Forward-looking Robust Portfolio Selection

Sara Cecchetti*† Laura Sigalotti*‡

July 31, 2012

Abstract

In this paper we develop a portfolio optimization strategy based on the extraction of option-implied distributions and the application of robust asset allocation. We compute the option-implied probability density functions of the constituents of the Euro Stoxx 50 Index. To obtain the corresponding risk-adjusted densities, we estimate the risk aversion coefficient through a Berkowitz likelihood test. The correlation structure among the stocks is computed via an *ad hoc* technique, which provides a correction term for the historical correlations. We implement a robust portfolio construction, in order to incorporate the uncertainty about the estimation error for the expected returns in the optimization procedure.

Contents

*Bank of Italy, Economic Outlook and Monetary Policy Department, Via Nazionale 91, 00184 Roma, Italy. We thank Antonio Di Cesare, Giuseppe Grande, Marcello Pericoli and Marco Taboga for their comments and suggestions. Any remaining error is our own. The views expressed in the article are those of the authors and do not involve the responsibility of the Bank of Italy.

†Email: sara.cecchetti@bancaditalia.it

‡Email: laura.sigalotti@bancaditalia.it

1 Introduction

Modern portfolio theory was introduced by Markowitz (1952), who laid the foundations of mean-variance optimization. This strategy, which is still widely used in the financial industry despite the many criticisms that have been leveled against it, is based on the estimation of the first two moments of the probability density function of the asset returns. The standard way to estimate these moments is based on the use of historical data (see, among others, De Miguel, Garlappi and Uppal (2009), for a review of various estimators). It is well known that the estimation errors that inevitably arise when using such backward-looking approach affect the calculation of the optimal portfolio, since mean-variance optimizers are very sensitive to small variations in expected returns.[1] An alternative methodology has been adopted by Kostakis, Panigirtzoglou and Skiadopoulos (2010), who developed a forward-looking approach based on the information extracted from option prices. They implement a strategy which uses option implied distributions of asset returns to calculate an optimal portfolio consisting of two assets, one risk-free and one risky. Once the risk-neutral implied distributions are extracted, they are converted into the related risk-adjusted (or real-word) distributions, that are required in the calculation of optimal portfolios. As the implied distributions reflect the market participants' expectations, this approach is inherently forward-looking and may provide more accurate estimates of the distribution, and the related moments, to be used in the asset allocation problem.

In this paper, we aim at extending the strategy developed in Kostakis, Panigirtzoglou and Skiadopoulos (2010) to a portfolio of many risky assets. Among the different techniques developed in the finance literature to extract risk-neutral densities (see Bedoui and Hamdi (2010) for a review of these methods), we use the mixture of lognormals approach. Once the risk-neutral distributions are estimated, we convert them into the risk-adjusted ones by assuming a power utility function, in which the risk-aversion coefficient is chosen so that it maximizes the forecasting ability of the risk-adjusted distributions (with respect to the future realizations of the underlying index), using an approach similar to the one followed by Bliss and Panigirtzoglou (2004).

It is worth noting that the generalization of the option-implied distribution approach to a portfolio of many risky assets is not straightforward. In fact, if we consider N stocks and N options written on these stocks, option prices convey information about the distribution of each stock price, but they do not provide any insight on the correlation structure among the N stock prices. The correlation structure could be inferred using only historical data, but this would give a matrix of covariances incoherent with the option implied variances of the N stocks. In order to tackle this issue, we focus on a collection of stocks satisfying two conditions:

- each stock is the underlying asset of a quoted option (actually, we need a family of call and put options for each stock, corresponding to different strikes);

[1] See, e.g., Ceria and Stubbs (2006) for details.

- there exists a family of options written on a index made up of all the stocks in our set, with publicly available index weights.

In order to meet these requirements, we consider a portfolio consisting of the 50 stocks of the the Euro Stoxx 50 Index. Using the market information about the index option prices, we derive the implied probability density function of the stock index and hence, in particular, its variance. Through an *ad-hoc* technique,[2] we integrate the historical covariance matrix of the stock returns with the implied volatilities of the stocks and the index and we derive a correlation structure that can be considered a proper estimation of the real-world covariance matrix of the stock returns. In particular, once we have computed the risk-adjusted standard deviations of the stock returns, the risk-adjusted standard deviation of the index return, and the historical correlation matrix of the stocks returns, we estimate a correction term for the historical correlation matrix which guarantees the coherence between the option-implied volatility of the index and the volatility of the corresponding portfolio of stocks.

Finally, we implement a robust portfolio construction (see Ceria and Stubbs (2006), for a detailed description) to cope with the estimation errors in the estimates of the expected returns and their error-magnification effect. Mean-variance portfolios tend to exacerbate the estimation error problem by significantly overweighting assets with an error to the upside of the expected returns and underweighting assets with an error to the downside. Robust optimization considers the estimation errors for the mean returns directly in the optimization problem itself, in order to perform an asset allocation that is robust to these errors. Namely, the robust optimization approach takes a confidence region for the estimated mean returns, and performs a portfolio optimization accounting for the worst case scenario that could be realized in that confidence region.

Once we have implemented the robust portfolio construction, we compare our results to the efficient frontiers obtained with different methodologies. In particular, we compute the *true* frontier, the Markowitz estimated frontier and the Markowitz actual frontier, that are, respectively: the efficient frontier computed using true expected returns (unobservable and derived through an equilibrium argument); the efficient frontier obtained with estimated expected returns and Markowitz mean-variance portfolio selection; the frontier obtained by calculating true expected returns of the portfolios on the Markowitz estimated frontier. The robust portfolio construction is used to get both a robust estimated frontier and a robust actual frontier, which are compared to the previous ones.

We also compare the out-of-sample performance of our robust portfolio to those of the index buy-and-hold strategy and of two alternative portfolios: the equally-weighted portfolio of De Miguel, Garlappi and Uppal (2009) (in which we build a portfolio with all the N stocks in the index, by assigning weight $1/N$ to each one) and a momentum strategy, in which we select the G stocks with positive returns in the previous year and allocate a portfolio with weight $1/G$ on these stocks. Our comparisons use two standard performance criteria: the out-of-sample Sharpe Ratio and the Certainty-Equivalent

[2] In the paper by Buss and Vilkov (2011), a similar approach is applied to a different context.

(CEQ) return for the expected utility of a mean-variance investor. We find that our robust portfolio has the best performance according to both criteria.

The rest of the paper is structured as follows. Section 2 describes the mixture of lognormals method that we apply to estimate the option-implied risk-neutral densities. Section 3 outlines the methodology adopted to convert the risk-neutral densities into the related risk-adjusted densities. In Section 4 we derive the option-implied risk-adjusted covariance matrix and Section 5 explains the robust optimization approach. The dataset used for the empirical application is described in Section 6 and the results are presented in Section 7. Finally, Section 8 presents the conclusions.

2 Option-Implied Risk Neutral Densities

2.1 The mixture of lognormals method

In the financial literature many methods have been proposed for the estimation of option-implied risk-neutral densities.[3] In particular, two methods have been widely studied and implemented. The first one, know as the "smile interpolation approach", is based on the interpolation of the implied Black-Scholes volatility smile and requires no parametric assumptions on the risk-neutral densities. The starting point of this method is the result by Breeden and Litzenberger (1978), who prove that the second derivative of the price of a call option with respect to the exercise price K, computed for a given value of the underlying price S, is equal to the risk neutral density f of the underlying asset, evaluated in S and discounted by the risk free rate:

$$\frac{\partial^2 C(K)}{\partial K^2} = e^{-r(T-t)} f(S).$$

In order to determine $f(S)$ from the previous formula, we need a sufficiently smooth expression for C: Shimko (1993) showed that fitting the implied Black-Scholes volatility smiles gives much better results than interpolating observed call prices directly.

An alternative approach is to assume a specific parametric form for the risk-neutral density (RND) functions and adjust it to the market data. The parameters of the RNDs are calibrated through the use of observed option prices and nonlinear least squares. The most commonly used functional form for the RNDs is a mixture of two or more lognormals; this choice gives a density function that is sufficiently flexible to reflect characteristics such as excess kurtosis, asymmetry, and even bimodality.

In this paper we follow the parametric approach to estimate the option-implied probability density functions (PDFs), calibrating a mixture of lognormals[4] for the PDFs of the 50 stocks and the stock index. In what follows we briefly describe the mixture of lognormals method.

[3] For a review and comparison of the different estimation methods, see Bedoui and Hamdi (2010).

[4] See Bahra (1996), Melick and Thomas (1997) and Sodernlind (1997) for a detailed description.

Let S_t be the price of an underlying asset which pays a continuous dividend D_t, which follows the dynamic

$$dD_t = \delta S_t dt, \tag{1}$$

δ being a positive constant. We also assume that the underlying asset follows the cum-dividend price dynamics of a generalized Black and Scholes model (BSMG), that is a geometric Brownian motion:

$$dS_t = \mu S_t dt + \sigma S_t dW_t, \tag{2}$$

where $\mu = r - \delta$, r is the continuously compounded risk-free rate and W_t is a standard Brownian motion.[5]

As well known, the price at time t of a call option on such underlying, with strike K and maturity T, is given by the expected value (under the risk-neutral probability measure) of the option at maturity, discounted at the risk-free interest rate:

$$C(S_t, K, T-t, r, \delta, \sigma) = S_t e^{-\delta(T-t)}\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2),$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function, f is the risk-neutral density of the price S_T and

$$d_1 = \frac{\log\left(\frac{S_t}{K}\right) + \left(r - \delta - \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}},$$

$$d_2 = d_1 - \sigma\sqrt{T-t}.$$

The random variable S_T follows a lognormal distribution with mean $m = (\log(S_t) + (r - \delta) - \sigma^2/2)(T-t)$ and variance $\sigma^2(T-t)$, i.e. its risk neutral density f satisfies

$$f(S_T) = l(S_T, m, \sigma) = \frac{1}{S_T\sqrt{2\pi\sigma^2(T-t)}} e^{-\frac{[\log(S_T)-m]^2}{2\sigma^2(T-t)}}.$$

A mixture of lognormals is defined as a convex combination of M lognormal densities in which the parameters m and σ can take different values:

$$q(S_T;\theta) = \sum_{i=1}^{M} \alpha_i l(S_T, m_i, \sigma_i), \tag{3}$$

where θ represents the unknown parameters α_i, m_i, σ_i, for $i = 1, \ldots, M$.[6]

If the price S_T of the underlying asset follows a distribution given by equation (3), then the corresponding option price, for a given strike K and time to maturity $T-t$, equals

$$C^{MIX}(S_t, K, r, T-t, \theta) = \sum_{i=1}^{M} \alpha_i [S_t e^{-\delta_i(T-t)}\Phi(d_{1,i}) - Ke^{-r(T-t)}\Phi(d_{2,i})],$$

[5] The ex-dividend dynamics of the underlying is $dS_t - dD_t = rS_t dt + \sigma S_t dW_t$.

[6] Of course, $\alpha_i > 0$, for each i, and $\sum_i \alpha_i = 1$.

where

$$d_{1,i} = \frac{\log\left(\frac{S_t}{K}\right) + \left(r - \delta - \frac{1}{2}\sigma_i^2\right)(T-t)}{\sigma_i\sqrt{T-t}},$$

$$d_{2,i} = d_1 - \sigma_i\sqrt{T-t}.$$

In this paper we use a mixture of two lognormals ($M = 2$). As we deal with market data for American options (both on the Euro Stoxx 50 Index and on the constituent stocks), the above formula has to be adjusted conveniently. In particular, for each option we calculate the early exercise premium by computing the Barone-Adesi correction term through an iterative procedure.[7] Then, we subtract the correction term from the market price of the American option to obtain the corresponding European option price. This price is now matched to the value C^{MIX} obtained as follows:

$$C^{MIX}(S_t, K, r, T-t, \alpha, \delta_i, \sigma_i) = \alpha C^{BSM}(S_t, K, T-t, r, \delta_1, \sigma_1)$$
$$+(1-\alpha)C^{BSM}(S_t, K, T-t, r, \delta_2, \sigma_2),$$

where $C^{BSM}(S_t, K, T-t, r, \delta, \sigma)$ denotes the standard Black and Scholes (1973) formula for options on a dividend-paying underlying. Finally, we estimate the parameters $\alpha, \delta_1, \delta_2, \sigma_1, \sigma_2$ that better replicate the observed market prices of the options.

3 Option-Implied Risk Adjusted Densities

In the previous paragraph we described the method we used to extract the option-implied risk neutral densities from market data. In order to perform a portfolio optimization we need to know the real-world probability density functions. To this end, we convert the RNDs into the corresponding risk-adjusted PDFs.

If there exists a representative agent with utility function $U(\cdot)$, then the link between the risk-neutral distribution measured at time t, $q_t(S_T)$, and the real world distribution $p_t(S_T)$ of the asset price S_T, is given by

$$p_t(S_T) = \frac{q_t(S_T)}{\zeta(S_T)}\left(\int \frac{q_t(x)}{\zeta(x)}dx\right)^{-1}, \tag{4}$$

where

$$\zeta(S_T) = e^{-r(T-t)}\frac{U'(S_T)}{U'(S_t)} \tag{5}$$

is the so-called pricing kernel. Equation (5) follows from the first-order condition of the intertemporal expected utility maximization problem of the representative agent (see Ait-Sahalia and Lo (2000) for a detailed discussion). We assume that the representative

[7] See Barone-Adesi and Whaley (1987) for details.

agent maximizes a power utility function,[8] defined as

$$U(W) = \frac{W^{1-\gamma} - 1}{1-\gamma}, \qquad \gamma \neq 1,$$

where γ is the coefficient of constant relative risk aversion (RRA) that must be estimated. According to our choice of the utility function, we compute the risk-adjusted density from the risk-neutral density as

$$p_t(S_T) = \frac{q_t(S_T) S_T^{\gamma}}{\int q_t(x) x^{\gamma} dx}. \tag{6}$$

In particular, we assume that the risk aversion coefficient is the same for the index and for all the constituent stocks, reflecting the view of a unique representative agent who decides how to allocate his investments among a class of assets. In the following paragraph we will describe the method we use to estimate the risk aversion coefficient γ, following the approach of Bliss and Panigirtzoglou (2004).

3.1 Option-Implied Risk Aversion coefficients estimates

In this paragraph we deal with the estimate of the optimal risk aversion coefficient γ. We consider options on the Euro Stoxx 50 Index, with monthly maturities ranging from 20 November 2009 to 17 June 2011, thus getting 20 admissible expiration dates. We fix a time-to-maturity of 78 days (0.21 years) and we consider the option prices in the dealing date 78 days ahead each maturity date. Having fixed a pair dealing date/maturity date, we extract the corresponding risk-neutral probability density functions using the mixture of lognormals method described in Section 2.1. For each value of the parameter γ, chosen in the range $\{0.5, 0.6, 0.7, \ldots, 9.8, 9.9, 10\}$, we convert the 20 RNDs into the corresponding risk-adjusted probability density functions, using equation (6). For each γ we evaluate the forecasting ability of the corresponding risk-adjusted PDF, then we select the value of the risk aversion coefficient which maximizes the consistency of future realizations of then index price with the option-implied probability density functions. In what follows we describe the method used to select γ, which is based on a Berkowitz likelihood ratio statistic.[9]

Having fixed a maturity date T, the corresponding dealing date t and a value of γ, we extract the risk-adjusted probability density function $p_t^{(\gamma)}(\cdot)$ for the index price S_T using the procedure explained in the previous paragraphs (see equation 6). If we denote by $q_t(\cdot)$ the option implied risk neutral density, we recall that $p_t^{(\gamma)}(\cdot)$ is given by

$$p_t^{(\gamma)}(S_T) = q_t(S_T) S_T^{\gamma} \Big(\int_0^{+\infty} q_t(x) x^{\gamma} dx \Big)^{-1}.$$

[8] This choice guarantees the integrability of $x \mapsto q_t(x)/U'(x)$ when q_t is a mixture of lognormals. If we had adopted an exponential utility function, for instance, the corresponding function would have had an infinite integral.

[9] See Berkowitz (2001) for details.

We want to test the hypothesis that the estimated PDFs $p_t^{(\gamma)}(\cdot)$ are equal to the true (unknown) PDFs $f_t(\cdot)$. At time t we forecast the future realizations of the index price at time T using the risk-adjusted density $p_t^{\gamma}(\cdot)$. At time T we observe the realized value of the price, which we denote here by X_t to remind that it has to be compared with its expected value computed at time t. The null hypothesis states that the realizations X_t are independent and that $p_t^{\gamma}(\cdot) = f_t(\cdot)$. If this hypothesis is verified, then the inverse probability transformations of the realizations

$$y_t = \int_{-\infty}^{X_t} p_t^{(\gamma)}(u)du \tag{7}$$

are $i.i.d. \sim U(0,1)$.

The Berkowitz likelihood ratio statistic allows to test jointly for uniformity and independence. To implement this methodology we consider the transformation of y_t

$$z_t = \Phi^{-1}(y_t) = \Phi^{-1}\Big(\int_{-\infty}^{X_t} p_t^{(\gamma)}(u)du\Big), \tag{8}$$

where $\Phi(\cdot)$ is the normal cumulative density function. In order to test the independence and standard normality of the z_t, Berkowitz (2001) uses the maximum likelihood to estimate the following autoregressive model of order 1

$$z_t - \eta = \theta(z_{t-1} - \eta) + \epsilon_t, \tag{9}$$

and tests restrictions on the parameters of the AR(1) using a likelihood ratio test. Under the null hypothesis, the model has the following parameters: $\eta = 0$, $\theta = 1$, $\sigma^2 := Var(\epsilon_t) = 0$. The log-likelihood function for this model[10] is given by

$$L(\eta, \sigma^2, \theta) = -\frac{1}{2}\log(2\pi) - \frac{1}{2}\log[\sigma^2/(1-\theta^2)] - \frac{\{z_1 - [\eta/(1-\theta)]\}^2}{2\sigma^2/(1-\theta^2)}$$
$$-[(\tau-1)/2]\log(2\pi) - [(\tau-1)/2]\log(\sigma^2) - \sum_{t=2}^{\tau}[\frac{(z_t - \eta - \theta z_{t-1})^2}{2\sigma^2}].$$

Under the null hypothesis, the likelihood ratio statistic

$$LR_3 = -2[L(0,1,0) - L(\widehat{\eta}, \widehat{\sigma}^2, \widehat{\theta})]$$

is distributed as a chi-squared with 3 degrees of freedom, ${\chi^2}_3$. The risk aversion parameter γ is chosen as the value which maximizes the p-value of the Berkowitz LR_3 statistic, which is a measure of the forecast ability of the correponding risk-adjusted PDFs.

As the presence of overlapping data may induce autocorrelation[11] and lead to the rejection of any value of γ, we selected a subsample of the time series, such that the maturity date of the n-th option is prior to the dealing date of the $(n+1)$-th option. We

[10] See Hamilton (1994), equation (5.2.9).
[11] See Bliss and Panigirtzoglou (2004).

found that the p-value of the Berkowitz test is maximized for $\widehat{\gamma} = 1.5$. This value is not far from the option-implied risk aversion estimates obtained by Bliss and Panigirtzoglou (2004), who compute the representative agent's risk aversion at different horizons, getting values between 1.97 and 7.91. They consider forecast horizons ranging from 1 to 6 weeks and suggest that the risk-aversion coefficient tends to decrease as the time lapse increases. Since our forecast horizon is 11 weeks, the optimal value of $\widehat{\gamma} = 1.5$ seems reasonable.

3.2 Risk Adjusted Expected Returns and Variances

Having estimated the risk aversion parameter γ, we assume that it reflects the representative agent's view for all the constituents of the Euro Stoxx 50 Index. Then, we use equation (6) to compute the risk-adjusted densities ${f_t}^i(\cdot) = p_{t,i}^{\widehat{\gamma}}(\cdot)$ for the underlying prices S_T^i, where $i \in \{1, \ldots, 50\}$, $t = 31$ March 2011 and $T = 17$ June 2011.

To perform the asset allocation we need to know the estimated expected returns and the estimated return variances, given by

$$\overline{\mu}_i = \mathbb{E}[\log(S_T^i) - \log(S_t^i)] = \int_0^{+\infty} \log(x) f_t^i(x) dx - \log(S_t^i) \tag{10}$$

and

$$\begin{aligned}({\sigma^{RA}}_i)^2 &= Var(\log(S_T^i) - \log(S_t^i)) \\ &= \int_0^{\infty} (\log(x))^2 f_t^i(x) dx - \left(\int_0^{+\infty} \log(x) f_t^i(x) dx \right)^2.\end{aligned} \tag{11}$$

The risk-adjusted density of the index price is denoted by f_t^{Ind}; its return variance equals

$$\begin{aligned}({\sigma^{RA}}_{Ind})^2 &= Var(\log({S_T}^{Ind}) - \log({S_t}^{Ind})) \\ &= \int_0^{+\infty} (\log(x))^2 {f_t}^{Ind}(x) dx - \left(\int_0^{+\infty} \log(x) {f_t}^{Ind}(x) dx \right)^2.\end{aligned} \tag{12}$$

4 Option-Implied Risk Adjusted Covariance

We compute the risk-adjusted variance-covariance matrix for the individual stock returns in our portfolio, using the (risk-adjusted) standard deviations of the stocks returns, the (risk-adjusted) standard deviation of the index returns and the historical correlation matrix of the stocks returns. The historical correlations are estimated using six years of daily returns $R_{i,t}$ and the exponential weighted moving average method.[12] We propose an *ad hoc* method to combine the implied volatilities with the historical correlations, through the computation of a correction coefficient β.

[12]See, among others, Mills and Markellos (2008).

4.1 Historical correlation matrix

An exponentially weighted moving average applies weighting factors which decrease exponentially as the observations become more and more distant in time. As in RiskMetrics,[13] we consider a smoothing constant $\lambda = 0.94$ and we compute the m-period[14] historical covariances between stocks i and j, measured at time $t+1$, as

$$\sigma^h{}_{ij,t+1} = \frac{\sum_{\tau=1}^{m}(R_{i,t+1-\tau} - \overline{R_i})(R_{j,t+1-\tau} - \overline{R_j})\lambda^{\tau-1}}{1 + \lambda + \lambda^2 + \ldots + \lambda^{m-1}}, \tag{13}$$

where for each stock $\overline{R_i} = \frac{1}{m}\sum_{t=1}^{m} R_{i,t}$. Hence, m-period historical variance of stock i, measured at time $t+1$, is given by

$$(\sigma^h{}_{i,t+1})^2 = \frac{\sum_{\tau=1}^{m}(R_{i,t+1-\tau} - \overline{R_i})^2\lambda^{\tau-1}}{1 + \lambda + \lambda^2 + \ldots + \lambda^{m-1}}.$$

The related correlations are given by

$$\rho^h{}_{ij,t+1} = \frac{\sigma^h{}_{ij,t+1}}{\sigma^h{}_{i,t+1}\sigma^h{}_{j,t+1}}, \tag{14}$$

where $\sigma^h{}_{i,t+1}$ is the historical standard deviation for stock i.

4.2 Implied Risk-Adjusted Correlation Matrix

Once we have calculated the historical correlations, we compute a perturbed correlation matrix which takes into account the option-implied estimates.[15] We denote by $\epsilon_{ij,t}$ the difference between the historical correlations of stocks i and j and our *risk-adjusted* correlations, computed at time t:

$$\widehat{\rho^{RA}}{}_{ij,t} = \rho^h{}_{ij,t} + \epsilon_{ij,t}. \tag{15}$$

We assume that the correction term can be written in the form

$$\epsilon_{ij,t} = \beta_t \times \frac{(\sigma^{RA}{}_{Ind,t})^2}{\sigma^{RA}{}_{i,t}\sigma^{RA}{}_{j,t}}, \tag{16}$$

where $\sigma^{RA}{}_{Ind,t}$ is the risk-adjusted implied standard deviation of the index, $\sigma^{RA}{}_{i,t}$ is the risk-adjusted implied standard deviation of stock i, and β_t is the unknown coefficient to be estimated. This choice for the correction term allows to take into account the option-implied risk-adjusted variances for both the stocks and the index. Moreover, the magnitude of the term $\epsilon_{ij,t}$ is comparable to the one of the historical correlations $\rho^h{}_{ij,t}$. In order to estimate β_t, we impose that the option-implied variance of the index

[13] See RiskMetrics (1996).

[14] We consider daily data from 8th July, 2005 to 1st April, 2011, so that m=1495.

[15] A similar approach, but in a different context, is used by Buss and Vilkov (2011).

is equal to the variance of the portfolio made of the constituent stocks, calculated with the option-implied stock variances:

$$
\begin{aligned}
({\sigma^{RA}}_{Ind,t})^2 &= \sum_i \sum_j w_i w_j {\sigma^{RA}}_{ij,t} \\
&= \sum_i {w_i}^2 ({\sigma^{RA}}_{i,t})^2 + \sum_i \sum_{j \neq i} w_i w_j {\sigma^{RA}}_{i,t} {\sigma^{RA}}_{j,t} {\rho^{RA}}_{ij,t},
\end{aligned}
\tag{17}
$$

where w_i are the index portfolio weights. Having assumed that formula (16) describes the correction term, the risk-adjusted correlations defined in (15) must be re-normalized:

$$
{\rho^{RA}}_{ij,t} = \frac{\widehat{\rho^{RA}}_{ij,t}}{\sqrt{\widehat{\rho^{RA}}_{ii,t}}\sqrt{\widehat{\rho^{RA}}_{jj,t}}}. \tag{18}
$$

Plugging equation (18) in (17), we get

$$
\begin{aligned}
({\sigma^{RA}}_{Ind,t})^2 = \sum_i \sum_j w_i w_j {\sigma^{RA}}_{i,t} {\sigma^{RA}}_{j,t} \\
\times \frac{{\rho^h}_{ij,t} + \beta_t \times \frac{({\sigma^{RA}}_{Ind,t})^2}{{\sigma^{RA}}_{i,t}{\sigma^{RA}}_{j,t}}}{\sqrt{{\rho^h}_{ii,t} + \beta_t \times \frac{({\sigma^{RA}}_{Ind,t})^2}{({\sigma^{RA}}_{i,t})^2}}\sqrt{{\rho^h}_{jj,t} + \beta_t \times \frac{({\sigma^{RA}}_{Ind,t})^2}{({\sigma^{RA}}_{j,t})^2}}}.
\end{aligned}
$$

Since β_t is the only unknown variable in the previous expression, we can calculate it by implementing a numerical method in Matlab. Once we have obtained β_t, we can compute the implied risk-adjusted correlations ${\rho^{RA}}_{ij,t}$ and consequently the risk-adjusted covariances

$$
{\sigma^{RA}}_{ij,t} = {\rho^{RA}}_{ij,t} {\sigma^{RA}}_{i,t} {\sigma^{RA}}_{j,t}.
$$

We proposed this *ad hoc* method to use all the available information, arising both from the options on the index and from the options on the 50 stocks. The option-implied risk-neutral densities do not provide any information about the correlation structure among the stocks, hence we decided to use as a starting point the historical correlation matrix given by an exponentially weighted moving average model. This historical correlation matrix is then modified by the introduction of a correction coefficient β, chosen so that it guarantees the matching between the option-implied volatility of the index and the volatility of the corresponding portfolio of stocks.

As an output of the numerical computations, we obtained that the optimal correction term is $\beta_t = 0.97$ per cent, for $t = 31$ March 2011.

5 Robust optimization

In this paragraph we show how we applied the robust portfolio construction to our setting. The robust optimization methodology has been developed to cope with estimation errors in the estimates of the expected returns and their error-magnification effects. It is well known that mean-variance portfolios tend to exacerbate the estimation error problem by significantly overweighting assets with an error to the upside and underweighting assets with an error to the downside. Robust optimization considers the estimation errors for the mean returns directly in the optimization problem itself, and aims at performing an asset allocation that is robust to these errors. Namely, the robust optimization approach fixes a confidence region for the estimated mean returns, and performs a portfolio optimization accounting for the worst case scenario that could be realized in that confidence region.

We want to select the minimum-variance portfolio over all asset allocations which guarantee an expected return bigger than or equal to a fixed target $\bar{r}$. Let R be the N-dimensional random variable describing the random returns of the N available stocks. Let μ and $\Sigma = (\Sigma_{ij})_{1\leq i,j\leq N}$ be the mean vector and the variance-covariance matrix of the stock returns. Having fixed $\bar{r} > 0$, we aim at solving the following problem:

$$\begin{aligned} \min_{w} \quad & w'\Sigma w \\ s.t. \quad & \mu' w \geq \bar{r} \\ & \mathbf{1}'w = 1. \end{aligned} \tag{19}$$

It is well known that the estimation of expected returns is a particularly challenging problem, since errors in estimated mean returns are the main determinants of optimal portfolio estimation risk (they account for most of the estimation error, in much bigger measure than errors in estimated variances). A possible way to face this issue consists in considering uncertainty in unknown parameters explicitly in the optimization problem. This approach, which is part of the robust optimization field, was introduced by Ben-Tal and Nemirovski (1997) for robust truss topology design. In what follows, we describe how robust optimization can be applied to our asset allocation problem, following the approach of Ceria and Stubbs (2006).

Since the true value of μ cannot be known with certainty, we model it as a random variable whose dispersion represents the possible estimation error. In particular, we assume that μ is normally distributed around its best guess π:

$$\mu \sim N(\pi, Q), \tag{20}$$

where the N-dimensional matrix Q represents the uncertainty of the guess. Actually it is not necessary to assume normality: the same arguments hold with minor changes if we assume that the distribution of μ is elliptical.

Since we do not have any *a priori* information on the covariance matrix Q of the dispersion error, we can simply assume that Q is the n-dimensional identity matrix. Given the best guess vector π, with probability $\eta \in [0, 1]$ the vector of actual expected returns lies inside the confidence region

$$E = \{\mu \in \mathbb{R}^n : (\mu - \pi)' Q^{-1} (\mu - \pi) \leq \kappa^2\}, \tag{21}$$

where $\kappa^2 = F_n^{-1}(1 - \eta)$ and F_n^{-1} is the inverse cumulative distribution function of a chi-squared distribution with N degrees of freedom. In the definition of E the variables

Q and κ represent respectively the shape and the size of the ellipsoid. In particular we set $\kappa^2 = \text{Var}(\pi_1^0, \ldots, \pi_n^0)$, i.e. κ^2 is the cross-sectional variance of the mean returns.

Our goal consists in performing a portfolio optimization which accounts for the worst-case scenario, hence we want to solve the following minimum-variance problem:

$$\begin{aligned} \min_{w} \quad & w'\Sigma w \\ \text{s.t.} \quad & \min_{\mu \in E} \mu' w \geq \bar{r} \\ & \mathbf{1}'w = 1. \end{aligned} \tag{22}$$

In the following paragraph we will describe the cutting plane algorithm we implemented to find a numerical solution of problem (22).

A key issue in the implementation of the robust asset allocation is the determination of the *best guess* vector π, which is the starting point of the robust optimization method. In this paper the vector π is set through an equilibrium argument.[16] If there were no estimation error, i.e. $Q_{ij} = 0$ for all $i, j \in \{1, \ldots, n\}$, then $\mu = \pi$ a.s. and the stock returns would follow the distribution

$$R \sim N(\pi, \Sigma).$$

In this scenario we could assume that all investors maximize a mean-variance trade-off and that the maximization is unconstrained; the investors' optimization problem would read as

$$\tilde{w} = \arg\max_{w} \{w'\pi - \lambda w'\Sigma w\},$$

where λ is the coefficient of risk aversion. By imposing the first order conditions we find the following link between the equilibrium returns π and the optimal weights $\tilde{w}$:

$$\pi = 2\lambda\Sigma\tilde{w}. \tag{23}$$

In our portfolio optimization we will take λ equal to the value of the risk aversion coefficient γ we obtained in section 3.1, i.e. the value which maximizes the forecast ability of the densities implied in the index options. The weights $\tilde{w}$ will be equal to the weights of the stock index.

In the empirical part of this work, we computed the robust efficient frontier in two different ways: first of all we calculated the actual robust efficient frontier, in which the starting point of the robust optimization was the equilibrium vector π defined in (23). Then we computed the estimated robust efficient frontier, in which the starting point was the vector $\widehat{\pi}$, obtained as follows: we randomly generated a time-series of normally distributed returns (with mean π and variance Σ) and we computed the average $\widehat{\pi}$ to use as an estimate of the expected returns.

---

[16]See, among others, Meucci (2010) for a detailed explanation.

5.1 The Cutting Plane approach

The solution to problem (22) is not straightforward, due the presence of the first constraint: the admissible weights w must satisfy the nonlinear condition $\mu' w \geq \bar{r}$ for all μ belonging to the uncountable set E. In order to find a numerical solution to this problem, we implement a *cutting-plane algorithm*.

1. First of all, we set $\mu_0 = \pi$ and we define $S_0 = \{\mu_0\}$. Then we solve the constrained problem

$$\begin{aligned} \min_{w} \quad & w'\Sigma w \\ \text{s.t.} \quad & \mu' w \geq \bar{r} \text{ for all } \mu \in S_0, \quad \mathbf{1}' w = 1. \end{aligned} \tag{24}$$

   Let w_0 be the solution of (24).

2. We want to determine the vector of expected returns which corresponds to the worst case scenario given the weights vector w_0. Namely, we determine $\mu_1 = \arg\min_\mu \{\mu' w_0 : \mu \in E\}$. Now, two outcomes can occur:

   - either $\mu_1' w_0 \geq \bar{r}$, which implies that w_0 is a feasible solution for problem (22). In this case the algorithm ends and $\tilde{w} = w_0$ is the optimal weights vector;
   - or $\mu_1' w_0 < \bar{r}$, which implies that w_0 is not an admissible solution for problem (22), since it violates the first constraint. In this case the procedure continues with the updated constraints set $S_1 = \{\mu_0, \mu_1\}$. We find the solution w_1 of the optimization problem

$$\begin{aligned} \min_{w} \quad & w'\Sigma w \\ \text{s.t.} \quad & \mu' w \geq \bar{r} \text{ for all } \mu \in S_1, \quad \mathbf{1}' w = 1. \end{aligned}$$

   Note that the vector μ_1 can be computed analitically. In fact, μ_1 is the solution of the minimization problem

$$\min_{\mu \in E} \mu' w_0. \tag{25}$$

   For a fixed $\mu \in E$, we set $v = Q^{-1/2}(\mu - \mu_0)$, so that $\mu = \mu_0 + Q^{1/2} v$. Now, μ belongs to E if and only if $(\mu - \mu_0)' Q^{-1} (\mu - \mu_0) \leq \kappa^2$; there follows that $\mu \in E$ is equivalent to $\|v\|_2 = v'v \leq \kappa^2$. Hence the minimum problem in (25) is equivalent to

$$\min_{\|v\|_2 \leq \kappa^2} (\mu_0 + Q^{-1/2} v)' w_0. \tag{26}$$

   Now, problem (26) can be solved analytically and its unique solution is

$$v_1 = -\kappa \frac{Q^{1/2} w_0}{\|Q^{1/2} w_0\|}.$$

   The corresponding vector μ_1 is the solution of problem (25):

$$\mu_1 = \mu_0 + Q^{1/2} v_1 = \mu_0 - \kappa \frac{Q w_0}{\sqrt{w_0' Q w_0}}. \tag{27}$$

3. When we reach the j-th iteration, we start with a constraints set $S_{j-1} = \{\mu_0, \ldots, \mu_{j-1}\}$ and a vector of weights w_{j-1} which solves the minimum-variance problem

$$\begin{aligned} \min_{w} \quad & w'\Sigma w \\ \text{s.t.} \quad & \mu' w \geq \bar{r} \text{ for all } \mu \in S_{j-1}, \quad \mathbf{1}' w = 1. \end{aligned}$$

   Then, we determine the worst-case expected returns $\mu_j = \arg\min_\mu \{\mu' w_{j-1} : \mu \in E\}$. We can find an explicit formula for μ_j arguing as above:

$$\mu_j = \mu_0 - \kappa \frac{Q w_{j-1}}{\sqrt{w_{j-1}' Q w_{j-1}}}.$$

   Then we look at the scalar product $\mu_j' w_{j-1}$ and we get one of the following outcomes:

   - if $\mu_j' w_{j-1} \geq \bar{r}$, then we have found the solution $\tilde{w} = w_{j-1}$ to problem (22) and the algorithm ends;
   - otherwise, if $\mu_j' w_{j-1} < \bar{r}$, then we iterate the procedure.

For a general discussion of cutting-set methods for robust convex optimization and their convergence see e.g. Mutapcic and Boyd (2009).

5.2 Constraints on the portfolio weights

In general, mean-variance efficient portfolios constructed using sample moments often assign extremely negative and positive weights to a number of assets. Since negative portfolio weights (short positions) are difficult to actually implement, many investors impose no short-sale constraints (i.e., portfolio weights must be nonnegative). This choice finds empirical support in the paper of Jagannathan and Ma (2003), that shows how imposing appropriate constraints improves the efficiency of the constructed optimal portfolios. In particular, the authors show that constraining portfolio weights to be nonnegative is equivalent to shrinking the sample covariance matrix (i.e., reducing its large elements) and then forming the optimal portfolio without any restriction on its weights. They show that each of the no short-sale constraints is equivalent to reducing the sample covariances of the corresponding asset with other assets by a certain amount. The intuition behind this result is due to the fact that assets that have high covariances with other stocks tend to get extreme negative portfolio weights. The paper also shows that imposing upper bounds on portfolio weights does not lead to a significant improvement in the out-of-sample performance of minimum risk portfolios when no short-sales restrictions are already in place, but a constraint from above can help in the practical construction of the portfolio. Summarizing the authors findings, we can say that constructing a minimum risk portfolio subject to the constraint that portfolio weights are positive (negative) is equivalent to constructing it without any constraint on portfolio weights after modifying the covariance matrix by a modification that shrinks the larger elements of the covariance matrix towards zero (towards one). The striking feature of minimum-variance portfolios with no-short-sale constraints is that in such portfolios investment is spread over only a few stocks, while imposing upper bounds on portfolio weights can ensure that optimal portfolios will contain a large enough collection of stocks.

In this paper we implement the asset allocation strategy in two different ways:[17] first, we impose no short-sale restrictions and then we constrain the portfolio weights to be in the interval $[-1, 1]$. In the first case our minimum-variance problem (22) thus becomes

$$\begin{aligned} \min_{w} \quad & w'\Sigma w \\ \text{s.t.} \quad & \min_{\mu \in E} \mu' w \geq \bar{r} \\ & \mathbf{1}' w = 1 \\ & w_i \geq 0 \qquad i =_1, \dots n, \end{aligned} \tag{28}$$

while in the second case we have

$$\begin{aligned} \min_{w} \quad & w'\Sigma w \\ \text{s.t.} \quad & \min_{\mu \in E} \mu' w \geq \bar{r} \\ & \mathbf{1}' w = 1 \\ & -1 \leq w_i \leq 1 \qquad i = 1, \dots n. \end{aligned} \tag{29}$$

The two sets of results are discussed in Section 7.

6 The Dataset

The dataset consists of the stocks which compose the Euro Stoxx 50 Index. The index covers 50 stocks from 12 Eurozone countries.[18] The 50 stocks in the index are generally very liquid. For each of them there are call and put options, with various strikes, quoted in the market (with the only exception of the stocks of the company CRH, for which there were no quoted options at the time we built our dataset). Since we were interested in a static asset allocation, we fixed a specific date (31 March 2011) as the initial time of our analysis, and for each stock we focused on the collection of options which expired in June 2011 (for most of the options the maturity date was June 17, while for the options on stocks corresponding to Italian companies - Generali, Enel, Eni, Intesa San Paolo, Telecom Italia, Unicredit - the expiration date was June 16). For each option the time-to-maturity was thus equal to 0.21 years.

[17]In addition, we implemented the portfolio optimization without any constraint on the portfolio weights, but the results were highly unsatisfactory due to the presence of extremly positive and extremely negative weights, as expected.

[18]Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal and Spain.

Having fixed the expiration date, for each stock we gathered the data regarding call and put options written on that stock (strike price, underlying price and market value of the option)[19] as well as the information about the options written on the Euro Stoxx 50 Index. In the case of the index, we built a time series made of 20 periods: we started from the family of options expiring on 17 June 2011 and quoted on 31 March 2011 and we went backwards for 20 months, taking market prices of options on Euro Stoxx 50 having time to maturity equal to 0.21 years. This time series was used to determine the risk aversion coefficient γ (see Section 3). The market prices of the option expiring on 17 June 2011 were also used in the construction of the variance/covariance matrix (see Section 4).

For the estimate of the historical correlation matrix, we used the daily prices of the 50 stocks composing the index from 8 July 2005 to 1 April 2011.

7 Empirical results

In order to measure the effect of the robust optimization methodology, we construct five efficient frontiers:

- estimated Markovitz efficient frontier: the efficient frontier obtained by using the estimated expected-return vector $\widehat{\pi}$ to compute the optimal portfolio according to the standard Markovitz mean-variance approach;
- estimated robust efficient frontier: the efficient frontier obtained by using the estimated expected-return vector $\widehat{\pi}$ to compute the optimal portfolio according to the robust optimization approach;
- true efficient frontier: the efficient frontier computed by using the true expected returns (unobservable, and represented by the equilibrium returns π) to compute the optimal portfolio;
- actual Markovitz efficient frontier: the efficient frontier obtained taking the portfolios on the estimated Markowitz efficient frontier and then calulating their expected returns using the true expected returns;
- actual robust efficient frontier: the efficient frontier obtained by using the true expected returns π to compute the optimal portfolio according to the Robust Optimization approach.

Figure 1 and Figure 2 illustrate these five efficient frontiers obtained when considering no short-sale restrictions and lower and upper bound restrictions on the portfolio weights, respectively. As expected, in both cases we can see that when we use robust optimization the actual and the estimated frontiers lie closer to each other with respect to the related frontiers obtained with the standard Markovitz mean-variance optimization. This is

[19] All the data are from Thomson Reuters Datastream.

Figure 1 – Efficient frontiers obtained with no short-sale constraint.

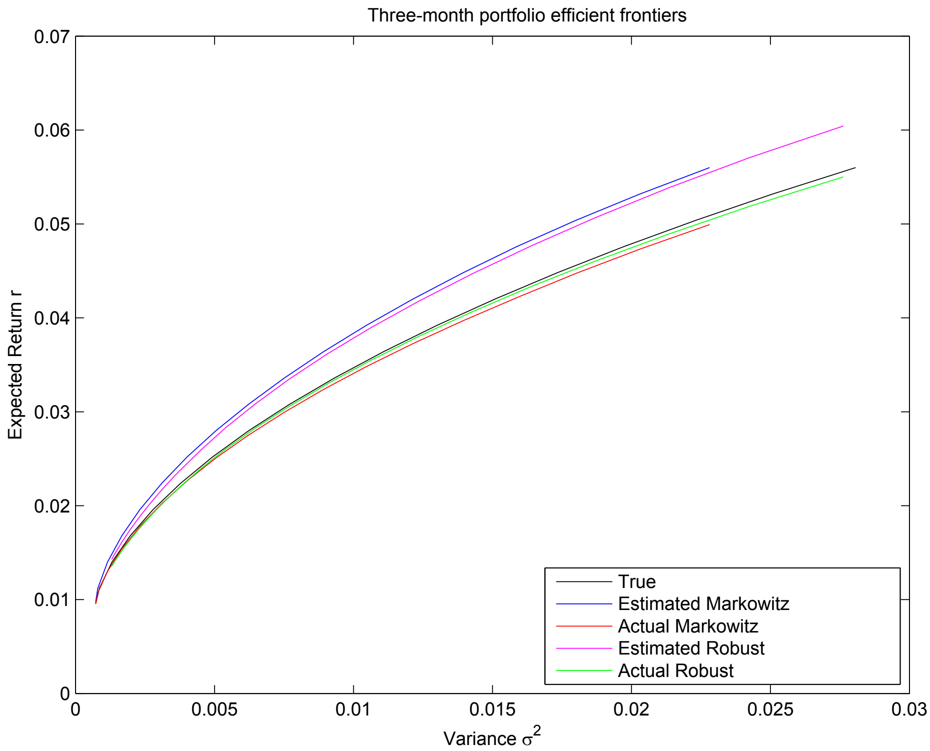

because, by construction, the objective function in the robust optimization problem is based on reducing the distance between the predicted and the actual frontier.[20] In addition, by incorporating the estimation errors in the portfolio construction process, we significantly reduce their effects on the optimal portfolio. Moreover, what is remarkable is that the robust estimated and actual efficient frontier are not only closer together, but also closer to the true efficient frontier.

Table 1 shows the number of stocks selected in the Robust and Markovitz optimization problems, without and with no short-sale constraint, respectively (i.e., the number of stocks which are assigned a non-zero weight). This table displays also the minimum and the maximum values for the asset weights. The results are obtained when setting the minimum expected annual return of the portfolio equal to 0.10, that is the historical mean return of the Euro Area Market General Index (computed by Datastream).

From Table 1 we can see that when we do not impose no short-sale constraint on the portfolio weights (± 1) in our optimization problems (both robust and Markovitz), we select all the available stocks in the index. In this case the weights reflect more extreme short and long positions in the Markovitz portfolio rather than in the robust portfolio, as we can see from the higher absolute value of both the minimum and the

[20]See Ceria and Stubbs (2006) for technical details.

Figure 2 – Efficient frontiers obtained without no short-sale constraint.

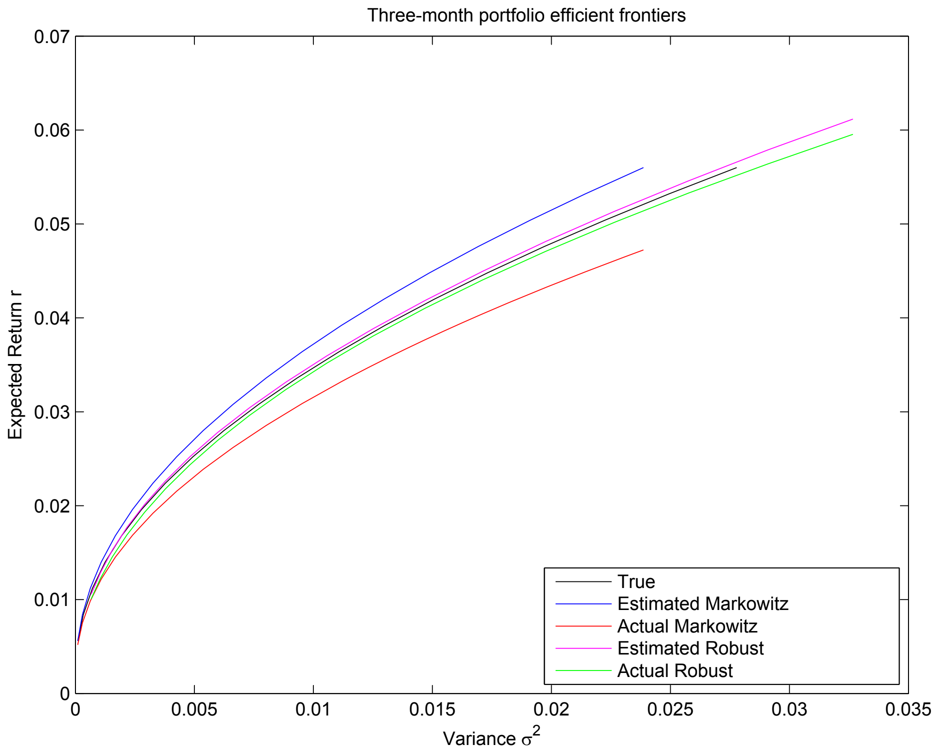

Table 1 – Number of stocks and weights for the different portfolios (1).
(1) LU refers to lower and upper bounds ±1, while NSS refers to no short-sale constraints.

	Number of stocks	Minimum weight	Maximum weight
Robust LU	50	-0.0368	0.0763
Markovitz LU	50	-0.1994	0.7003
Robust NSS	33	0.0057	0.1091
Markovitz NSS	13	0.0018	0.4212

maximum weights. When we impose no short-sale constraint, on the other hand, the robust portfolio selects a more consistent number of stocks than the Markovitz portfolio. Moreover, the robust portfolio assignes more homogeneous weights to the selected stocks.

Finally, in order to evaluate the performance of our model we consider two criteria:

- the out-of-sample *Sharpe Ratio*, defined as the sample mean of out-of-sample excess returns (over the risk-free rate), divided by their sample standard deviation:

$$SR_{out} = \frac{(\mu_{real} - r_f)^T w}{\sqrt{w^T \Sigma_{real} w}} \tag{30}$$

where μ_{real} are the realized returns of the stocks in the portfolio, r_f is the risk-free rate, w is the weights vector, and Σ_{real} is the variance of the realized portfolio returns;

- the *Certainty-equivalent (CEQ) return*, defined as the risk-free rate that an investor is willing to accept rather than adopting a particular risky portfolio strategy:

$$CEQ = (\mu_{real} - r_f)^T w - \frac{\gamma}{2}(w^T \Sigma_{real} w), \tag{31}$$

where γ is the risk aversion.

We compare the value of the two measures obtained with the robust optimization method[21] (both with the lower-upper bounds on the weights and with the no-short sales constraints) to those obtained by implementing three simpler strategies:

- *Index Buy-and-Hold*: this strategy considers the returns that would have been achieved by investing directly on the Euro Stoxx 50 Index on 31 March 2011 and selling it on 17 June 2011;[22]
- the $1/N$ *Portfolio Strategy*: this strategy, proposed by De Miguel, Garlappi and Uppal (2009), consists in selecting a portfolio with all the N stocks in the index, by assigning weight $1/N$ to each one;
- a *momentum strategy*: this strategy consists of selecting only the G stocks with positive returns in the previous year and allocating a portfolio with weights $1/G$ on those stocks.

Table 2 shows the out-of-sample annualized Sharpe Ratios and Certainty-Equivalent returns obtained in the different cases.

From Table 2 we can see that the robust portfolio selected when imposing lower and upper bounds shows the best out-of-sample performance with respect to both criteria, giving the highest Sharpe Ratio and the highest certainty-equivalent-return.

8 Conclusions

In this paper we propose a methodology which extends the forward-looking appraoch of Kostakis, Panigirtzoglou and Skiadopoulos (2010) to the case of a portfolio consisting

[21] We set the target minimum annual expected return in our optimization problem equal to 0.10, as in Table 1.

[22] Note that the Euro Stoxx 50 weights remained unvaried between these dates.

Table 2 – Sharpe Ratio for the different portfolios

	Robust Portfolio LU	Robust Portfolio NSS	Index Buy-and-Hold	Equally weighted Portfolio	G-Stocks Portfolio
Sharpe Ratio	2.67	1.82	-1.41	1.81	0.73
CEQ	0.384	0.2549	-0.2502	0.268	0.0978

of a variety of risky assets. We deal with the empirical implementation of a static asset allocation problem, using information extracted from the market prices of options written on the Euro Stoxx 50 Index and its constituents.

A key issue that arose in the generalization of this option-implied distribution approach to the case of N risky assets is related to the computation of the correlation structure among the assets, which has to be coherent with the related option-implied variances. To this aim we develop an ad-hoc procedure to integrate the historical stock returns with the option implied distributions of the stocks and the index.

The asset allocation is performed using the robust optimization technique. This methodology has been developed to cope with estimation errors in the estimates of expected returns and their error-magnification effect. Robust Optimization, in fact, considers the estimation errors for the mean returns directly in the optimization problem itself, and performs an asset allocation that is robust to these errors. The starting point of the robust asset allocation was our *best guess* vector for the asset returns (obtained here through an equilibrium argument). The methodology then accounted for the worst case scenario which could occur in a confidence region containing the best guess. We introduced two sets of constraints in the optimization problem: first, we imposed a no short-sale condition and then we constrained the portfolio weights to lie in the interval $[-1, +1]$.

In order to evaluate the results of our static asset allocation, we set a comparison among five efficient frontiers constructed using different methodologies (estimated Markovitz efficient frontier, estimated robust efficient frontier, true efficient frontier, actual Markovitz efficient frontier and actual robust efficient frontier).

We finally compare the out-of-sample performance of our robust portfolios to those of the Index Buy-and-Hold strategy and of two alternatives portfolio: the *1/N portfolio* of De Miguel, Garlappi and Uppal (2009) (in which we build a portfolio with all the N stocks in the index, by assigning weight $1/N$ to each one) and a *momentum strategy*, in which we select only the G stocks with positive returns in the previous year, and we allocate a portfolio with weight $1/G$ on these stocks. For this comparison we use two standard performance criteria: the out-of-sample Sharpe Ratio and the Certainty-Equivalent (CEQ) return for the expected utility of a mean-variance investor. In this case, our strategy leads to good results, showing the best out-of-sample performance according to both criteria. In future work, the methodology will be tested on a larger dataset to evaluate its robustness and effectiveness.

We decided to apply our methodology to a family of stocks which were the constituents of a quoted index, so that we could use the information extracted from market data on the index to infer a coherent correlation structure among the stocks. A further analysis will focus on the development of a procedure which can integrate option-implied variances with the related covariances without requiring the existence of an overall quoted index, so that the methodology will be much more flexible.

References

[1] Ait-Sahalia, Y., Lo, A.W. (2000). *Nonparametric risk management and implied risk aversion.* Journal of Econometrics 94, 9-51.

[2] Bahra, B. (1996). *Probability distributions of future asset prices implied by option prices.* Bank of England Quarterly Bulletin, 299-311, August.

[3] Barone-Adesi, G., Whaley, R.E. (1987). *Efficient analytical approximation of American option values.* Journal of Finance 42, 301-320.

[4] Bedoui, R., Hamdi, H. (2010). *Implied Risk-neutral probability density functions from option prices: a comparison of estimation methods.* Working Paper, February 2010.

[5] Ben-Tal, A., Nemirovski, A.S. (1997). *Robust truss topology design via semidefinite programming.* SIAM Journal on Optimization, 7(4):991-1016, 1997. 465-474.

[6] Berkovitz, J. (2001). *Testing density forecasts with applications to risk management.* Journal of Business and Economics Statistics 19, 465-474.

[7] Black, F., Scholes, M. (1973). *The Pricing of Options and Corporate Liabilities.* Journal of Political Economy 81 (3), 637-654.

[8] Bliss, R. R., Panigirtzoglou, N. (2004). *Option-implied risk aversion estimates.* Journal of Finance 59, 407-446.

[9] Breeden, D., Litzenberger, R. (1978). *Prices of State Contingent Claims Implicit in Option Prices.* Journal of Business 51, 621-651.

[10] Buss, A., Vilkov, G. (2011). *Option-Implied Correlation and Factor Betas Revisited.* Working Paper.

[11] Ceria, S. Stubbs, R. A. (2006). *Incorporating Estimation Errors Into Portfolio Selection: Robust Portfolio Construction.* Axioma Research Paper No. 003.

[12] De Miguel, V., Garlappi, L. and Uppal, R. (2009). *How inefficient are simple asset allocation strategies?.* Review of Financial Studies 22, 1915-1953.

[13] De Miguel, V., Garlappi, L. and Uppal, R. (2009). *Optimal Versus Naive Diversification: How inefficient is the 1/N Portfolio Strategy?.* Review of Financial Studies 22, 1915-1953.

[14] Hamilton, J. D. (1994). *Time Series Analysis.* Princeton University Press.

[15] Idzorek, T. (2004). *A Step-by-Step Guide to the Black-Litterman Model.* Axioma Research Paper No. 003.

[16] Jagannathan, R., Ma, T. (2003). *Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps.* The Journal of Finance. August 2003

[17] Kostakis, A., Panigirtzoglou, N. and Skiadopoulos, G. (2010). *Asset Allocation with Option-Implied Distributions: A Forward-Looking Approach.* Working Paper, February 2010.

[18] Markovitz, H. (1952). *Portfolio selection.* The Journal of Finance, VII(1):77-91. March 1952.

[19] Meucci, A. (2010). *The Black-Litterman Approach: Original Model and Extensions.* The Encyclopedia of Quantitative Finance, Wiley (2010).

[20] Mills, T.C., Markellos, R.N. (2008). *The econometric modelling of financial time series.* CUP - Cambridge University Press, 2008, third edition, 59.

[21] Mutapcic, A., Boyd, S. (2009). *Cutting-set methods for robust convex optimization with pessimizing oracles.* Optimization Methods and Software, Vol. 24, No. 3, June 2009, 381-406.

[22] $Riskmetrics^{TM}$ (1996). $Riskmetrics^{TM}$ *- Technical Document.* Fourth edition, 1996.