LAURENT CANTALUPPI*

The Reverse Optimization

1. Introduction

The problem of finding mean/variance optimal portfolios has first been formulated by Markowitz in his famous paper. The problem can be mathematically expressed as a quadratic-programming problem (QP), i.e., a problem with a quadratic objective function and linear restrictions. The input parameters of the problem are the estimated returns of the assets, the covariances of the assets and the target return of the portfolio. The result of the QP consists in the weights of the assets in the mean/variance optimal portfolio. One of the main difficulties in applying mean/variance optimization lies in the input parameters. In particular, estimating future returns with a certain confidence is a difficult task, and the result of the QP is very sensitive to these estimated returns.

This leads to the formulation of the reverse problem. Given a portfolio structure, i.e., portfolio weights and given the covariances between the assets, what are the (implied) returns that make this portfolio mean/variance optimal? Unfortunately this problem does not have a unique solution. The properties of these implied returns, however, give interesting insight into the methodology. In particular we show that the implied returns are in linear relationship with the marginal contributions to risk of the assets. As a consequence, the ranking of the implied returns remains constant for all solutions. This is a central property of the reverse optimization model, since it allows checking the consistency of the given portfolio with estimated future performance of asset classes, even when these estimates are not quantitative.

We then examine two major applications of the reverse optimization. In the first application we complete the input parameters of the reverse optimization problem by giving the expected return values of two asset classes, and let the other values be computed. We show how to analyze and interpret the results, taking into account the investment restrictions that are usually present. The second application is in the field of relative optimization. We first define "relative optimization" and derive some additional theoretical properties of the reverse optimization in this special case. In particular we show that the marginal contributions to risk of all asset classes are zero for the benchmark itself. In our concrete example, a benchmark is defined on a given universe of asset classes and the goal is to outperform the benchmark. No quantitative return estimates of the asset classes are available and a passive/active management strategy is followed. Asset classes that are estimated to outperform the benchmark get a higher

* I thank the referees Thomas Stucki, Markus Rudolf and Karl Frauendorfer for their stimulating comments and suggestions. Laurent Cantaluppi, Cantaluppi & Hug, Freudenbergstrasse 142, CH - 8044 Zürich, Tel: +41 - 1 - 366 66 70, E-Mail: cantaluppi@chsoft.ch.

weight than they have in the benchmark, and conversely, asset classes that are estimated to underperform the benchmark get a lower weight than they have in the benchmark. Neutral asset classes get the same weight as in the benchmark. The magnitudes of the deviations from the benchmark weights are not optimized in a mean/variance sense, but are determined by other means. We show that such a strategy can lead to a portfolio that is not consistent with the return assumptions. More precisely, the return ranking implied by the given portfolio can be in contradiction with the performance estimation of the asset classes that lead to it. We also show how to use the results of the reverse optimization in order to eliminate these inconsistencies. Finally we show that such inconsistencies cannot occur if the weights of only two asset classes are changed relative to the benchmark's weights.

2. The Mean/Variance Optimization Problem

2.1 Mathematical Formulation

Let us formulate the simplest problem of mean/variance optimization with n asset classes, using the following notation. Let w_i the weight of asset class i in the portfolio, r_i its expected return, R the expected target return of the portfolio and C the covariance matrix. The optimization problem can be written

$$\text{minimize} \quad \sqrt{\sum_{i,j} w_i C_{ij} w_j} \qquad (1)$$

$$\text{subject to} \quad \sum_i w_i = 1 \qquad (2)$$

$$\sum_i w_i r_i = R. \qquad (3)$$

In the usual case, where short positions are not allowed, the following restrictions must be added to the problem.

$$w_i \geq 0 \qquad \text{for all i.} \qquad (4)$$

Formula (1) defines the risk as the standard deviation of returns of the portfolio; equation (2) is referred to as the budget constraint and equation (3) as the target return constraint. With this formulation the input parameters are the expected returns of the asset classes r_i, the covariance matrix C, and the target return for the portfolio R. The result of the optimization consists in the weights w_i of the asset classes in the optimal (efficient) portfolio. Other restrictions that are normally part of the problem are neglected in this formulation. We will show later on how to take them into account.

2.2 Input Parameters

Much has been written about the difficulties in setting the input parameters of the optimization problem. They can be summarized as follows:

- The number of input parameters is very large; in particular the size of the covariance matrix is a quadratic function of the number of asset classes. For example a problem with 20 asset classes requires the input of 230 parameters!
- The estimation of returns is an extremely difficult problem.
- The sensitivity of the solution to variations in returns is high; it actually is 10 times higher than the sensitivity to the risk estimations, and 20 times higher than the sensitivity to the correlation estimations[1].

In practical situations using a covariance matrix based on the historical return time series of the asset classes can reduce the difficulty due to the number of input parameters. Unfortunately the difficulty in estimating future returns cannot be avoided so easily, and non-quantitative return estimates cannot be used directly.

3. The Reverse Optimization Problem

3.1 Motivation

The difficulties in estimating the input parameters for the mean/variance optimization problem motivates the study of the following reverse problem: given the portfolio structure and the (historical) covariance matrix, what are the expected returns that make this portfolio mean/variance optimal. This is known as the reverse optimization problem. In fact this is a more natural formulation than the direct mean/variance optimization, since a portfolio is always given (as the existing client portfolio or normal portfolio) and the estimated returns are often only available with large estimation errors, or are not available at all. The goal of the reverse optimization is to find out if the given allocation is consistent with the (not necessarily quantified) performance estimates of the different asset classes. For example the implied return of a given asset class could be 20%/year, and such a high expected return might be considered very unlikely for this asset class, and consequently necessitate a decrease in its weight in the portfolio.

3.2 Mathematical Formulation and Properties

The reverse optimization problem can be mathematically formulated the following way: the input parameters are the weights w_i of the asset classes in the portfolio to be analyzed, and the covariance matrix C, the result of the optimization consists in the expected returns of the asset classes r_i. Before writing down the equations of the problem let us change the objective function (1) to $\frac{1}{2}\sum_{i,j} w_i C_{ij} w_j$. Clearly an optimal solution for one objective function is also optimal for the other, since the risk is always positive, the square root function is a monotone increasing function as well as the multiplication by the positive factor ½. The sole purpose of this transformation is to simplify the formulation of the reverse optimization problem. In a first step we will neglect the nonnegativity constraints (4). The reverse optimization problem is formulated by writing down the optimality conditions of the direct optimization problem. First define the Lagrange function of the direct optimization problem by letting λ_B and λ_R be the Lagrange multipliers of equations (2) and (3) respectively[2]. The Lagrange function is then

$$L(w, \lambda_B, \lambda_R) = \frac{1}{2}\sum_{i,j} w_i C_{ij} w_j - \lambda_B\left(\sum_i w_i - 1\right) - \lambda_R\left(\sum_i w_i r_i - R\right) \tag{5}$$

Differentiating this Lagrange function with respect to w_i and setting the result equal to zero gives the following optimality conditions

$$\lambda_B + \lambda_R r_i = (Cw)_i \qquad \text{for all i.} \tag{6}$$

Differentiating the Lagrange function with respect to λ_B and λ_R simply gives back the budget and return constraints (2) and (3). One more optimality condition can be derived from the financial properties of the problem. We are only interested in the "upper part" of the efficient frontier, i.e., the part that has a nonnegative slope. This nonnegativity is reflected in the Lagrange multiplier λ_R, so that we can add the condition

$$\lambda_R \geq 0. \tag{7}$$

Equations (6) define a system of n equations and n + 2 variables (the n implied returns r_i and the Lagrange multipliers λ_B and λ_R) that has either no feasible solution or infinitely many solutions. This should not be surprising, because if the asset weights w_i are optimal for a given set of expected returns r_i in the direct optimization problem, the same asset weights w_i are also optimal for a set of expected returns $A + Br_i$, where $B > 0$. Such a transformation is simply a linear transformation in the risk/return graph, so that a portfolio that was on the efficient frontier before the transformation remains on it after the transformation. We find these two degrees of freedom in the reverse optimization problem, in the form of two more variables than equations. Formally, suppose that the two sets of Lagrange multipliers λ_B^1, λ_R^1 (respectively λ_B^2, λ_R^2) define the implied returns r_i^1 (respectively r_i^2). Subtracting equation (6) with the second set of multipliers from the same equation (6) with the first set of multipliers, and rearranging, we get

$$r_i^2 = \frac{\lambda_R^1}{\lambda_R^2} r_i^1 + \frac{\left(\lambda_B^1 - \lambda_B^2\right)}{\lambda_R^2} \qquad \text{for all i,} \qquad (8)$$

showing the linear relationship between the implied returns r^1 and r^2 as a function of the Lagrange multipliers.

Notice that the implied returns r_i are the returns that would be the input of the direct optimization problem (1) – (3), i.e., the expected returns in the **base currency** of the portfolio, the algorithm has no way of separating these total returns into asset and currency returns!

3.3 Invariant Ranking of the Implied Returns

The value $(Cw)_i$ is the derivative of the objective function (risk) with respect to the weight w_i, i.e., the marginal contribution to risk of asset class i. A one-unit increase in the weight w_i increases the portfolio risk by $(Cw)_i$. Equations (6) tell us that there is a linear relationship between the implied return of an asset class and the marginal contribution to risk of that asset class. But since the return is a linear function, the return of an asset class is also its marginal contribution to return. It is then intuitively clear that the marginal contribution to risk and the marginal contribution to return of the asset classes must be in the same linear relationship at optimality. This relationship is directly related to the slope of the efficient frontier at the point under scrutiny (the given portfolio, which is optimal). If they were not, the expected return of the given portfolio could be increased and its risk simultaneously decreased, contradicting its optimality. This could be achieved by increasing the weight of an asset class with relatively low contribution to risk and simultaneously decreasing the weight of an asset class with relatively high contribution to risk. As a consequence the ranking of the implied returns is completely determined by the values $(Cw)_i$. If an asset class i has a lower marginal contribution to risk than an asset class j, i.e., $(Cw)_i < (Cw)_j$, we must have that $r_i < r_j$. This is a key property of the reverse optimization, which even though it does not produce a unique solution, gives solutions that obey a strict ranking rule. This fact opens the path towards applications of mean/variance optimization without the need for quantitative return estimates.

3.4 Influence of Restrictions

Remember, however, that equations (6) are only valid if the weights of the asset classes i are unconstrained. Taking into consideration the non-negativity constraints (4) will help avoiding mistakes in the interpretation of the implied returns. Let us associate a Lagrange multiplier λ_i with each equation in (4), we get instead of (6) the following equations

$$\lambda_B + \lambda_R r_i + \lambda_i = (Cw)_i \qquad \text{for all i,} \qquad (9)$$

$$\text{with} \quad \lambda_i \geq 0. \qquad (10)$$

If the weight of the asset class i is not at its lower bound, i.e., if $w_i > 0$, then $\lambda_i = 0$, and we have the same result as before. If $w_i = 0$, however, then $\lambda_i > 0$, and (9) can be interpreted as follows: the implied return of asset class i is at most in the same linear relationship to its marginal contribution to risk as that of the other asset classes. If an asset class has a zero weight, its implied return must be interpreted as the largest expected total return, such that the weight of this asset class in the optimal portfolio is zero. In other words, if the expected total return of the asset class is smaller than the implied return, the resulting weight is zero, and if the expected total return is larger, the optimal weight is positive. Therefore if an optimization is performed, and then a reverse optimization with the resulting portfolio, the implied returns are generally not identical with the original expected total returns. Actually we will get identical returns only if the optimization was performed without any restrictions, or if no restriction is active for the optimal portfolio. This result can be easily generalized to the case of lower and upper bounds for each asset class.

4. Reverse Optimization with Supplementary Return Input

4.1 Description

As we have seen before, the reverse optimization cannot deliver absolute implied returns without supplementary input. This fact is somewhat disappointing; on the other hand the constant ranking property of the implied returns shows that the solutions are not completely arbitrary. As we already mentioned, the optimality conditions given by equations (6) define a linear system with n equations and $n+2$ variables. This suggests that it could have one unique solution if the estimated returns of two asset classes were given as supplementary input. Although estimating returns was to be avoided in the first place, this can be considered as a substantial progress, since return estimates might be available with some confidence for particular asset classes, such as money market or bond investments in the base currency of the investor. Remember, however, that the system of equations (6) is only valid if no restrictions are present; estimating the return of an asset class that has a zero weight in the given portfolio does not help!

What happens if the input returns violate the implied ranking? We simply get no feasible solution from (6) and (7); otherwise we get one unique solution. Moreover, two sets of implied returns computed from different input returns are in a linear relationship as described in Section 3.2. Once the implied returns are computed, the plausibility of the result can be analyzed. The next section provides a numerical example of such an analysis.

4.2 Numerical Example

Table 1 shows an example for a CHF portfolio that is invested internationally. The asset classes and their respective weights in the portfolio are given in the first two columns. In the first reverse optimization the expected total returns are set for the money market and bond market in CHF at values of 3% and 5% respectively. The Swiss stock market must obtain a return of 17.44% to "justify" ist weight of 20%. The money market in FRF remains at a weight of 0% as long as its expected return remains under 5.25%. For the second reverse optimization the total returns are set for the money market and bond market in CHF at values of 1% and 6% respectively. Observe that all implied returns for this second run can be computed by multiplying the implied returns of the first run by 2.5 and then subtracting 6.5, they are in a linear relationship. As a consequence the ranking of the implied returns is the same for both runs, in accordance with the theory. If we try to break this ranking, as is the case in the third reverse optimization, we do not get any feasible solution. This fact says that if the expected return of the CHF money market is higher than the expected return of the CHF bond market, it cannot be optimal to own more bonds than cash in CHF, considering that cash has a lower risk than bonds, and a similar correlation pattern. In other words the marginal contribution to risk of cash is lower than that of bonds. The expected returns must then be in the same relationship, otherwise we can increase the expected return of the portfolio

Table 1: Implied Expected Returns with Two Given Expected Returns

Asset Class	Portfolio weights	Given returns	Implied returns	Given returns	Implied returns	Given returns	Implied returns
Stocks France	0		17.34		36.86		
Stocks Germany	0		14.76		30.41		
Stocks Japan	0		10.15		18.88		
Stocks Switzerland	20		17.44		37.10		
Stocks U.K.	0		15.54		32.36		N
Stocks USA	10		20.58		44.95		O
Bonds GBP	0		11.00		21.00		
Money Market GBP	0		7.69		12.71		S
Bonds DEM	0		6.98		10.95		O
Money Market DEM	0		4.73		5.32		L
Bonds FRF	15		8.06		13.65		U
Money Market FRF	0		5.25		6.63		T
Bonds CHF	35	5.00	5.00	6.00	6.00	2.00	I
Money Market CHF	20	3.00	3.00	1.00	1.00	4.00	O
Bonds USD	0		15.42		32.06		N
Money Market USD	0		13.72		27.81		
Bonds JPY	0		9.36		16.91		
Money Market JPY	0		8.45		14.64		

and simultaneously decrease its risk by increasing the weight of the cash position and decreasing the weight of the bond position.

As already mentioned in Section 2 the optimal asset weights in the direct mean/variance optimization problem (1) – (3) are extremely sensitive to the expected returns. However, the objective function itself (1) is much less sensitive to the expected returns. As a consequence the reverse optimization problem is very stable, i.e., the implied returns are a robust function of the portfolio weights. This is illustrated in Figure 1, displaying the implied return of the asset class "Stocks Switzerland" as a function of its portfolio weight, in the case where the expected returns of the money market and bond market in CHF are set at values of 3% and 5% respectively.

5. Reverse Optimization in the Relative Case

5.1 Description and Properties

We now describe an application of the reverse optimization in the field of relative optimization. Before analyzing a concrete example we will briefly explain relative optimization and derive some properties of the reverse optimization for this special case.

Mathematically the formulation of the mean/variance optimization problem in the relative case is identical with that of the absolute case (1) – (3). The difference lies in the interpretation of the input data and the results. In the absolute case the covariance matrix C is the covariance of absolute returns and in the relative case it is the covariance of returns relative to a predefined benchmark[3]. We use the term "relative return" to refer to the excess of the asset class return over the benchmark

Figure 1: Implied Return of the Swiss Stock Market as a Function of its Portfolio Weight

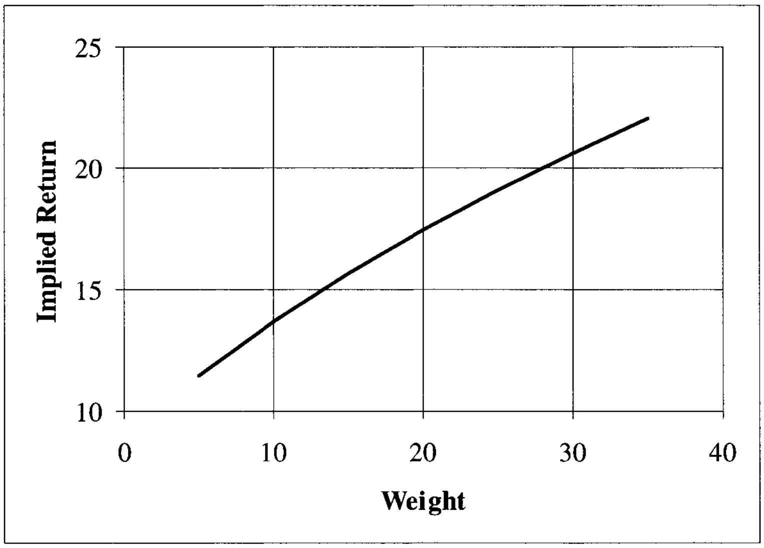

return. The riskiness of relative returns is usually called the tracking error with respect to the benchmark. Notice that the question on whether the implied returns are absolute or relative returns is irrelevant: from equations (6) we have seen that implied returns can be arbitrarily translated by a constant. By taking this constant equal to the benchmark return, we can therefore arbitrarily switch between absolute and relative returns.

Let us now derive some properties of the benchmark. Let the weights of the asset classes in the benchmark be $w^B{}_i$. Of course the tracking error of the benchmark itself is zero per definition. The benchmark is therefore an unconditional solution to the direct optimization problem (1) – (3). The differential of (1) with respect to w_i is therefore equal to zero for all asset classes i, i.e., the marginal contribution to tracking error of each asset class must also be zero, or $(Cw^B)_i = 0$. As a consequence the implied returns of the benchmark itself are undefined, since the right hand side of (6) is zero for all i, and by taking $\lambda_B = 0$ and $\lambda_R = 0$, we get that any implied returns satisfy equations (6).

5.2 The Passive / Active Management Strategy

Let us now examine in detail the following special case. A benchmark is given on a universe of asset classes. A portfolio is managed on the same universe, using the benchmark as a reference portfolio. The goal is to outperform the benchmark, while keeping a small tracking error to the benchmark. The strategy is a passive/active management. We assume that the portfolio manager has no quantitative return estimates for the asset classes, but has expectations for some asset classes relative to the benchmark. The asset classes can be divided into three groups:

- those that are estimated to outperform the benchmark,

- those that are estimated to underperform the benchmark,
- those for which no estimation is available.

A direct mean/variance optimization cannot be used, and the portfolio manager relies on other techniques to define the weights of the asset classes in the portfolio. The strategy conforms to the following rules:

- Asset classes that are estimated to outperform the benchmark get a higher weight than their respective benchmark weight.
- Asset classes that are estimated to underperform the benchmark get a lower weight than their respective benchmark weight.
- Asset classes for which no return estimation is available get the same weight as they have in the benchmark.

5.3 The Reverse Optimization under the Passive / Active Management Strategy

The following question arises: are the implied returns always consistent with the assumptions? In other words do we get "higher" implied returns for overrepresented asset classes, and "lower" implied returns for underrepresented asset classes? If this is not the case, how can we improve the given portfolio? Of course the magnitude of the implied returns is not known, but remember that their ranking is constant over all solutions.

Suppose that the implied returns are not consistent with our assumptions. In other words, the implied return of at least one overrepresented asset class i is lower than the implied return of at least one underrepresented asset class j. The allocation is then inconsistent with the performance expectations of the asset classes, let us see why. The marginal contribution to tracking error of asset class i is lower than the marginal contribution to tracking error of asset class j. Thus we can decrease the tracking error of the portfolio by increasing the weight of asset class i and simultaneously decreasing the weight of asset class j. But this will also increase the expected return of the portfolio, since by assumption, asset class i is expected to outperform the benchmark and asset class j is expected to underperform it. We have therefore improved the portfolio in the mean/variance sense, since we decreased its tracking error while increasing its expected return. This process can be repeated until the portfolio is consistent with the performance expectations of the asset classes. We show an example in the next section.

5.4 Numerical Example

Consider a USD portfolio invested in stocks in the following markets: Australia, Austria, Canada, France, Germany, Italy, Japan, Netherlands, Spain, Switzerland, United Kingdom, and United States. These asset classes as well as their respective weights in the benchmark are given in Table 2. The classes "Stocks Australia" and "Stocks Switzerland" are expected to underperform the benchmark and are given a correspondingly lower weight. Conversely, the classes "Stocks Germany" and "Stocks Japan" are expected to outperform the benchmark and are given a correspondingly higher weight. No expectation is available for the other classes and their weights are identical to their respective benchmark weights. The implied returns are computed as marginal contributions to risk. Notice that the marginal contribution to risk of "Stocks Australia" is higher than that of "Stocks Germany", and as a consequence the implied return of "Stocks Australia" is higher than that of "Stocks Germany", which is in contradiction with our expectations, since we estimated "Stocks Australia" to underperform the benchmark and "Stocks Germany" to outperform it. Further decreasing the weight of "Stocks Australia" from 2% to 0% and increasing the weight of "Stocks Germany" from 10% to 12% will decrease the tracking error of the portfolio, and increase its expected return, giving a better portfolio in the mean/variance sense. Computing the implied returns for this modified portfolio shows

Table 2: Implied Expected Returns Computed as Marginal Contributions to Risk in the Case of Benchmarking

Asset Class	Benchmark weights	Portfolio weights	Implied returns	Portfolio weights	Implied returns
Stocks Australia	3	**2**	0.10	**0**	–0.03
Stocks Austria	2		–0.05		0.06
Stocks Canada	5		–0.09		–0.11
Stocks France	7		–0.07		–0.02
Stocks Germany	8	**10**	–0.11	**12**	0.05
Stocks Italy	4		0.14		0.06
Stocks Japan	18	**22**	0.56	**22**	0.44
Stocks Netherlands	6		–0.02		0.07
Stocks Spain	4		–0.06		–0.13
Stocks Switzerland	8	**3**	–0.61	**3**	–0.53
Stocks U.K.	12		–0.07		–0.06
Stocks USA	23		–0.13		–0.12

consistency with our assumptions, i.e., the two outperforming classes "Stocks Germany" and "Stocks Japan" have a higher implied return (marginal contribution to tracking error) than the two underperforming classes "Stocks Australia" and "Stocks Switzerland".

Notice that the marginal contributions to risk are local linear approximations of a quadratic function. The weight adjustments on the basis of these numbers must therefore be performed with successive small steps. Sharper adjustments could lead to undesired results.

5.5 Special Case: Weight Change for Two Asset Classes only

We now examine the case where we have exactly one asset class that is expected to outperform the benchmark and one asset class that is expected to underperform the benchmark. We consider the weights of the resulting portfolio as the differences to the benchmark weights, i.e., we write $w = w^B + \Delta$. With this notation, equations (6) of the reverse optimization problem can be reformulated as follows

$$\lambda_B + \lambda_R r_i = (C\Delta)_i \qquad \text{for all } i \tag{11}$$

since $(Cw^B)_i = 0$.

Since the portfolio weights differ from the benchmark weights for two asset classes only, the vector Δ has all its components equal to zero except say $\Delta_i = \delta$ and $\Delta_j = -\delta$, with $\delta > 0$. Subtracting equation j from equation i in (11), we get

$$\begin{aligned}\lambda_R (r_i - r_j) &= (C\Delta)_i - (C\Delta)_j \\ &= \delta(C_{ii} - C_{ij} - C_{ji} + C_{jj}) \geq 0,\end{aligned}$$

since C is positive semi-definite.

The overrepresented asset class i always gets a higher implied return than the underrepresented asset class j. Notice, however, that nothing can be said about the magnitude of $(C\Delta)_k$ for $k \neq i, k \neq j$. This result is always consistent with our assumptions, since the portfolio was constructed with the expectation that asset class i would outperform the benchmark and asset class j would underperform the benchmark. The reverse optimization is of no particular help in this case, since it does not give any indication about the optimal magnitude of the weight change δ. This must be determined by setting it so that an acceptable tracking error is achieved.

6. Conclusion

The reverse optimization algorithm can be a powerful mean/variance analysis tool, especially when no quantitative return expectations (or incomplete return expectations) are available. It also helps providing a better understanding of the direct mean/variance optimization. However it has some limitations that the analyst must understand if he does not want to get to the wrong conclusion. In particular the role of restrictions is absolutely essential. The most important feature is the invariant ranking of the implied returns. This feature helps checking the consistency of a solution (portfolio) with possibly non-quantitative return estimates. Of course the use of reverse optimization does not exclude the use of direct optimization. In fact the combination of both provides a superior analysis tool.

Footnotes

[1] See CHOPRA and ZIEMBA (1993) for details.
[2] For a comprehensive treatment of optimization problems, and in particular Lagrange functions, see GILL et al (1981).
[3] For a detailed description of the relative optimization problem, see for example RUDD (1980), TREYNOR and BLACK (1973), or LEIBOWITZ et al (1993).

References

CHOPRA, V. and W. ZIEMBA (1993): "The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice", Journal of Portfolio Management, Spring, pp. 6–11.

GILL, P. E., W. MURRAY and M. H. WRIGHT (1981): Practical Optimization, Academic Press.

LEIBOWITZ, M. L., L. N. BADER and S. KOGELMAN (1993): "Optimal Portfolios Relative to Benchmark Allocation", The Journal of Portfolio Management, Summer, pp. 18–29.

RUDD, A. (1980): "Optimal Selection of Passive Portfolios", Financial Management, Spring, pp. 57–66.

TREYNOR, J. and F. BLACK (1973): "How to Use Security Analysis to Improve Portfolio Selection", Journal of Business, January, pp. 66–86.