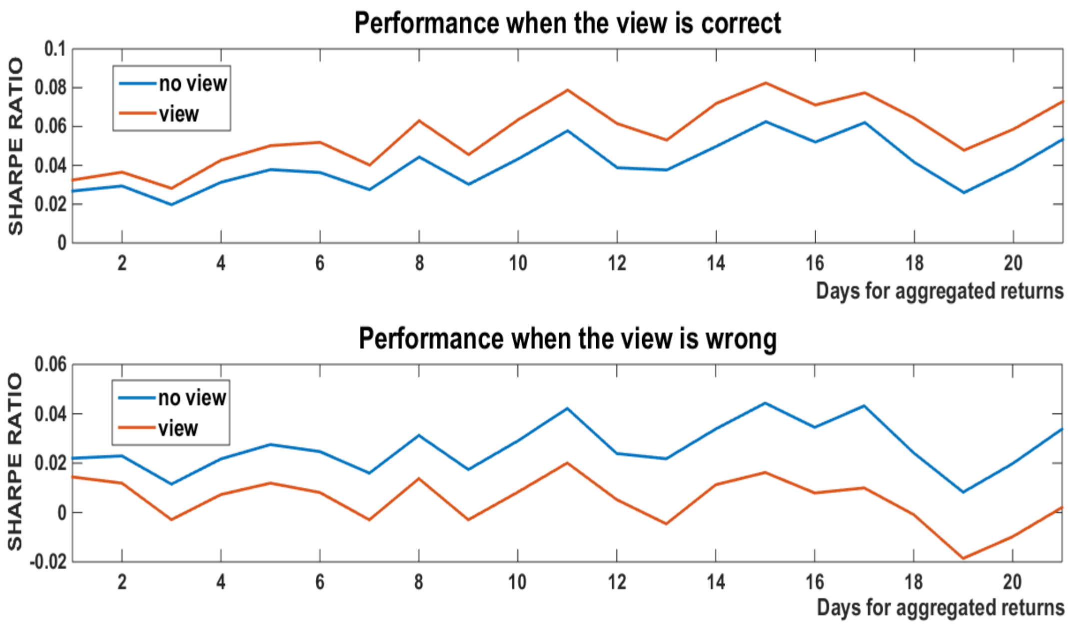

Figure 4 Comparing performances: portfolio with the view (view) vs. neutral allocation (no view). The view is correct (upper plot) or wrong (lower plot) on average in every single period. Shortselling is allowed.

First consideration reveals that implementing the view is convenient only when it is going to be correct: both expected return and CVaR worsen sharply when the view is wrong. Allowing short positions the portfolio is much more exposed to the occurrence of the view. The more the view states an underperformance of an asset, the more it brings the respective weight into deep negative territory the worse is the performance for incorrect views.

As a second step of the performance analysis through simulation, we want to introduce some "noise" around the realization of manager's view. Distribution of simulated returns is generated as before but the adjustment for the occurrence of the view (or not) happens only in the cumulated returns for different time horizons ranging from 1 to 7 periods after the asset allocation has been established.

This analysis is actually more useful and reasonable in practice for asset managers: they do not know exactly when their opinion could occur in the market even if they may hope to be right soon. Therefore we will not assume a distribution reflecting (or not) the view for every single period, but we are going to analyse the behaviour of our portfolios when the view occurs or not "cumulatively" within a certain amount of time. We create 7 cases when the view is correct (and 7 when it is not) for cumulated returns from 1 to 7 periods. The view implemented covers the same two assets but, given the longer term to be evaluated, the spread between performances is chosen at 20bps to appreciate a sensible change of performances in a longer term. First picture displays these results when short positions are not allowed

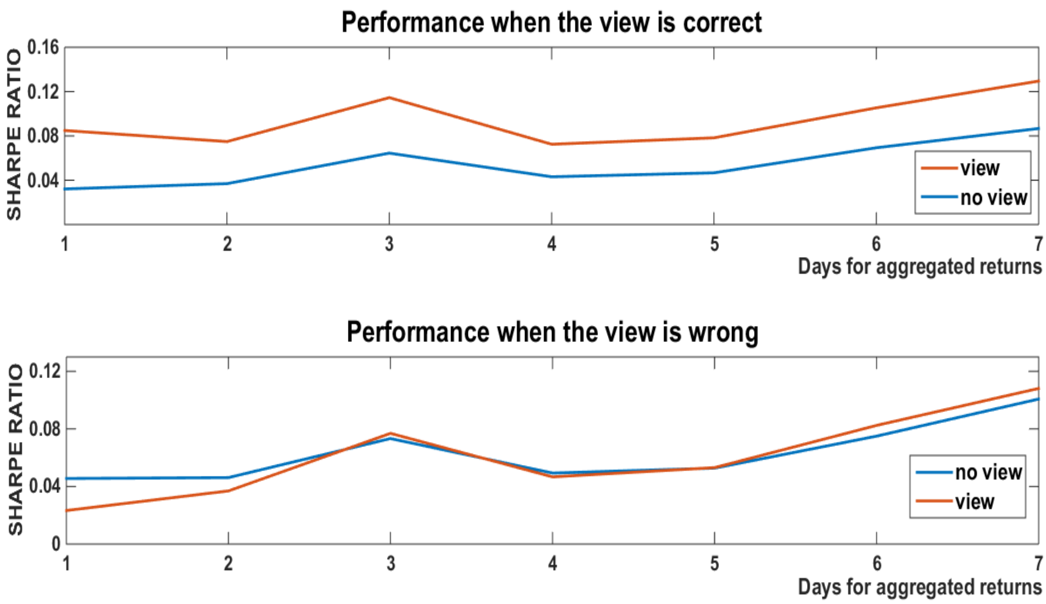

Figure 5 Comparing performances: portfolio with the view (view) vs. neutral allocation (no view). The view is correct (upper plot) or wrong (lower plot) on the aggregated return. Shortselling is not allowed.

Portfolio conditioned on the opinion performs better than the neutral one when the view is correct all the way through the time horizon. The behaviour is different when the view is wrong: Actually, the unconditioned portfolio shows better results in the early time horizons while it tends to be recovered by the allocation embedding the view in the longer term. It may seem reasonable that the mistake about the view worsen the adjusted Sharpe Ratio in the early term just because the allocation was suited for different market outcomes.

We have already noticed that, when short positions are allowed, the performance is deeply impacted by a mistake on the view. Evidence supporting this intuition is again detected.

The following picture displays the results when we allow negative weights on assets:

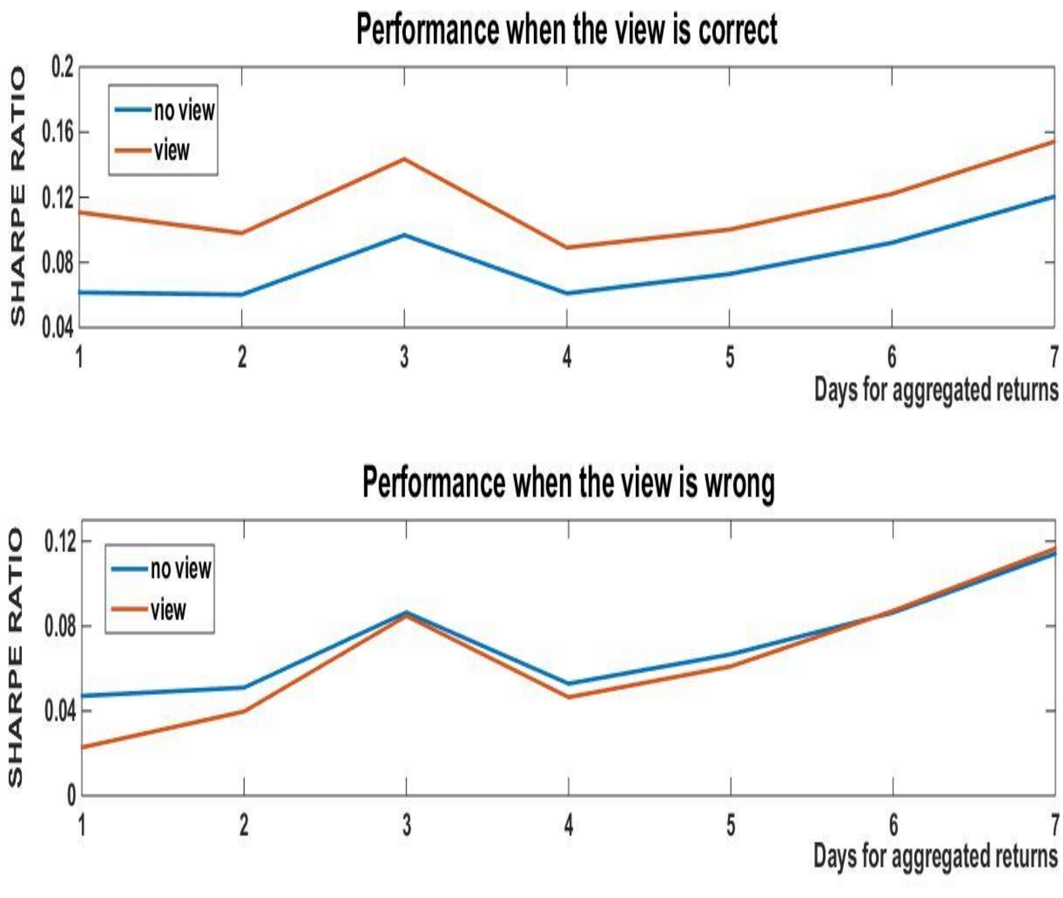

Figure 6 Comparing performances: portfolio with the view (view) vs. neutral allocation (no view). The view is correct (upper plot) or wrong (lower plot) on the aggregated return. Shortselling is allowed.

As already anticipated when market realizations do not match expectations stated by the view, the neutral portfolio outperforms the conditioned one. Like in the previous case with no shortselling, portfolio embedding the view starts to perform better after several periods even though the opinion continues to be wrong in the cumulative returns. Shortselling make the asset allocation conditioned on the view more focused and suited on the view itself. Adjusted Sharpe ratio is consistently higher for portfolios reflecting manager's opinion when shortselling is allowed and the view is correct.

Measure of risk expressed in CVaR is the key driver for better risk-adjusted returns in portfolio conditioned on the view. Either the view is right or wrong the adjusted Sharpe ratio is almost entirely determined by the different return associated with every scenario. Of course, a correct view improves the return. The interesting feature of portfolios with view is the unaltered CVaR measure whether the view is reflected by the market or not. As a matter of better diversification among assets, the view is likely to reshape efficiently the allocation minimising the CVaR for a target return. Very aggressive portfolios with weighty views and high target returns are still well diversified, while for the very same conditions the portfolio without the view collapses at the extreme of the efficient frontier concentrating 100% of the money in only one asset.

Portfolios without view have consistently higher CVaR measures along the time horizons, no matter if the view reveals to be right or wrong. The most important feature to be considered is that, at a given point in time, the CVaR measure is kept at the same level for different scenarios where the view is either correct or wrong.

3.2 Confidence level, performance and relative indicator function

The ground-breaking technique introduced by Black-Litterman motivated practitioners to focus on the view to introduce in their portfolio strategy. The formulation of the view needs to select a certain confidence level on the opinion. Few researches have tried to analyse this parameter aiming to disentangle a convenient procedure to assess an optimum confidence level.

During our analysis, we found out that, different confidence levels associated with the same view has very relevant effects on portfolio allocation and therefore, performances. The intuition behind this analysis concerns the non-normality of market returns: the more the prior is far from normality assumption, the more likely is the case to find portfolio allocations (still consistent with the prior) that over-performs on the risk adjusted return only by calibrating a certain confidence level.

First, we will show how the adjusted Sharpe Ratio is affected by different weights associated to different percentages of confidence levels, namely every (integer) percentage value ranging from 1% to 100%.

The simulation is held according to the second typology previously employed: the view is right or wrong within cumulated returns on 7 periods. The view is again asset4 over-performing asset5 by 20bps.

The following plot shows the performance of portfolios at different confidence level when the view is correct or it is wrong.

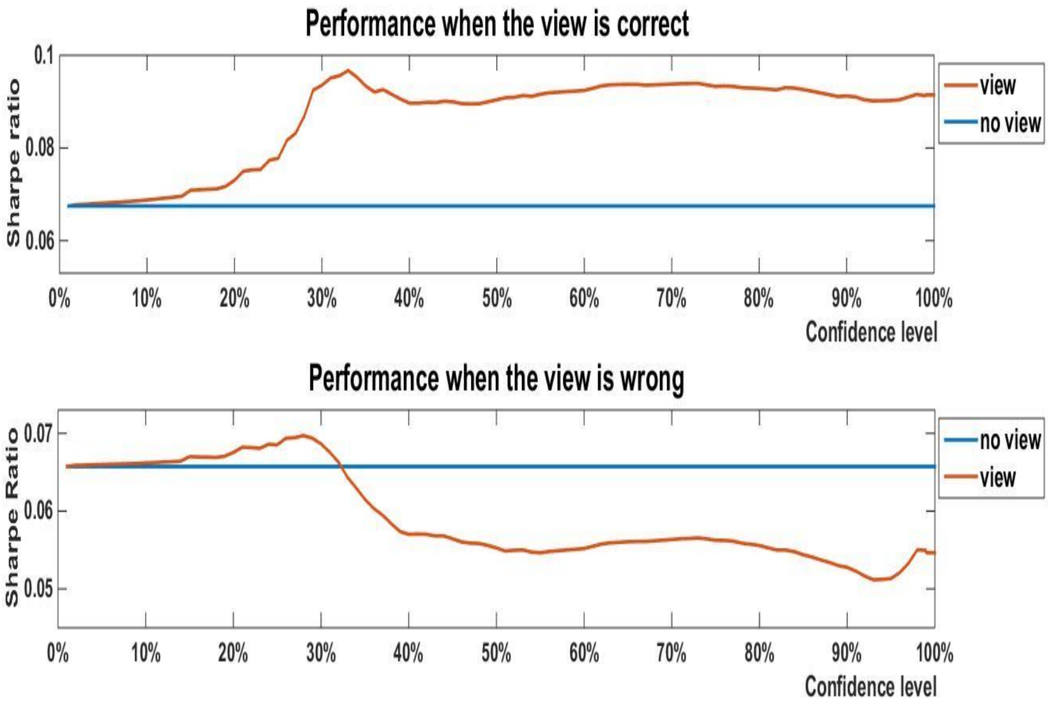

Figure 7 Comparing performances across the range of confidence levels. The view is correct or not on the aggregated return.

From the origin of the two curves, we see that the CVaR-based Sharpe ratio of the portfolio with view overperforms the unconditioned one for low percentage of confidence level. This is most likely due to the better diversification of portfolios embedding the extra information of the view, so better reacting to temporary shifts in parameters estimated for the optimization. Eventually increasing the confidence level of a wrong view will lead to an allocation conditioned on wrong information. In our example right over 30% of CL, the conditioned portfolio starts to underperform the neutral one.

In this example and in several others simulations we observed that the confidence level in which the conditioned portfolio reach the maximum performance for a correct view is strictly higher than the CL associated with the maximum performance achieved by the portfolio with view when the information is wrong:

$$\mathrm{CL}^*_{|\mathrm{CORRECT}} > \mathrm{CL}^*_{|\mathrm{WRONG}}$$

Maximisation of the performances according to this parameter may seem a too ambitious task. The following analysis aims to improve the efficiency in the choice of this "discretional" parameter. The intuition comes from studies to assess the confidence mirrored by the final allocation. In particular, a research from Idzorek (2005) formulated an easy way to drive out this object:

$$\mathrm{CL} = \frac{\widetilde{\mathrm{w}} - \mathrm{w}_{\mathrm{mkt}}}{\mathrm{w}_{100\%} - \mathrm{w}_{\mathrm{mkt}}}$$

Where $\widetilde{\mathrm{w}}$ is the current allocation observed in the portfolio, $\mathrm{w}_{\mathrm{mkt}}$ is the neutral allocation (portfolio with no view) and $\mathrm{w}_{100\%}$ is the allocation resulting by giving a 100% of confidence to the view. In this way, we can roughly assess at which distance from the two extremes (full neutral) the confidence really is. Actually, we want to employ this procedure to perform a sensitivity analysis based on an unambiguous value of how weights in portfolio behave for different CLs. This value will be called relative indicator. As a final step, we are making considerations on how different confidence levels affect multivariate skewness and kurtosis of the portfolio.

The following plot shows the relationships between portfolios' performances observed the previous paragraph and the values assumes by the relative indicator.

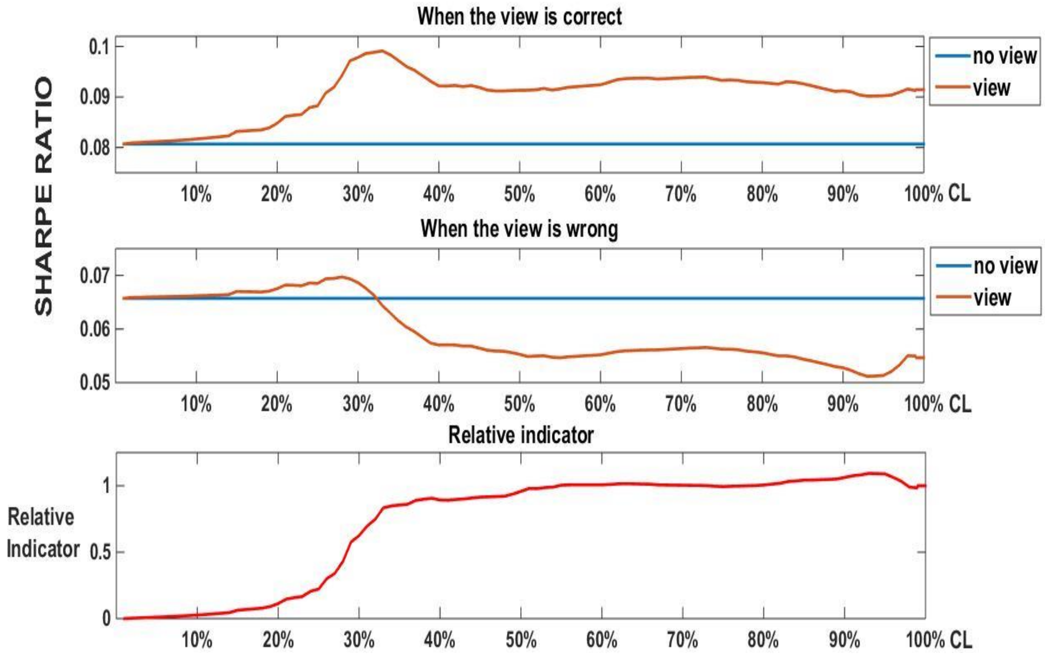

Figure 8 Portfolio performances across the range of confidence levels and Relative indicator (lower plot). The view is correct or not on the aggregated return.

To summarise the results in the first two plots we compare the performance already described above, with the third plot presenting the behaviour of the relative indicator.

We have noticed in this and in many other simulations that the performance tends to be higher and increasing for ranges of confidence level for which the relative indicator function has a cut-off on the slope, becoming flatter. As we will discuss further in the next paragraph, the market anomalies could push the relative indicator function to values above 1 for CL close to 100%.

Remembering that our simulated performance is measured for cumulated returns (for which the view reveals to be correct or not), we want to investigate how the "price" to be paid when you are wrong evolves along time, namely cumulating more and more returns for different confidence levels.

We will select two assets with the same expected return and implement a view claiming a spread between the two assets' average return of about 10bps.

The evidence from Monte Carlo simulation shows that for short time horizon (3 periods) in which to aggregate return, the performance of the conditioned portfolio when the view is correct improves as long as the CL is increased. Moreover, it overperforms the neutral portfolio. The opposite happens when the view is wrong: the neutral portfolio is consistently better than the conditioned one, especially for high confidence levels. On the other hand, increasing the time horizon we see that the performance of the portfolio with the view tends to converge to the same level no matter if the view is correct or not.

The following picture shows the details for a short horizon of 3 periods. Shortselling is not allowed.

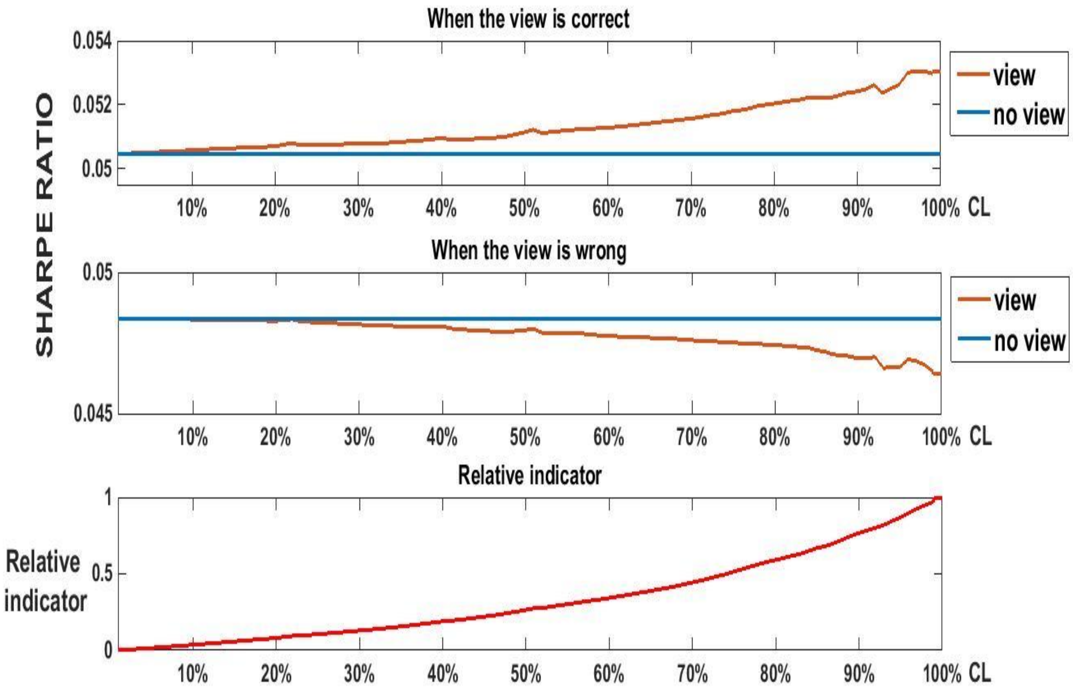

Figure 9 Portfolio performances across the range of confidence levels and Relative indicator (lower plot) in **short time horizon**. The view is correct (upper plot) or not (middle plot) on the aggregated return.

Clearly, in this short horizon a portfolio calibrated on a wrong view suffers and underperforms the neutral one. On the other hand, what it is interesting is that increasing the time horizon even if the view is wrong, the Black-Litterman portfolio achieves a better performance than the unconditioned one. The longer the time horizon, the more is convenient to insert a view in the allocation according to this model. In the following example, we represent the very same situation shown right above, but on a 10 period's time horizon.

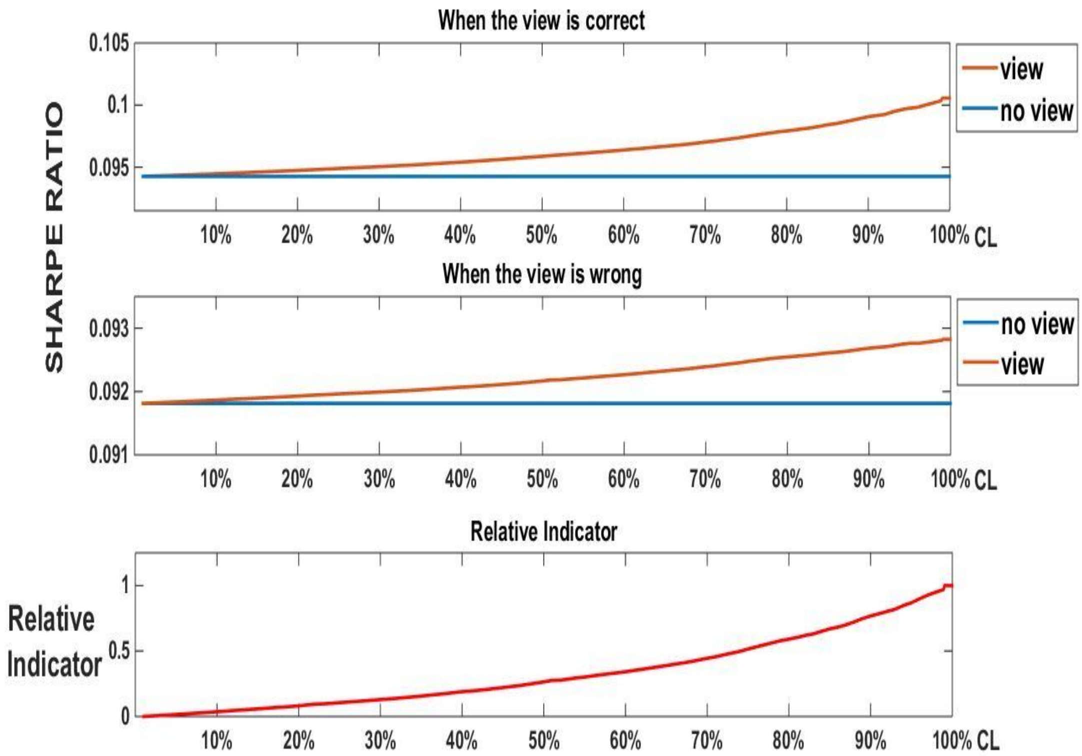

Figure 10 Portfolio performances across the range of confidence levels and Relative indicator (lower plot) in **long time horizon**. The view is correct (upper plot) or not (middle plot) on the aggregated return.

Even when the view is wrong, the conditioned portfolio performs better than the neutral one. When the view is wrong in the second graph, we see that the performance increases for higher confidence levels just right in the case of a correct view. The absolute value of the adjusted Sharpe Ratio is obviously higher for a correct view due to a consistently higher return.

The power of the model exploiting the properties of the CVaR minimisation is that it can control for the risk even when the market opinion is wrong. Moreover, as we have stated before, posterior distribution of asset returns with view lead to a better diversified allocation. Higher time horizons are able to blend the mistake in view's statement keeping the benefits of including more information in asset allocation.

In this precise simulation, the conditioned portfolio starts to over-perform the neutral one even for a wrong view for a time horizon equal to 6 periods.

In the case just presented, the relative indicator evolves as a monotonic increasing function, improving the intuitiveness of considerations. For higher confidence level, a wrong view is a real burden for the performance. For longer time horizons, the performance would be consistently better than the neutral portfolio's one: it tends to converge to the performance given by a correct view.

So far so good, we can add an extra consideration on concavity of relative indicator function. The relative indicator function presents a positive concavity all over its dominion and more important it is strictly increasing.

The performance ends up being improved increasing the confidence level on the view. It is consistent with the previous case: here the function does not know a cut-off in the slope and the performance is free to go up as far as the confidence level increases.

To support the previous statement, we want to analyse the performance and its relationship with the Relative Indicator function for higher time horizons: the performance when the view is wrong should not be any far from the risk-adjusted return achieved for a correct view scenario. In order to do that, the following picture shows what happens for a time horizon equal to 22 periods.

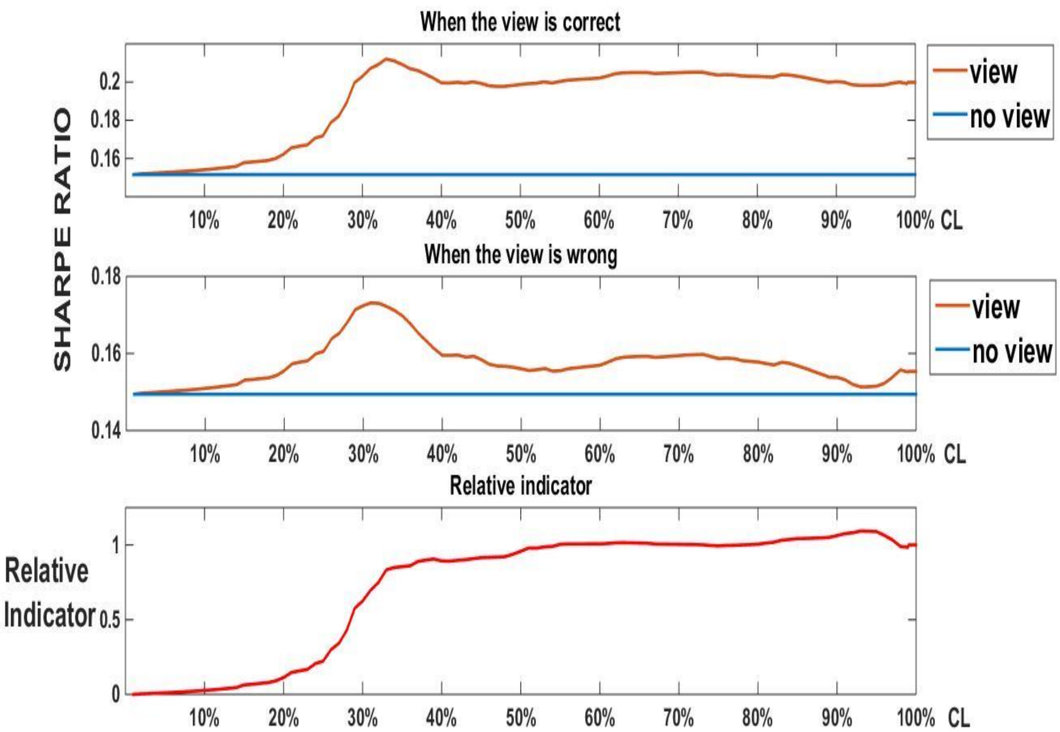

Figure 11 Portfolio performances across the range of confidence levels and Relative indicator (lower plot) in **long time horizon**. The view is correct (upper plot) or not (middle plot) on the aggregated return.

In 22 periods, time horizon the performance of the conditioned portfolio is strictly higher than the neutral one and again it is confirmed that the view can improve the risk-adjusted performance. The time you have to "wait" obviously depends on how wrong the opinion was, facing the following market realizations, and how "extreme" the view was: for instance stating that an underperforming asset's return would skyrocket, beating all the others assets. The information so far available in the market would blend the extreme view but obviously having set an overweight for the asset that continues to have a weak return, can end up being very painful.

The merit of this powerful feature to achieve better performances than the neutral portfolio is mainly due to the capability to limit the risk, especially for the threat of the lower tail of the distribution. In Appendix A**,** we see for a last example how the risk measure behaves and we provide some considerations about risk contribution of the assets covered by the view and its relationship with relative indicator function.

For the cases when the view is just not accurate, we have now the evidence to state that it would not make any sense to change the weights on the portfolio once you have realised that the view is wrong. At first the manager would regret for having missed the outlook (neutral portfolio is actually performing better at this point) but being patient and checking for cumulated returns (even though the view carry on being wrong), the conditioned portfolio would overperform the neutral one. Changing the asset allocation straight away would import heavy transaction costs plus the risk to be wrong one more time, keeping on regretting the passive market portfolio.

3.3 Risk contribution

The previous example has been analysed in appendix A performing a sensitivity analysis of risk measure CVaR. The result showed how the conditioned portfolio was much less risky than the neutral one, while no relevant changes related to the realisation of the view have been detected.

The following analysis on the same case study aims to investigate the risk contribution brought by the assets covered by the relative view of the case example of the appendix A. Risk contribution is usually measured in terms of VaR and it answers the question “ what percentage of portfolio risk (VaR) is due to a certain asset?”.

The contribution to VaR employs the *beta* of the asset with respect to the entire portfolio and it is defined as below:

$$Contribution\ to\ VaR = \beta_i w_i VaR_P$$

Where w_i is the weight of the asset, β_i is the beta of the asset and VaR_P is portfolio VaR

Again the contribution to VaR is tested for a complete range of confidence levels. In the following graph the related result.

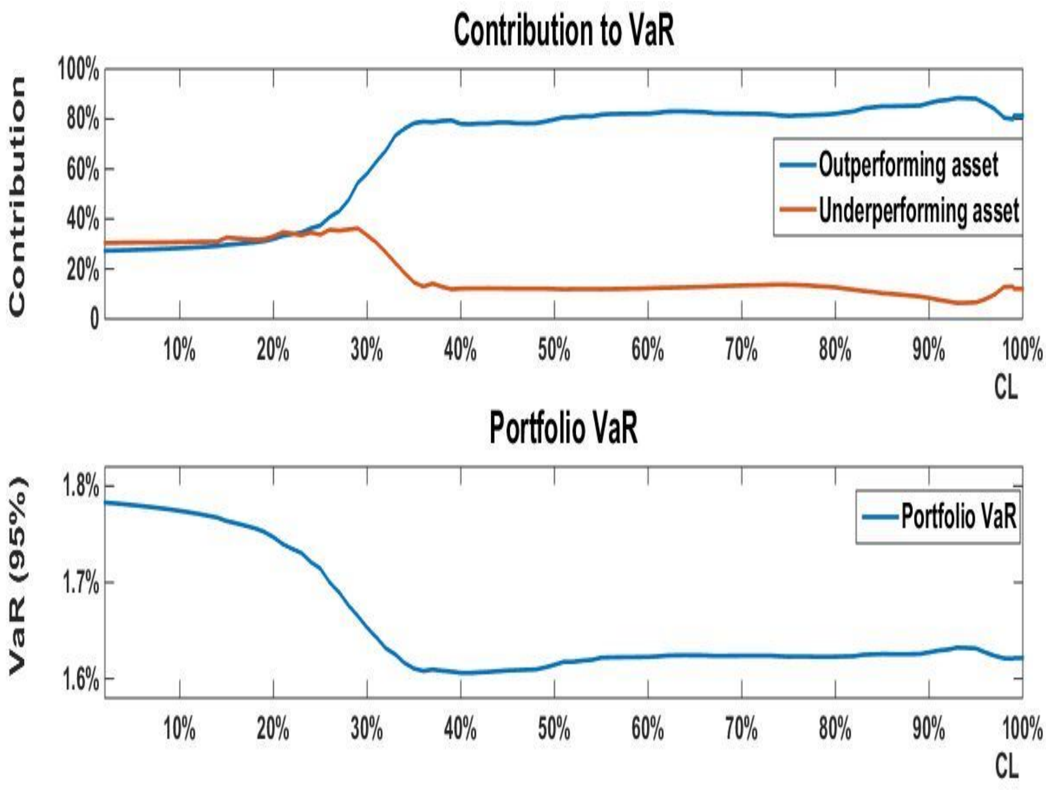

Figure 12 Contribution to VaR due to the two assets covered by the view (above plot). Portfolio VaR (95%) across the range of confidence levels (below plot).

Recalling appendix A, the portfolio VaR does not behave differently from the CVaR measure, increasing the confidence level on the view. Again, as the relative indicator function starts to be flat, VaR measure does not change relevantly for higher confidence levels.

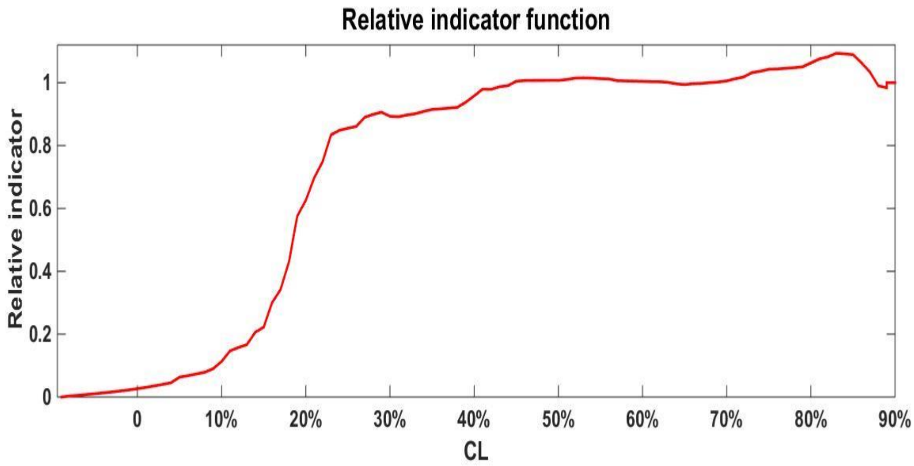

Contribution to risk related by the two assets follows similar patterns according to the confidence level. The asset, which is seen as outperforming in the relative view, will have an increasing weight as the confidence level approach 100%. The opposite would happen to the asset seen as underperforming, namely less weight and decreasing contribution to risk.

It is important to discuss the result according to how this contribution to risk changes. Again, as the relative indicator starts to flatten the variation of contribution to risk from the zero level of confidence has already reached its final level and remains almost unaltered until the 100% confidence. In other words, when this happens, the view has already saturated its “impact” in reshaping the portfolio.

The relative indicator function is again confirmed to be a synthetic indicator to evaluate the strength of the view’s impact on the portfolio. More technically the slope of the relative indicator function best describes the condition of the portfolio and signal that a certain level of confidence in the view of the portfolio would achieve better performances in terms of less risk, no matter how wrong the view is found out to be.

As a last consideration of this paragraph, we want to highlight that, when the relative indicator function starts to be flat the portfolio reaches the minimum level of risk (in terms of VaR and CVaR). As we have seen in the previous paragraph, in that spot, the risk-adjusted performance achieves the highest level, other things being equal about the realisation of the view.

3.4 Relative indicator and further moments

Relative indicator is function of portfolio weights and it is useful to disentangle how different confidence levels can actually affect the portfolio performances.

The powerful feature of this indicator we are going to discuss in this paragraph concerns its relationship with the adjusted Sharpe ratio, but more relevant with the multivariate skewness and kurtosis of the portfolio.

We can actually start from the multivariate third and fourth portfolio moments. We have already provided the formulas to compute these further moments and we will report again only the final formula for the relevant value to be estimated.

For the skewness we have

$$S_s = \mathrm{w}' M_3 (\mathrm{w} \otimes \mathrm{w})$$

while for the kurtosis, we have

$$S_K = \mathrm{w}' M_4 (\mathrm{w} \otimes \mathrm{w} \otimes \mathrm{w})$$

Where M_3 and M_4 are the co-skewness and the co-kurtosis matrices respectively.

We can now compute the kurtosis and the skewness for every efficient allocation produced by different CLs.

For practical reason we will present the analysis on the case seen before, where the relative indicator function was monotonically increasing; the graph is shown again below to easily make considerations.

The following plot shows how relative indicator function, portfolio skewness and portfolio kurtosis correlate for different levels of confidence associated to the view.

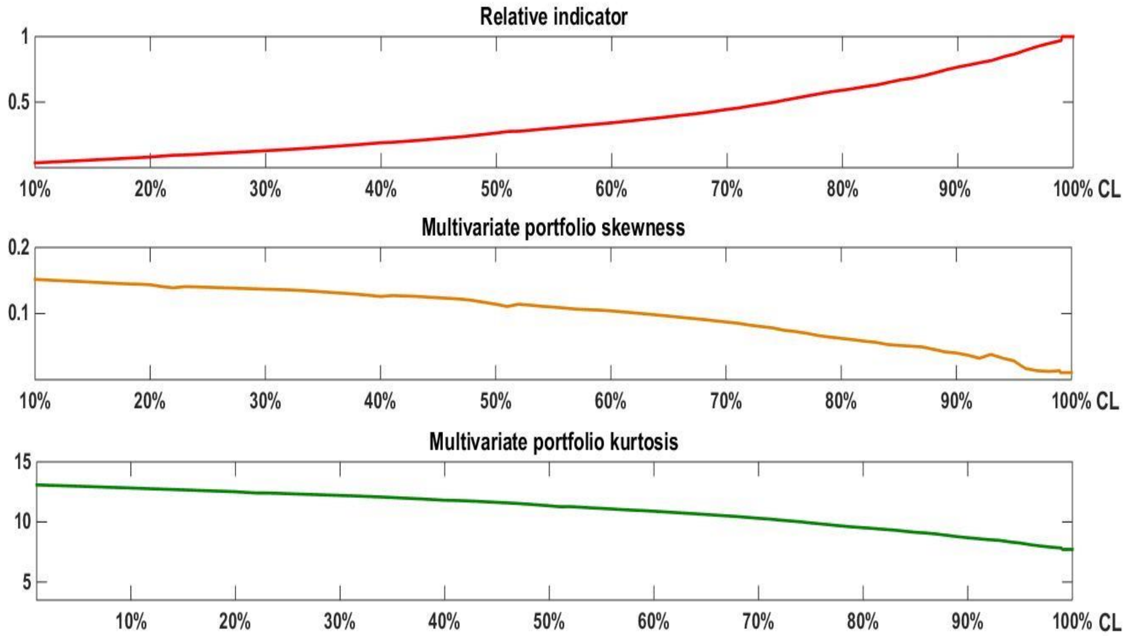

Figure 14 From above: Relative indicator. Multivariate portfolio skewness across the range of confidence levels. Multivariate portfolio kurtosis across the range of confidence levels.

Showing below the performances and relative indicator function described in the paragraphs above to allow an easier comparison :

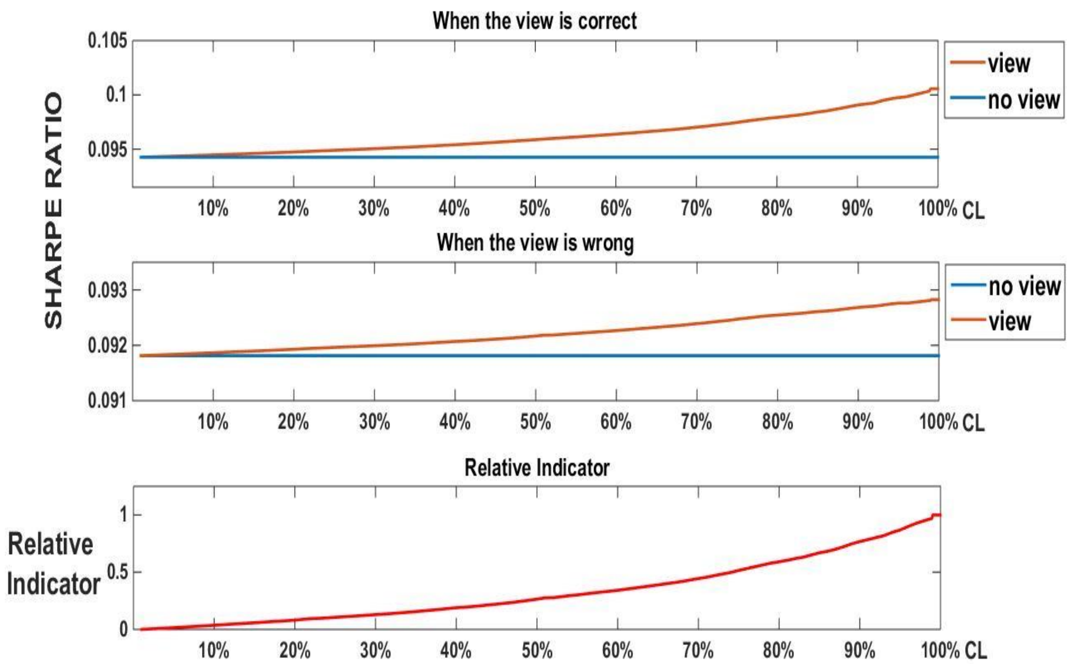

Figure 15 Portfolio performances across the range of confidence levels and Relative indicator (lower plot) in **long time horizon**. The view is correct (upper plot) or not (middle plot) on the aggregated return.

By increasing the confidence level, we see the third and the fourth moment reaching lower levels. This is partially not surprising given that when the CL is 100%, the posterior embeds the view itself, which is by construction uniformly distributed (so null skewness and excess kurtosis equal to -6/5).

We have seen that for long time horizon the risk-adjusted performance tends to be positively correlated with the relative indicator. The asset allocation summarised by the relative indicator has therefore a positive impact on portfolio performance subordinated to the level of kurtosis that the weights load to the portfolio.

The very last purpose of such a model is to make the asset allocation facing affectively the non-normality of asset's returns without losing in performance (not losing too much, at least).

The major concern for a portfolio manager is in this case the excess kurtosis, the main cause of sharp and sudden losses. Remembering the accuracy in modelling the Monte Carlo simulation employing Extreme Value Theory to have a brighter precision about what happens in the tails of the distribution, we can draw reasoned considerations on how extreme events affects the performance.

In the previous case study, we have dealt with strictly increasing relative indicator, positively correlated with the risk-adjusted performance and negatively with the portfolio kurtosis. This frame is just as things are supposed (and hoped) to be.

We propose again the plot showing performance when the view is correct, when it is wrong and the relative indicator function. The view is widely trend following: the manager decides to state an over performance of an asset over another while the very same pattern has been confirmed by historical evidence. This might clearly be a momentum strategy.

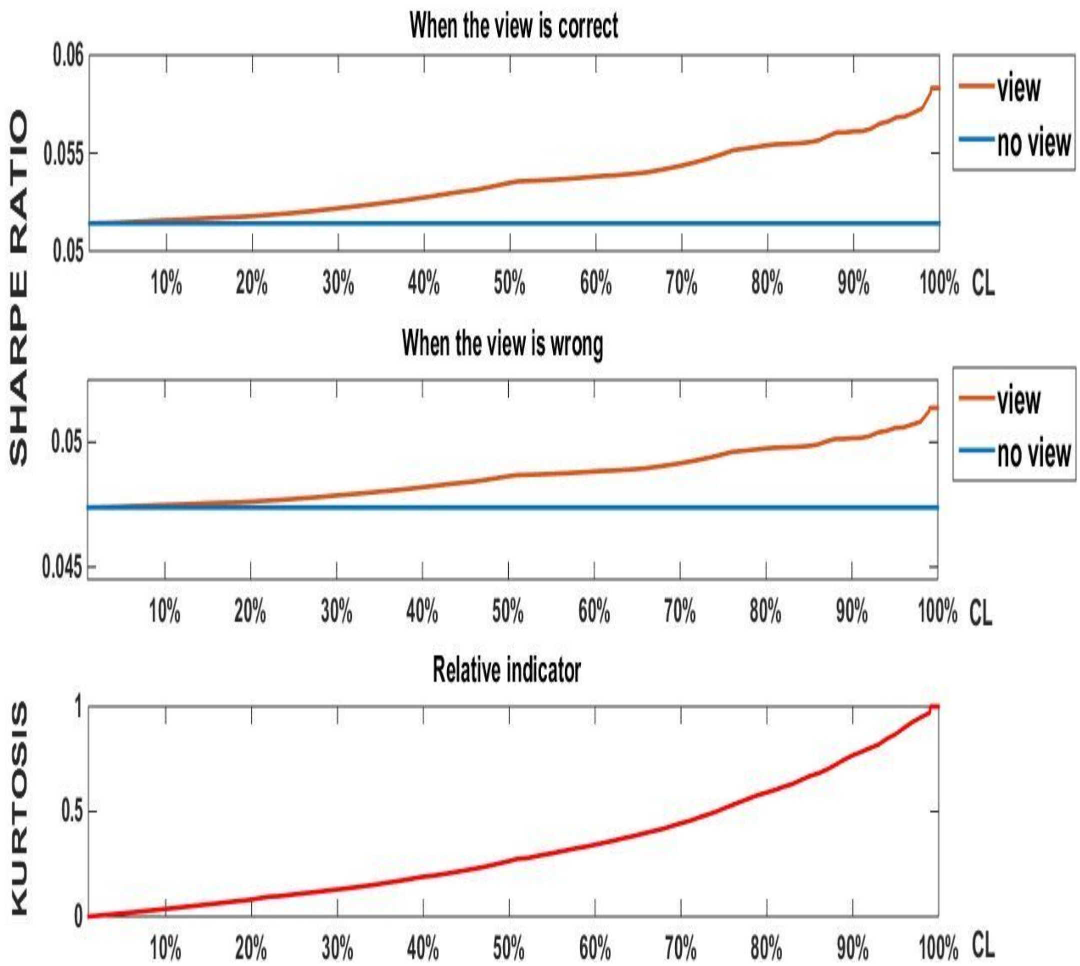

Figure 16 Portfolio performances across the range of confidence levels and Relative indicator (lower plot) in **long time horizon**. The view is correct (upper plot) or not (middle plot) on the aggregated return.

The relative indicator function looks less regular and intuitive: this is not particularly surprising given the non-normality of market data[10].

The simulation is run aggregating returns for a four period's horizon. Our main concern is the way the portfolio faces excess kurtosis. We have already described how the assets in the portfolio are leptokurtic. Here the problem is about calibrating the allocation in order to face the issue.

Before taking further considerations, we present the 2 graph of performances as before, but this time compared with portfolio kurtosis.

[10] Discontinuity points and non-strictly convex neighbourhood make less smooth the computation of an efficient frontier. We want to remember that in mean-CVaR framework the efficient frontier converges to the mean-variance frontier in case of normally distributed returns. Moreover sometimes discontinuity points and non-strictly convex neighbourhood make less smooth the computation of an efficient frontier resulting in a fragmented evolution of efficient weights along the way to 100% of confidence level.

We clearly see that the performance becomes positively related to the level of kurtosis.

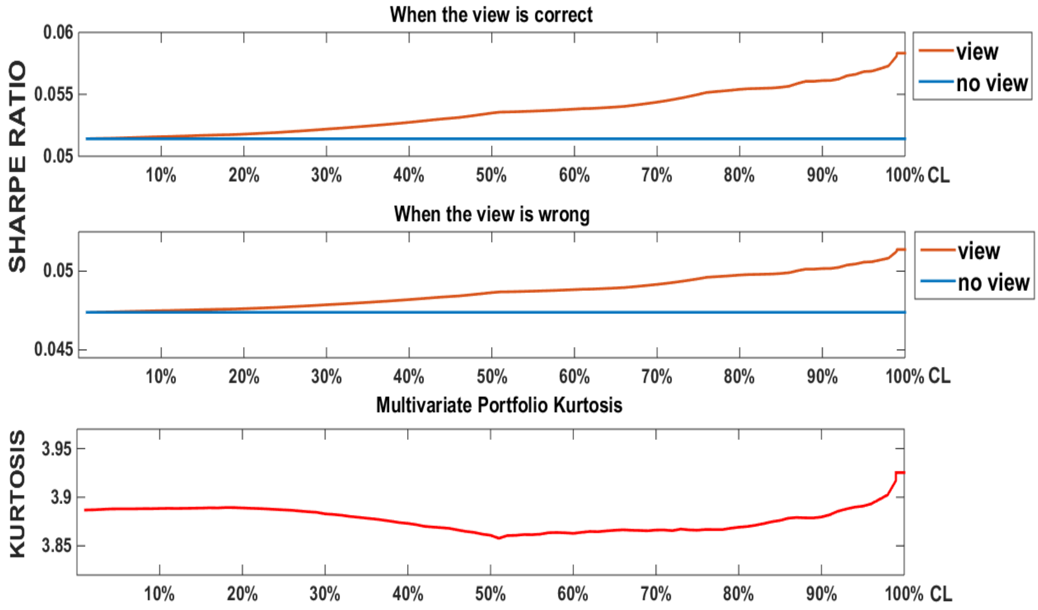

Figure 17 Portfolio performances across the confidence levels and related portfolio kurtosis (lower plot) in **long time horizon**. The view is correct (upper plot) or wrong (middle plot) for the aggregated return.

The momentum strategy benefits from the input of a coherent view. Whether or not the opinion finds its realisation, the conditioned portfolio performs better than the one without the view. The portfolio structure has not been reshaped by the view. It has just strengthened the momentum opportunity.

What we want to highlight is the low level of (excess) kurtosis. The asset allocation has found a way to limit the disrupt of extreme events in the tail's distribution. On top of that, we see a slight pick of kurtosis for very high levels of confidence. The optimised asset allocation starting from the posterior finds support on the market anomaly of excess kurtosis.

Conclusions

Active portfolio management is currently facing the threat posed by passive investment strategies. Actively managed funds generally ask for a higher Total Expense Ratio compared to a passive investor. This is why the disrupt of ETF product has led the balance of the competition to their advantage. What differentiates active portfolio management from the passive philosophy is the contribution from the portfolio manager: he/she would formulate investment opinions in order to achieve excess returns and the portfolio allocation would mirror those views. The question is whether it is convenient to insert the view or not.

The model that we have analysed, aims to process information from the market (passive side) and the managers' view (active) in order to reshape the market distribution of returns blending both sides. The advanced feature is to manage view implementation without assuming normally distributed asset returns. Actually, we do not assume any specific distribution.

Non-normally distributed returns are actually one the major challenges for investment management. Traditional portfolio theory assumes a Gaussian distribution to describe the market but the reality, especially during distressed periods has proved to be different. To who might follow a different convictions the risk of sharp drawdowns is high. A model that counts for this reality is actually necessary.

To answer the question on view's benefit we have tested the efficient allocation from the Mean-CVaR framework on not normally distributed simulated returns. CVaR optimisation and Monte Carlo simulations adjusted for Extreme Value Theory allowed us to analyse the non-normality of the market.

The main contestants are the neutral (no view) portfolio and the allocation conditioned on the view. The simulation is held in order to have the view occurring in the mean of the distribution of returns or for cumulative returns over a certain time horizon.

Conditioned portfolio delivers a better performance when the view is correct on average in every period. Anyway, due to an effective tail risk cut, CVaR measures do not differ considerably whether the view is correct or not. As a matter of facts, portfolios conditioned on the view end up being much better diversified. Moreover we proved how effectively the allocation from the model can face the threats hidden in the fat-ails of returns' distributions.

Time is a crucial component to answer our question. What if the market will not mirror the view along the way? Should we change the view? We can test the view checking for cumulated returns over a certain time horizon: if the manager was right, this would make the conditioned portfolio over-performing the neutral one by far, even for short time horizons. The brilliant feature to insert an investment opinion in our model relies on the fact that even if the view keeps on being wrong, the better construction of the conditioned portfolio would push the performance over the neutral's one. Time plays at our advantage. Changing the view would import heavy transaction costs and the risk to be wrong one more time. Better to wait passively for our mistake to be diluted.

In a certain way, it is the proof that active management can lead by avoiding transaction costs and by focusing on resilient and long term views.

How intensely the view has reshaped the asset allocation deserves some considerations too. Our relative indicator measures the impact of the view on the portfolio according to its confidence level. We have seen that this indicator increases at the same pace along the way to 100% confidence on the view. The point, in which the view has just exhausted its effect, might be the confidence level to achieve better performances.

Keeping in mind an intuitive way to asses a "convenient" confidence level, we have to know that patience and low transaction costs are what really pay out at the end.

Appendix A – Risk Levels

The analysis' result on whether to insert a view in the portfolio or not, clearly leant in favour of the portfolio conditioned on the view. The superior performance found the principal merit on the capability to limit the risk, especially when the view is actually wrong. We now see for the last example presented how the risk measure evolves by changing the confidence level, when the view is correct or not.

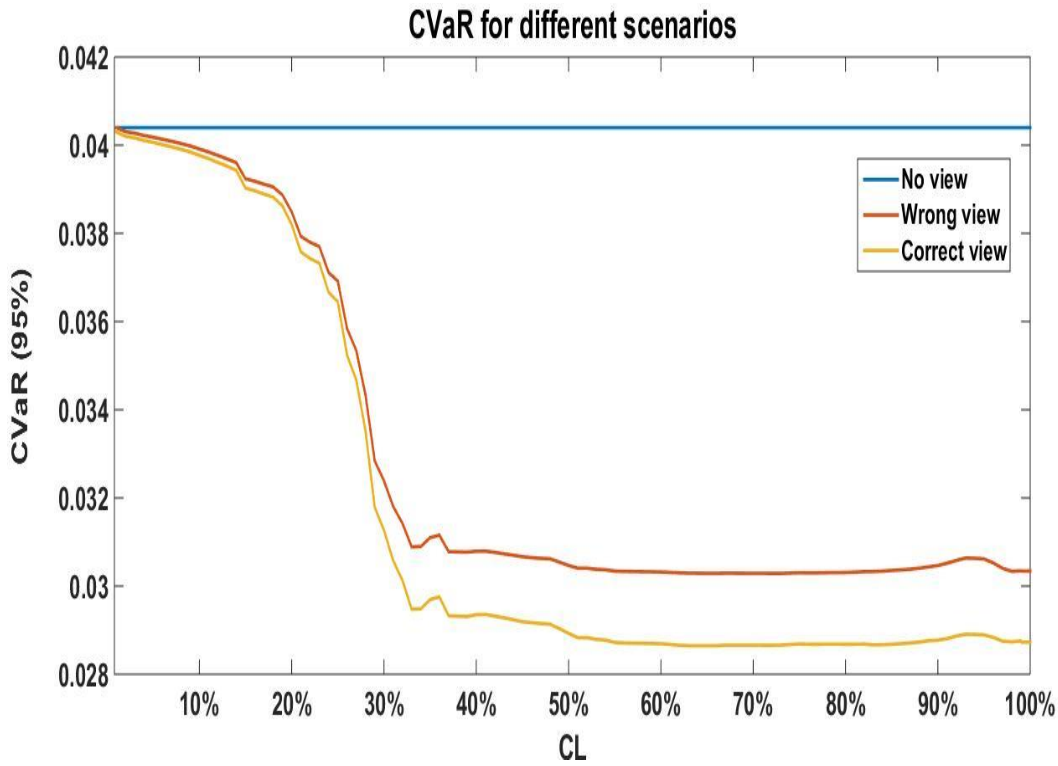

Figure 18 CVaR measures for neutral portfolio and for conditioned portfolio when the view is correct or it is wrong.

The graph clearly shows a huge gap between risk measures of the neutral portfolio (in blue) compared to the levels achieved by the conditioned ones, especially when the CL increases. First conclusion brings us to repeat that portfolios conditioned on the view can count on lower risk. Second, whether the view is correct or not the distance between the two risk measures is not so relevant, while for low confidence levels, it is actually null. Third, the evolutions of the risk measure follow inversely the relative indicator function. The slope of CVaR function finds a cut-off where also the relative indicator function starts to be flat.

Appendix B– MATLAB CODES

```
   % RELATIVE INDICATOR
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%

   %INPUT:
   % - WW is the matrix Nx101 of asset weights for the 100
different       confidence levels [0:1:100];

   %OUTPUT:
   % - IND = Relative indicator;

   [J,N]=size(WW);
   IND=zeros(J,1);

   for i=1:J
    ind= (WW(i,:)-WW(1,:))/(WW(101,:)-WW(1,:));
    IND(i:1)= ind
   end


% Multivariate Skewness and Kurtosis
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%
%INPUT:
   % - R = asset returns
%OUTPUT:
   % - SKW = Portfolio Skewness
   % - KURT = Portfolio Kurtosis

length = size(R,2)
T = size(R,1)
Co_Kurt=[];
for i=1:length
    for j=1:length
        kurtcok=[];
        for k=1:length
            for l=1:length
                kc=0;
                for t=1:T
                    kc=kc+((R(t,i)-mean(R(:,i)))*(R(t,j)-
mean(R(:,j)))*(R(t,k)-mean(R(:,k)))*(R(t,l)-mean(R(:,l))));
                end
                kurtcok(k,l)=kc/(T-3);
            end
        end
        Co_Kurt=[Co_Kurt kurtcok]; %Co_Kurt= Co-Kurtosis Matrix
    end
end
```

```
Co_Skew=[];
for i=1:length
    skewcoskew=[];
    for j=1:length
        for k=1:length
            sc=0;
            for t=1:T
                sc=sc+((R(t,i)-mean(R(:,i)))*(R(t,j)-
mean(R(:,j)))*(R(t,k)-mean(R(:,k))));
            end
            skewcoskew(j,k)=sc/(T-2);
        end
    end
    Co_Skew=[Co_Skew skewcoskew]; %Co_Skew= Co-Skewness Matrix
end


% Compute Portfolio kurtosis
    KRON = kron(kron(w',w'),w');
    KURT = w*Ku_Cok*KRON;

% Compute Portfolio skewness
    kRON = kron(w',w'); %
    SKW = w*Sk_CoS*kRON;  %
```

Codes to compute posterior distribution borrowed from Meucci (2006) and available on:

ARPM – ARPM website. 2016. ARPM – ARPM website. [ONLINE] Available at: http://www.arpm.co. [Last access 20 September 2016]

Code to simulate scenarios is inspired to the documentation open script: "Using Extreme Value Theory and Copulas to Evaluate Market Risk" available on:

MathWorks – Makers of MATLAB and Simulink. 2016. MathWorks – Makers of MATLAB and Simulink. [ONLINE] Available at: http://www.mathworks.com. [Last access 20 September 2016]

Bibliography

Acerbi, C., Szekelyy, B., 2014, "Backtesting Expected Shortfall"

Aparicio, F., Estrada, J., 1997, "Empirical distributions of stock returns: Scandinavian securities markets", Carlos III University, Madrid

Autchariyapanitkuly, K., Chanaimz, S., Sriboonchitta, S., 2014, "Portfolio optimization of stock returns in high-dimensions: A copula-based approach", Thai Journal of Mathematics

Bouye, E., Durrleman, V., Nikeghbali, A., Riboulet, G., and Roncalli, T., 2000, "Copulas for Finance A Reading Guide and Some Applications", Groupe de Recherche Operationnelle Credit Lyonnais Paris

Campenhout, J., Cover, T., 1981, "Maximum Entropy and Conditional Probability", IEEE TRANSACTIONS ON INFORMATION THEORY, VOL. IT-27

Cai, Z., Wang, X., 2008, "Nonparametric estimation of conditional VaR and expected shortfall", Journal of Econometrics, Volume 147, Issue 1, Pages 120–130

Caticha, A., Preuss, R., 2003, "Maximum entropy and Bayesian Data Analysis: Entropic priors", Phys.Rev.

Charpentier, A., Fermanian, J. D., Scaillet, O., 2006, "The Estimation of Copulas: Theory and Practice", Ensae-Crest and Katholieke Universiteit Leuven; BNP-Paribas and Crest, HEC Genève and Swiss Finance Institute

Cheng, S. X., 2008, "Nonparametric Estimation of Expected Shortfall", JOURNAL OF FINANCIAL ECONOMETRICS 6 (1): 87-107.

Durrleman V., Nikeghbali, A., Roncalli, T., 2000, "Copulas Approximation and New Families", Groupe de Recherche Operationnelle Credit Lyonnais Paris

Embrechts, P., Resnick, S., Samorodnitsky, P.,1999, "Extreme Value Theory as a Risk Management Tool", North American Actuarial Journal, Volume 3, Number 2

Fabozzi, F., Foccardi, S., Kolm, A., 2008, "Incorporating trading strategies in the Black-Litterman framework", The Journal of Trading 1(2)

Forghieri, S., 2014, "Portfolio Optimization using CVaR", LUISS Thesis

Glosten, L. R., Jagannathan R., Runkle, D. E., 1993, "On the Relation between Expected Value and the Volatility of the Nominal Excess Return on Stocks." The Journal of Finance, Vol. 48

Golan, A., 2008, “Information and entropy econometrics”, Department of Economics, American University

Jondeau, E., Rockinger, M., 1999, “The tail behavior of stock returns: Emerging versus mature markets” Mimeo, HEC and Banque de France

Jondeau, E., Rockinger, M., 2007, “Financial Modelling Under Non-Gaussian Distributions”

Konno, H., Yamazaki, H., 1992, “Mean-Absolute Deviation Portfolio Optimization and Its Applications to Tokyo Stock Market”, Management Science Vol. 37 No. 5, pp.519-531.

Idzorek, T., 2005, “A step-by-step guide to the black-litterman model:Incorporating user-specified confidence levels”, Ibbotson Associates

Ledoit, O., Wolf, M., 2003, “Improved Estimation of the Covariance Matrix of Stock Returns With an Application to Portfolio Selection”, Journal of Empirical Finance

Longin, F., M., 1996, “The Asymptotic Distribution of Extreme Stock Market Returns”, Journal of Business

Mashal, R., Zeevi, A., 2002. "Beyond Correlation: Extreme Co-movements between Financial Assets." Columbia University, New York

Martellini, L., V. Ziemann. 2007 “Extending Black-Litterman analysis beyond the mean-variance framework”, The Journal of Portfolio Management

McNeil, A. Frey, R., 2000 "Estimation of Tail Related Risk Measure for Heteroscedastic Financial Time Series: An Extreme Value Approach." Journal of Empirical Finance. Vol. 7, pp. 271–300

Meucci, A., 2006, “Beyond Black-Litterman: Views on Non-normal Markets”, Risk Magazine, February 2006, pp. 87-92

Meucci, A., 2006, “Beyond Black-Litterman in practice: a five-step recipe to input views on non-normal markets”, Risk Magazine, September 2006, pp. 114-119

Meucci, A., 2008, “Fully flexible views: theory and practice”, Risk 21(10) 97–102.

Meucci, A., D. Ardia, S. Keel., 2011, “Fully flexible views in multivariate normal markets”, from SSRN id:1742559

Neftci, S., N., 2000, “Value at risk calculations, extreme events, and tail estimation” Journal of Derivatives, pages 23–37.

Nystrom, K., Skoglund, J., 2002, "Univariate Extreme Value Theory, GARCH and Measures of Risk."

Nystrom, K., Skoglund, J., 2002, "A Framework for Scenario-Based Risk Management."

O'Toole, R., 2013, "The Black–Litterman model: A risk budgeting perspective", Journal of Asset Management, Volume 14, Issue 1, pp 2–13

Ortobelli, S., Biglova, A., Rachev, S., Stoyanov, S, 2010, "Portfolio selection based on a simulated Copula", Journal of Applied Functional Analysis , Vol. 5 Issue 2, p178-193

Pasquale, G., Marca, A., Deglialberti, C., Mignacca, D., 2012, "Dynamic Libility Driven Investing Using Black Litterman", from SSRN id:2297795

Patton, A., 2004, "On the Out-of-Sample Importance of Skewness and Asymmetric Dependence for Asset Allocation", JOURNAL OF FINANCIAL ECONOMETRICS

Rockafellar, R.T., S. Uryasev, 2000, "Optimization of conditional value-at-risk, Journal of Risk

Rockafellar, R.T., S. Uryasev. 2002, "Conditional Value-at-Risk for General Loss Distributions", Journal of Banking & Finance

Uryasev, S., Palmquist, Krokhmal, P., 2001, "Portfolio optimization with conditional Value-at-risk objective and constraints"

Villiers, J. D., 2013, "Non Linear Dependence Structures: A Copula Opinion Approach In Portfolio Optimization" , ESSEC Business School

Wang, I., Zhan Zheng, P. X., 2010, "Rebalancing, Conditional Value at Risk, and t-Copula in Asset Allocation", Duke University

Xu Yang, 2014, "Some Extensions of the Black-Litterman Model", Instituto Nacional de Matematica Pura e Aplicada

Webliography

ARPM – ARPM website. 2016. ARPM – ARPM website. [ONLINE] Available at: http://www.arpm.co. [Last access 20 September 2016]

Home :: SSRN. 2016. Home :: SSRN. [ONLINE] Available at: http://www.ssrn.com. [Last access 20 September 2016]

MathWorks – Makers of MATLAB and Simulink. 2016. MathWorks – Makers of MATLAB and Simulink. [ONLINE] Available at: http://www.mathworks.com. [Last access 20 September 2016]

Quant at Risk. 2016. Quant at Risk. [ONLINE] Available at: http://www.quantatrisk.com. [Last access 20 September 2016]

EXECUTIVE SUMMARY

The debate around passive and active investment management is going through revolutionary times. Active portfolio managers need to face an increasing competition from passive investment products made by ETF and other instrument replicating every sort of index or bucket of assets. The benefit brought by portfolio managers is their views about the market: manager's insight and projections on what could happen in the future are supposed to be correct enough to achieve higher returns.

There are many ways to express a view in practical terms, in order to affects the investment strategies. When it comes to managing a portfolio, the model employed to process the information from the market mixed and the input provided by the portfolio managers ends up being crucial.

This work analyses an advanced version of the traditional and widely used Black-Litterman Model: we actually want to extend the model, trying to remove the heavy assumption of dealing with normally distributed returns. Market turmoil has shown how unreasonable the assumption is and how painful it could it be to stick into inaccurate convictions.

The model is able to blend effectively the market neutral information with an exogenous opinion from the investment manager in a posterior distribution. The principal task is to face market anomalies, namely non-normally distributed returns. The algorithm borrowed from Meucci (2006) is able to process the information in a consistent way with respect to the non-normality (skewness and excess kurtosis of returns) without assuming any kind of underlying distributions on market data. Skewness and kurtosis of univariate returns distributions are left unaltered. When the view is introduced with a zero confidence the efficient allocation is again the market allocation. On the other hand, for full confidence the univariate distribution of the asset covered by the view will tend to a uniform distribution, namely the distribution of the view.

The availability of a reliable posterior distribution let us proceed to compute the efficient allocation. For a target return, the CVaR is minimised and it provides us a pure measure of downside-risk. At the same time Mean-CVaR framework would fulfil the need, coming from asset management industry, to cut tail risk on stressed market conditions.

The performance is tested and stressed on Monte Carlo simulation adjusted for Extreme Value Theory. A distinctive feature of this work concerns tail risk modelling and hedging. Monte Carlo simulation is achieved employing t-copula to fit the asset correlation. Every time series of returns is first modelled according to an asymmetric GARCH model. From the computed residuals, the CDF of returns distribution is inferred. Here the adjustment for tail risk is made through an EVT to model the tail distribution. The CDF of residuals is used to draw samples: the tails (5% percentile) of this sample distribution are therefore modelled as a Generalised Pareto distribution, while the remaining mass of the distribution is fitted through a Gaussian Kernel.

The following picture helps to visualize the result for the EVT process (figure 1):

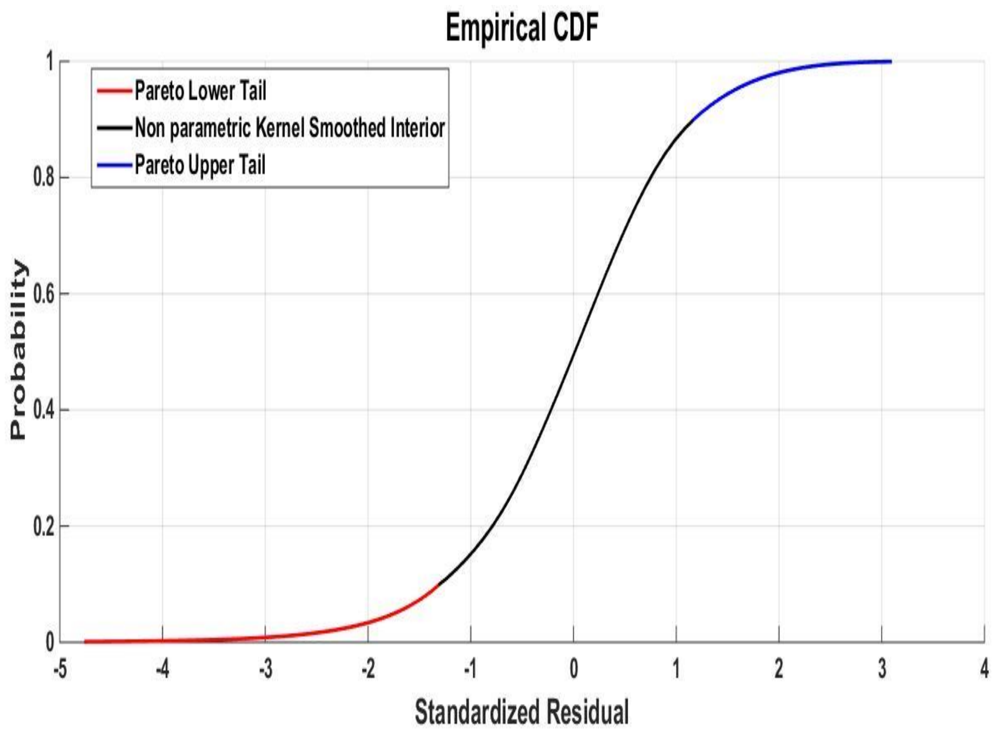

Figure 1 Empirical CDF of residuals

where the Generalised Pareto is defined by

$$F(y) = 1-(1+\frac{\varsigma y}{\beta})^{-1/\varsigma}$$

A multivariate t-copula will now fit this sample to import correlation among assets. Returns are simulated and asset allocation tested.

Testing portfolios conditioned on the view should assess whether it is convenient or not to input a view: more intuitively, how “painful” it would be to insert a wrong view. The Monte Carlo simulation adjusted for tail risk is going to be calibrated in order to replicate a correct or a wrong view, depending on the case. Being consistent with the Mean-CVaR framework, portfolio performance is assessed through a CVaR based Sharpe ratio. This modified version is the ratio of Mean expected return over CVaR.

We have conducted the performance analysis with daily returns according to two cases. For the first typology, the view is revealed to be correct or wrong on the distribution of the single market realization, checking also what happens on the cumulated returns.

Five hypothetical assets compose our portfolio; returns are simulated and resampled. We decided to insert a view on the portfolio. For practical reasons, our views are going to be relative: modelling non-normal markets with a focus on excess kurtosis is more likely to be a concern for sophisticated investors. Hedge funds widely use these trades, but also their returns to investors are frequently leptokurtic. Relative value trades based on a certain relative view are actually very common strategies.

According to this perspective, we found out that portfolios conditioned on the view without shortselling perform much better than the neutral one when the information is correct. On the other hand, the burden brought by a wrong view seems to be quite light. CVaR optimisation allows cutting the risk, namely CVaR does not change considerably whether the view is correct or not, the shift is due to lower returns. The following picture describes the risk-adjusted performance the statement above (figure 2).

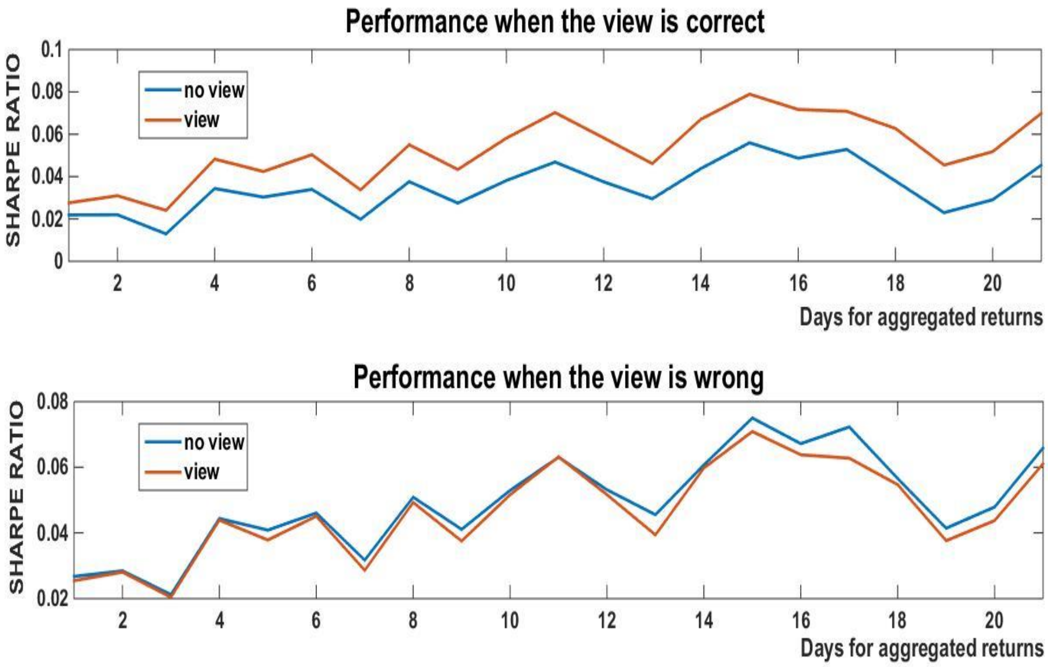

Figure 2 Comparing performances: portfolio with the view (view) vs. neutral allocation (no view). The view is correct (upper plot) or wrong (lower plot) on average in every single period. Shortselling is not allowed.

The story is different allowing for shortselling: the benefit brought by a correct view roughly equals (in the amount) the penalty carried by a wrong view. The decision to input a view is much more risky and a doubtful solution as shown by the following picture (figure 3).

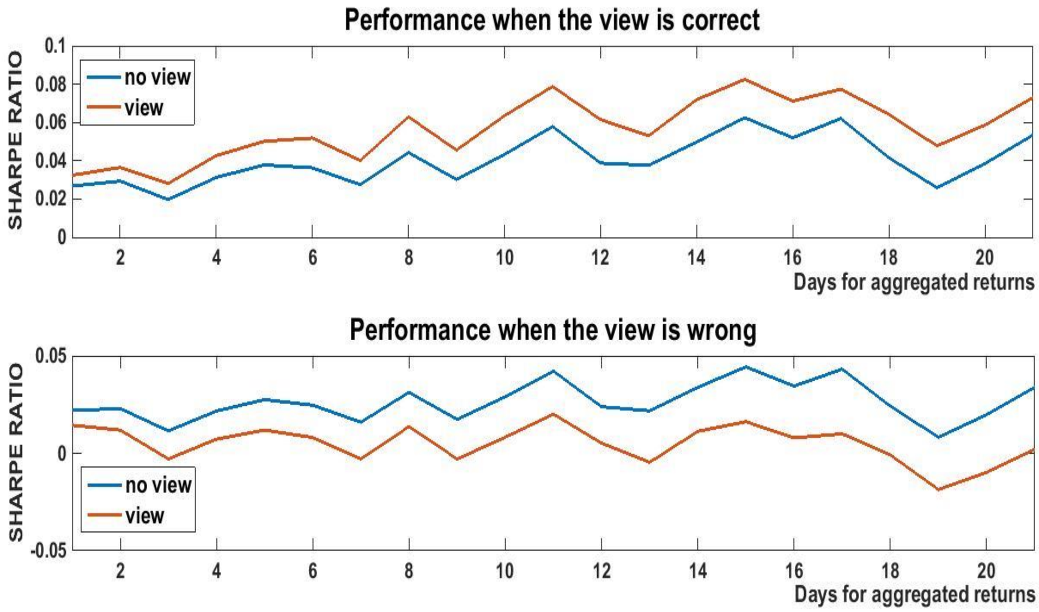

Figure 3 Comparing performances: portfolio with the view (view) vs. neutral allocation (no view). The view is correct (upper plot) or wrong (lower plot) on average in every single period. Shortselling is allowed.

The second kind of analysis strictly concerns the cumulative returns for different time horizon (mainly from 1 to 7 periods). Noise is added to the single period return distribution: the adjustment for the occurrence of the view (or not) happens only in the cumulated returns for different time horizons. This analysis is actually more useful and reasonable in practice for asset managers: they do not know exactly when their view could occur in the market, even if they may hope to be right soon. From this perspective we have time playing to our advantage. It is convenient to introduce a view, and even if it is wrong the portfolio would rapidly converge to the benchmark result drawn by the neutral portfolio (figure 4).

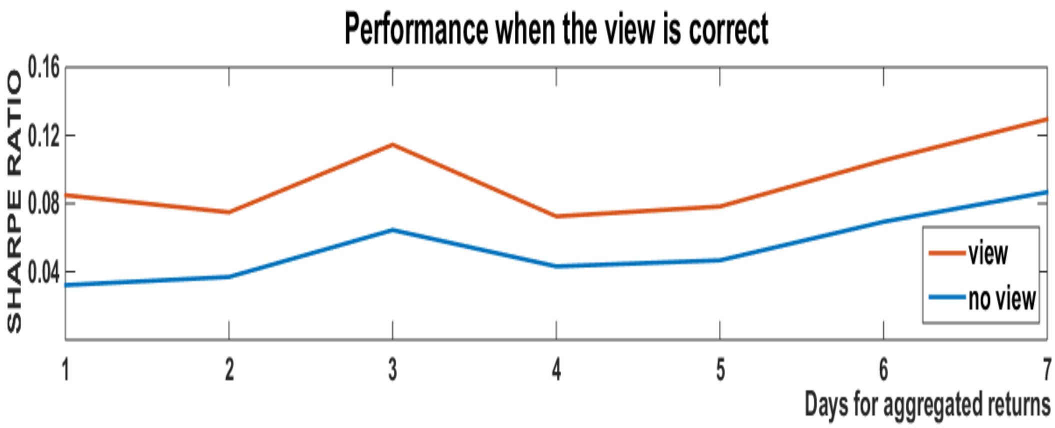

Figure 4 Performances by portfolio with the view (view) vs. neutral allocation (no view). The view is correct (upper plot) or wrong (lower plot) on the aggregated return. Shortselling is not allowed.

Shortselling affects only the time to recover and to converge to the neutral portfolio performance, making it slowly (figure 5).

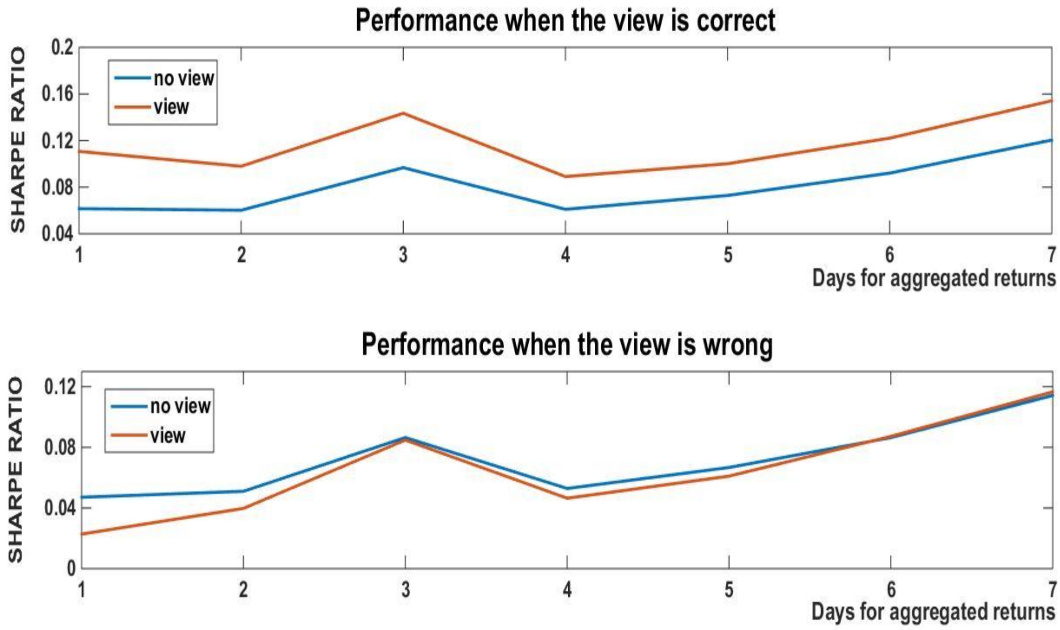

Figure 5 Performances by portfolio with the view (view) vs. neutral allocation (no view). The view is correct (upper plot) or wrong (lower plot) on the aggregated return. Shortselling is allowed.

Maintaining the focus on aggregated returns, we have analysed how the confidence level associated to a view affects the allocation and therefore the results. To have an unambiguous measure of the relative strength of a view, we have formulated a relative indicator defined as

$$\mathrm{CL} = \frac{\tilde{\mathrm{w}} - \mathrm{w}_{\mathrm{mkt}}}{\mathrm{w}_{100\%} - \mathrm{w}_{\mathrm{mkt}}}$$

Where $\tilde{\mathrm{w}}$ is the current allocation observed in the portfolio, $\mathrm{w}_{\mathrm{mkt}}$ is the neutral allocation (portfolio with no view) and $\mathrm{w}_{100\%}$ is the allocation resulting by giving a 100% of confidence to the view. In this way, we can roughly assess how strongly the view is affecting the asset allocation.

The relative indicator goes from 0 to 1, as the confidence level associated to the view increases. Describing this relative indicator as a function, we have noticed that the maximum performance is reached when the relative indicator function find its maximum slope (in absolute value). Usually the marginal strength of a view grows for higher confidence levels: the relative indicator function would be monotonically increasing with positive concavity. This is why the best performance is achieved for confidence equal to 100% when the view is correct (figure 6).

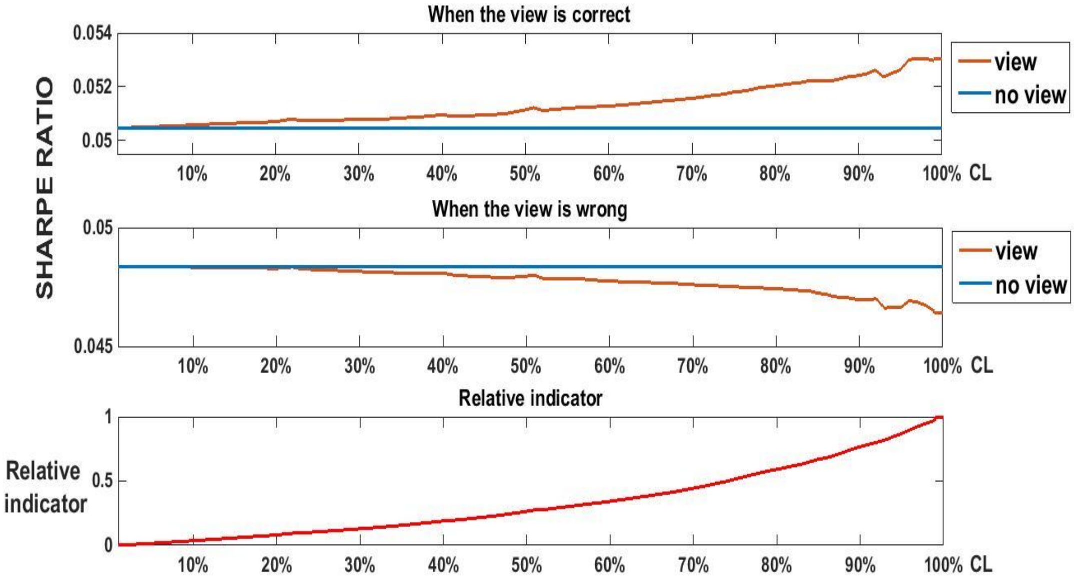

Figure 6 Portfolio performances across the range of confidence levels and Relative indicator (lower plot) in **short time horizon**. The view is correct (upper plot) or not (middle plot) on the aggregated return.

What surprises is that even when the view is wrong, the risk adjusted performance starts to show better results than the neutral portfolio simply enlarging the time horizon in which to cumulate returns. In other words, even if on average you are wrong, you had better waiting to have the mistake "diluted": inserting the view is again convenient compared to not doing so (figure 7).

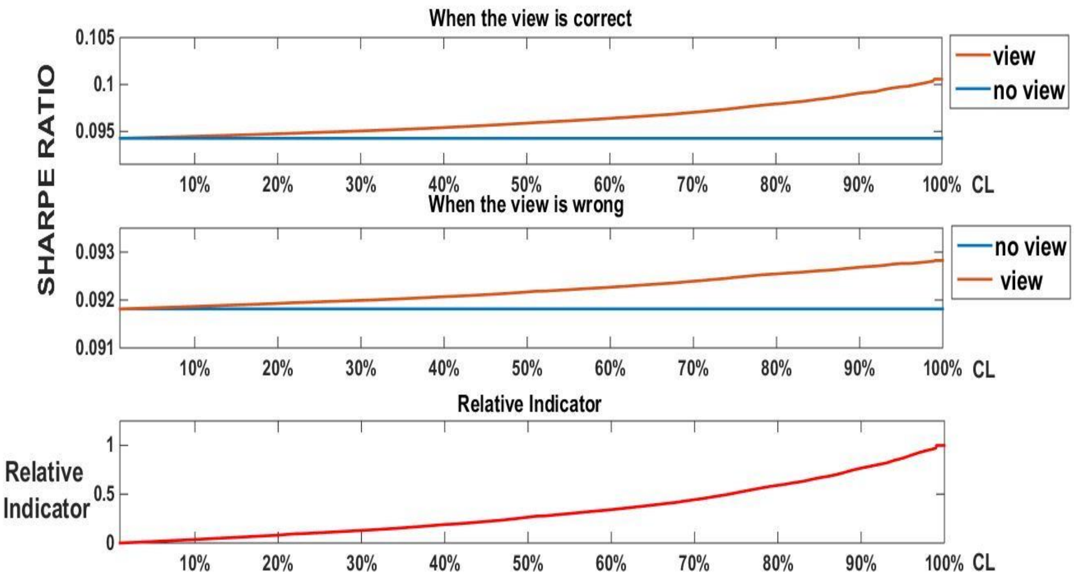

Figure 7 Portfolio performances across the range of confidence levels and Relative indicator (lower plot) in **long time horizon**. The view is correct (upper plot) or not (middle plot) on the aggregated return.

Some views from a portfolio manager can be extreme, namely very unlikely to happen given the asset's return distribution. In these cases, it is likely to deal with portfolios where the view has a very strong impact even for low confidence levels.

We remember that the relative indicator provides us a measure of how the portfolio allocation has been altered by the view at different confidence levels. An anomaly could be signalled by the fact that the relative indicator function reaches values very close to 1 for low confidence levels, 30% for instance. It is basically what happens to in the following asset allocation (figure 8):

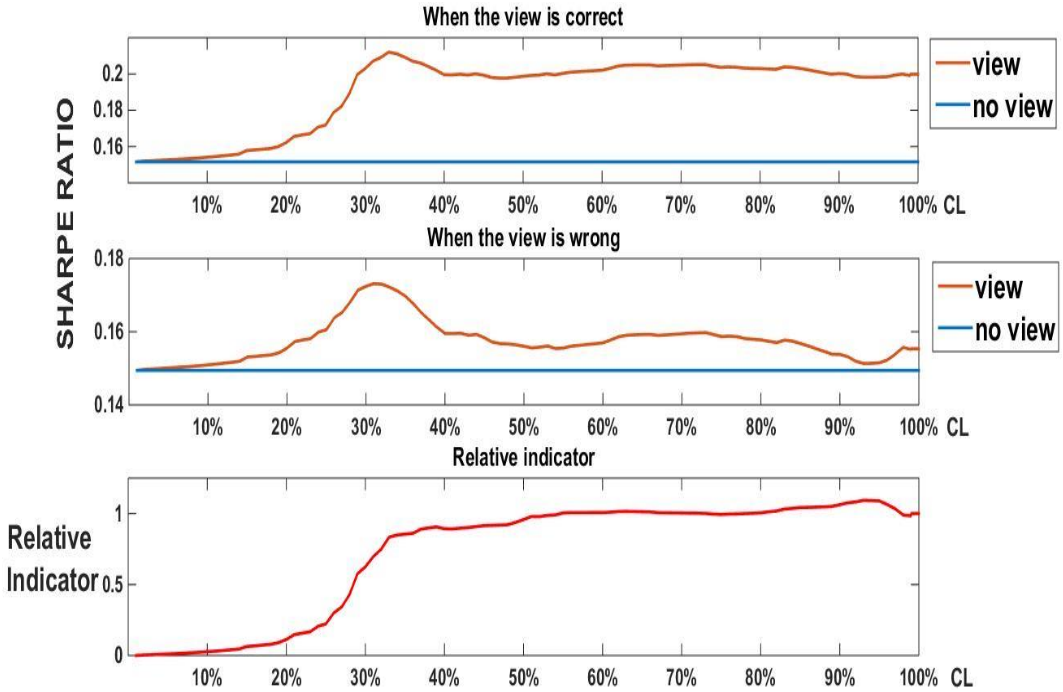

Figure 8 Portfolio performances across the range of confidence levels and Relative indicator (lower plot) in **long time horizon**. The view is correct (upper plot) or not (middle plot) on the aggregated return.

Either when the view is correct or not in a time horizon of 22 periods, the top performance is achieved for a confidence level close to 30%. As the picture shows, the slope of the relative indicator function knows a cut-off. It would be excessively ambitious to calibrate a model to optimise the confidence level, mainly because we have to see how wrong the view was.

Anyway, the information carried by the relative indicator can help us in determining some areas in which the confidence level delivers portfolios with superior performances.

As we have stated before the better performance achieved by a conditioned portfolios compared to the neutral one is due to more efficiency in terms of risk.

The following picture shows how the CVaR portfolio measure is lower for conditioned portfolios, compared to the neutral allocation, depending on the confidence and no matter if the view is correct (figure 9).

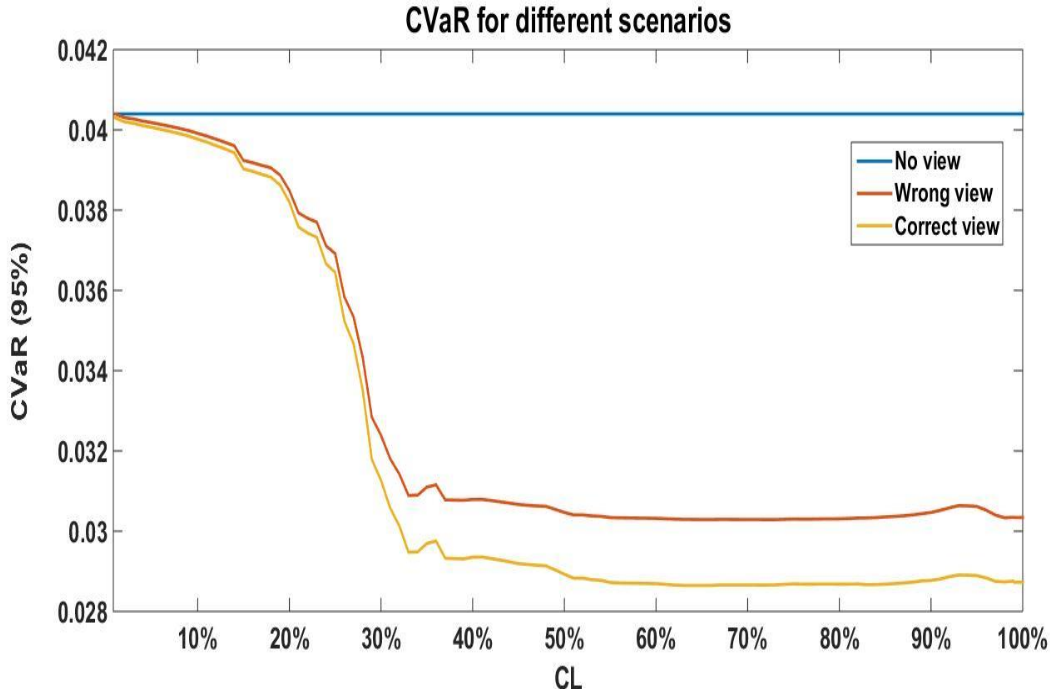

Figure 9 CVaR measures for neutral portfolio and for conditioned portfolio when the view is correct or it is wrong.

The final step of the analysis goes through the multivariate fourth moment of the portfolio. Relevant excess kurtosis is actually detected in the overall portfolio allocation. The purpose is to analyse how efficiently the allocation can face multivariate portfolio kurtosis..

Multivariate portfolio kurtosis is computed as

$$S_K = \mathrm{w}' M_4 (\mathrm{w} \otimes \mathrm{w} \otimes \ \mathrm{w})$$

where M_4 is the co-kurtosis matrix and w are the portfolio weights.

Asset allocation does not suffer from kurtosis even when the portfolio show relevant leptokurtosis: usually for a monotonically increasing relative indicator function with positive concavity, we see the portfolio kurtosis going down for higher confidence levels. This finding is consistent with the fact that for 100% confidence level the uniform distribution[11] of the view substitutes the asset's univariate return distribution. On the other hand, market anomalies can lead to optimised allocation where the kurtosis can remain stable or can slightly increase for higher confidence levels. Even in this case the risk adjusted performance increases with the confidence levels (figure 10).

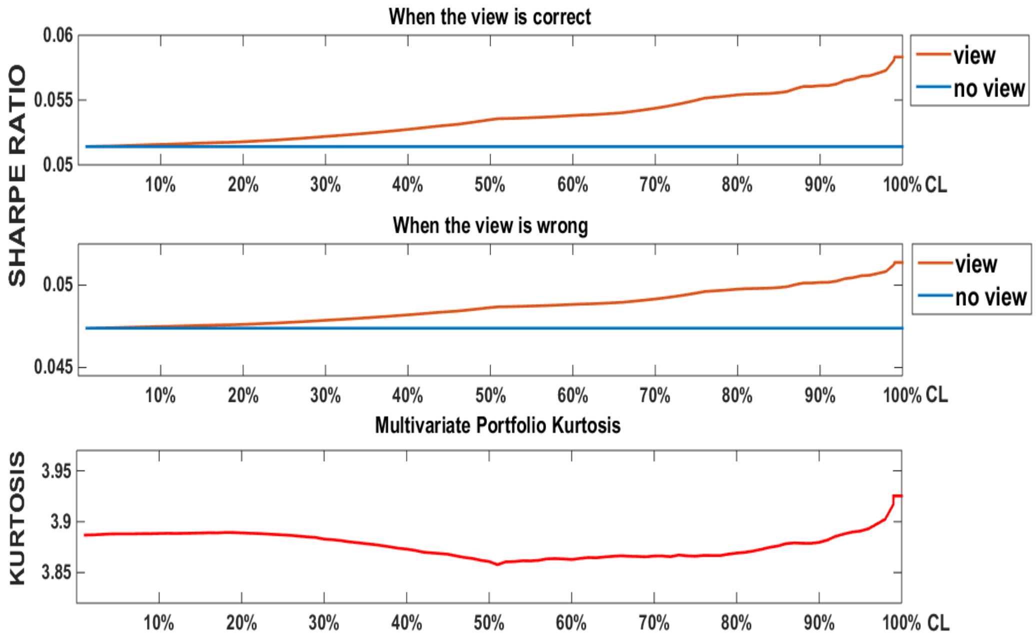

Figure 10 Portfolio performances across the confidence levels and related portfolio kurtosis (lower plot) in **long time horizon**. The view is correct (upper plot) or wrong (middle plot) for the aggregated return.

The performance delivered by a conditioned portfolio is still higher than the neutral one. The same feature is achieved even when the view is wrong, for a sufficiently long time horizon.

One last consideration on transaction costs. We have already highlighted the evidence supporting the advice not to change the weights on the portfolio once the manager has realised that the view is wrong. At first, the manager would regret for having missed the expectations (neutral portfolio is actually performing better at this point) but being patient and counting for cumulated returns, even though the view carries on being wrong, the conditioned portfolio would overperform the neutral one. Changing the asset allocation straight away would import heavy

[11] Uniform distribution has negative excess kurtosis (-6/5)

transaction costs, plus the risk to be wrong one more time, regretting the passive market portfolio repeatedly.

Exogenous information is processed according to this model to better face the non-normality of the market. Leptokurtosis is handled by the Mean-CVaR optimisation. Neutral portfolio is under-diversified: quite often it is concentrated in few assets, “shortening” the capability of the efficient frontier. The view brings benefits to asset allocation and, even if it may be wrong, the better-diversified portfolio would overperform the neutral one. Patience and minimum transaction costs are what really pay out at the end.