Lecture 9: Benchmark tracking problems

Prof. Dr. Svetlozar Rachev

Institute for Statistics and Mathematical Economics
University of Karlsruhe

Portfolio and Asset Liability Management

Summer Semester 2008

Copyright

The material is based on the text-book:

Svetlozar T. Rachev, Stoyan Stoyanov, and Frank J. Fabozzi

Advanced Stochastic Models, Risk Assessment, and Portfolio Optimization: The Ideal Risk, Uncertainty, and Performance Measures

John Wiley, Finance, 2007

Prof. Svetlozar (Zari) T. Rachev
Chair of Econometrics, Statistics
and Mathematical Finance
School of Economics and Business Engineering
University of Karlsruhe
Kollegium am Schloss, Bau II, 20.12, R210
Postfach 6980, D-76128, Karlsruhe, Germany
Tel. +49-721-608-7535, +49-721-608-2042(s)
Fax: +49-721-608-3811
http://www.statistik.uni-karslruhe.de

Introduction

An important problem for fund managers is comparing the performance of their portfolios to a benchmark. The benchmark could be:

- a market index or
- any other portfolio or
- liability measure in the case of defined benefit pension plans.

There are two types of strategies that managers follow: active or passive.

- An *active portfolio strategy* uses available information and forecasting techniques to seek a better performance. It is an expectation about the factors that could influence the performance of an asset class. The goal of an active strategy is to outperform the benchmark after management fees by a given number of basis points.
- A *passive portfolio strategy* involves minimal expectational input and instead relies on diversification to match the performance of some benchmark. A passive strategy, commonly referred to as indexing, assumes that the marketplace will reflect all available information in the price paid for securities.

Introduction

- There are various strategies for constructing a portfolio to replicate the index but the key in these strategies is designing a portfolio whose tracking error relative to the benchmark is as small as possible.
- *Tracking error* is the standard deviation of the difference between the return on the replicating portfolio and the return on the benchmark.

$\Rightarrow$ The benchmark tracking problem can be formulated as an optimization problem.

Introduction

- We'll consider the benchmark tracking problem from a very general viewpoint, replacing the traditional tracking error measures by a general functional satisfying a number of axioms.
- We call this functional a metric of relative deviation because it calculates the relative performance of the portfolio to the benchmark.
- Our approach is based on the universal methods of the theory of probability metrics. As a result, the optimization problems which arise are a significant generalization of the currently existing approaches to benchmark tracking.

The tracking error problem

- In lecture 8, the optimal portfolio problems have one feature in common in that the risk measure, or the dispersion measure, concerns the distribution of portfolio returns without any reference to another portfolio.
- In contrast, benchmark-tracking problems include a benchmark portfolio against which the performance of the managed portfolio will be compared.
- The arising optimization problems include the distribution of the *active portfolio return* defined as the difference $r_p - r_b$ in which r_p denotes the return of the portfolio and r_b denotes the return of the benchmark.
- If the active return is positive, this means that the portfolio outperformed the benchmark and, if the active return is negative, then the portfolio underperformed the benchmark.

The tracking error problem

- In the ex-post analysis, we observe a specified historical period in time and try to evaluate how successful the portfolio manager was relative to the benchmark.
- In this case, there are two time series corresponding to the observed portfolio returns and the observed benchmark returns.
- A measure of the performance of the portfolio relative to the benchmark is the average active return, also known as the portfolio alpha and denoted by α_p.

The tracking error problem

- Alpha is calculated as the difference between the average of the observed portfolio returns and the average of the observed benchmark returns,

$$\hat{\alpha}_p = \bar{r}_p - \bar{r}_b,$$

where

$\hat{\alpha}_p$	denotes the estimated alpha
$\bar{r}_p = \frac{1}{k}\sum_{i=1}^{k} r_{pi}$	denotes the average of the observed portfolio returns $r_{p1}, r_{p2}, \ldots, r_{pk}$
$\bar{r}_b = \frac{1}{k}\sum_{i=1}^{k} r_{bi}$	denotes the average of the observed benchmark returns $r_{b1}, r_{b2}, \ldots, r_{bk}$

The tracking error problem

- A widely used measure of how close the portfolio returns are to the benchmark returns is the standard deviation of the active return, also known as tracking error.
- When it is calculated using historical observations, it is referred to as the ex-post or backward-looking tracking error.
- If the portfolio returns are equal to the benchmark returns in the specified historical period, $r_{pi} = r_{bi}$ for all i, then the observed active return is equal to zero and, therefore, the tracking error will be equal to zero.

$\Rightarrow$ The closer the tracking error to zero, the closer the risk profile of the portfolio matches the risk profile of the benchmark.

The tracking error problem

- In the ex-ante analysis, portfolio alpha equals the mathematical expectation of the active return,

$$\begin{aligned} \alpha_p &= E(r_p - r_b) \\ &= w'\mu - Er_b, \end{aligned} \tag{1}$$

  where $r_p = w'X$ in which w denotes the vector of portfolio weights, X is a random vector describing the future assets returns, and $\mu = EX$ is a vector of the expected assets returns.

- The tracking error equals the standard deviation of the active return,

$$TE(w) = \sigma(r_p - r_b),$$

  where $\sigma(Y)$ denotes the standard deviation of the r.v. Y.

$\Rightarrow$ Tracking error in this case is referred to as ex-ante or forward-looking tracking error.

The tracking error problem

- If the strategy followed is active, then the goal of the portfolio manager is to gain a higher alpha at the cost of deviating from the risk profile of the benchmark portfolio; that is, the manager will accept higher forward-looking tracking error.

$\Rightarrow$ Active strategies are characterized by high alphas and high forward-looking tracking errors.

- If the strategy is passive, then the general goal is to construct a portfolio so as to have a forward-looking tracking error as small as possible in order to match the risk profile of the benchmark. As a consequence, the alpha gained is slightly different from zero.

$\Rightarrow$ Passive strategies are characterized by very small alphas and very small forward-looking tracking errors.

The tracking error problem

- Strategies which are in the middle between the active and the passive ones are called enhanced indexing.[1]
- In following such a strategy, the portfolio manager constructs a portfolio with a risk profile close to the risk profile of the benchmark but not identical to it.
- Enhanced indexing strategies are characterized by small to medium-sized forward-looking tracking errors and small to medium-sized alphas.
- The optimal portfolio problem originating from this framework is the minimal tracking error problem.

[1] Loftus (2000) classifies the three strategies for equity portfolio strategies as follows: indexing, 0 to 20 basis points; 50 to 200 basis points; and, active management, 400 basis points and greater.

The tracking error problem

- The *minimal tracking error problem* has the following form

$$\begin{array}{rl} \min\limits_{w} & \sigma(r_p - r_b) \\ \text{subject to} & w'e = 1 \\ & w'\mu - Er_b \geq R_* \\ & w \geq 0, \end{array} \tag{2}$$

  where R_* denotes the lower bound of the expected alpha.

- Its structure is very similar to the mean-variance optimization problems. The difference is that the active portfolio return $r_p - r_b$ is used instead of the absolute portfolio return r_p.
- The goal is to find a portfolio which is closest to the benchmark in a certain sense, while setting a threshold on the expected alpha. In this case, the "closeness" is determined by the standard deviation.

The tracking error problem

- By varying the limit R_*, we obtain the entire spectrum from passive strategies (R_* close to zero), through enhanced indexing (R_* taking medium-sized values), to active strategies (R_* taking from medium-sized to large values).
- The set of the optimal portfolios generated by problem (2) is the set of efficient portfolios which, if plotted in the plane of expected alpha versus tracking-error, form the corresponding efficient frontier. (See *Figure 1* for illustration.)

The tracking error problem

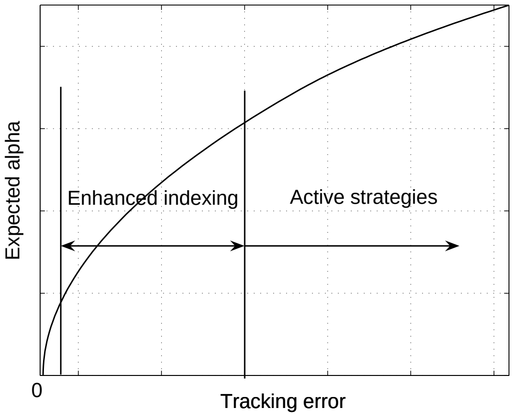

Figure 1. The efficient frontier generated from the minimal tracking error problem. The passive strategies are positioned to the left of enhanced indexing strategies.

The tracking error problem

- If the investment universe is the same as or larger than the universe of the benchmark portfolio, then the global minimum tracking error is equal to zero. Then the optimal portfolio coincides with the benchmark portfolio.
- The global minimum tracking error portfolio represents a typical passive strategy.
- Increasing the lower bound on the expected alpha we enter the domain of enhanced indexing.
- Increasing further R_* leads to portfolios which can be characterized as active strategies.

The tracking error problem

- Recall that a serious disadvantage of the standard deviation is that it penalizes in the same way the positive and the negative deviations from the mean of the r.v.
- Then the tracking error treats in the same fashion the underperformance and the outperformance, while our attitude towards them is asymmetric.
- We are inclined to pay more attention to the underperformance.

$\Rightarrow$ From an asset management perspective, a more realistic measure of “closeness” should be asymmetric.

The tracking error problem

- Our aim is to restate the minimal tracking error problem in the more general form

$$\begin{array}{rl} \min\limits_{w} & \mu(r_p, r_b) \\ \text{subject to} & w'e = 1 \\ & w'\mu - Er_b \geq R_* \\ & w \geq 0, \end{array} \tag{3}$$

  where $\mu(X, Y)$ is a measure of the deviation of X relative to Y.

- Due to this interpretation, we regard μ as a functional which metrizes relative deviation and we call it a relative deviation metric or simply, r.d. metric.

The tracking error problem

- If the portfolio r_p is an exact copy of the benchmark, i.e. it contains exactly the same stocks in the same amounts, then the relative deviation of r_p to r_b should be close to zero.[2]

The converse should also hold but, generally, could be in a somewhat weaker sense.

- If the deviation of r_p relative to r_b is zero, then the portfolio and the benchmark are indistinguishable but only in the sense implied by μ. They may, or may not, coincide with probability 1.

[2] It would not be equal to zero due to transaction costs.

Relation to probability metrics

- The benchmark-tracking problem given by (3) belongs to a class of problems in which distance between random quantities is measured.
- In order to gain more insight into the properties that μ should satisfy, we relate the benchmark-tracking problem to the theory of probability metrics, explained in *Lecture 3*.

Relation to probability metrics

- A functional which measures the distance between random quantities is called a probability metric.
- These random quantities can be random variables (such as the daily returns of equities, the daily change of an exchange rate, etc.) or stochastic processes (such as a price evolution in a given period), or much more complex objects (such as the daily movement of the shape of the yield curve).
- The probability metric is defined through a set of axioms, given in *Lecture 3*; that is, any functional which satisfies these axioms is called a probability metric.

Relation to probability metrics

- From the standpoint of the theory of probability metrics, the benchmark-tracking problem given by (3) can be viewed as an approximation problem.
- So that we are trying to find a r.v. r_b in the set of feasible portfolios which is closest to the r.v. r_b and the distance is measured by the functional μ.
- This functional should satisfy the properties stated in *Lecture 3*, or some versions of them, in order for the problem to give meaningful results.
- Let us reexamine the set of properties Property 1 - Property 3 to verify if some of them can be relaxed while application to a specific problem.

Relation to probability metrics

- *Property 1 and Property 3, we leave intact.* The reason is that anything other than Property 1 is just nonsensical and Property 1 together with Property 3 guarantee nice mathematical properties, such as continuity of μ. Property 3 alone makes sense because of the interpretation that we are measuring distance.
- *Property 2 can be dropped.* The rationale is that, in problem (3) the assumption of asymmetry is a reasonable property because of our natural tendency to be more sensitive to underperformance than to outperformance relative to the benchmark portfolio.

Relation to probability metrics

The nature of the problem may require the additional properties.

- Let us consider two equity portfolios with returns X and Y.
- Suppose that we convert proportionally into cash 100 a % in total of both portfolios where $0 \leq a \leq 1$ stands for the weight of the cash amount.
- As a result, the two portfolios returns scale down to $(1-a)X$ and $(1-a)Y$ respectively.
- Since both random quantities get scaled down by the same factor, we may assume that the distance between them scales down proportionally. Actually, we assume that the distance scales down by the same factor raised by some fixed power s,

$$\nu(aX, aY) = a^s \nu(X, Y) \text{ for any } X, Y \text{ and } a, \geq 0.$$

- If $s = 1$, then the scaling is proportional.

Relation to probability metrics

- The reason we are presuming a more general property is that different classes of r.d. metrics will originate and depending on s they may have different robustness in the approximation problem.
- This property we call positive homogeneity of degree s. It is similar to the homogeneity property of ideal probability metrics.

Relation to probability metrics

- As a next step, consider an equity with return Z which is independent of the two portfolios returns X and Y.
- Suppose that we invest the cash amounts into equity Z. The returns of the two portfolios change to $(1 - a)X + aZ$ and $(1 - a)Y + aZ$, respectively, where $0 \leq a \leq 1$ denotes the weight of equity Z in the portfolio.

How does the distance change?

- Certainly, there is no reason to expect that the distance should increase. It either remains unchanged or decreases.

Relation to probability metrics

- In terms of the functional, we assume the property

$$\nu(X + Z, Y + Z) \leq \nu(X, Y)$$

  for all Z independent of X, Y.
- Any functional ν satisfying this property, we call weakly regular, a label we borrow from the probability metrics theory. In fact, this is the weak regularity property of ideal probability metrics.
- If the distance between the two new portfolios remains unchanged for any Z irrespective of the independence hypothesis, then we say that ν is translation invariant.

Relation to probability metrics

- Note that if the positive homogeneity property and the weak regularity property hold, then the inequality

$$\nu((1-a)X + aZ, (1-a)Y + aZ) \leq \nu(X, Y), \qquad a \in (0,1) \quad (4)$$

  holds as well and this is exactly the mathematical expression behind the conclusion in the example.
- While the weak regularity property may seem more confined than postulating directly (4), we assume it as an axiom because (4) is tied to the interpretation of the random variables as return on investment.
- Suppose that this is not the case and X, Y denote the random wealth of the two portfolios under consideration and Z denotes random stock price.

Relation to probability metrics

- Furthermore, assume that the present value of both portfolios are equal and that we change both portfolios by buying one share of stock Z.
- Then the random wealth of both portfolios becomes $X + Z$ and $Y + Z$, respectively, and, because of the common stochastic factor Z, we expect the relative deviation to decrease; that is $\nu(X + Z, Y + Z) \leq \nu(X, Y)$.

$\Rightarrow$ The weak regularity property is the fundamental property we would like to impose.

Relation to probability metrics

- In order to state the last axiom, suppose that we add to the two initial portfolios other equities, such that returns of the portfolios become $X + c_1$ and $Y + c_2$, where c_1 and c_2 are some constants.
- We assume that the distance between the portfolios remains unchanged because it is only the location of X and Y that changes. That is,

$$\nu(X + c_1, Y + c_2) = \nu(X, Y)$$

  for all X, Y and constants c_1, c_2. We call this property location invariance.
- As a corollary, this property allows measuring the distance only between the centered portfolios returns.

Relation to probability metrics

- We demonstrate how such a functional ν can be constructed, for example, from a given probability metric.
- Suppose that $\mu(X, Y)$ is a given probability metric and denote by g the mapping
$$g : X \to X - EX.$$
- The mapping takes as an argument a random variable with a finite mean and returns as output the random variable with its mean subtracted, $g(X) = X - EX$.
- The mapping g has the property that shifting the random variable X does not change the output of the mapping,
$$g(X + a) = X - EX = g(X).$$

Relation to probability metrics

- Let us define the functional ν as

$$\nu(X,Y) = \mu(g(X), g(Y)). \tag{5}$$

- Thus, ν calculates the distance between the centered random variables $X - EX$ and $Y - EY$ by means of the probability metric μ.
- As a consequence, the functional ν defined in (5) is location invariant,

$$\begin{aligned} \nu(X + c_1, Y + c_2) &= \mu(g(X + c_1), g(Y + c_2)) \\ &= \mu(g(X), g(Y)) \\ &= \nu(X, Y). \end{aligned}$$

- The definition in equation (5) can be written in a more compact form without introducing an additional notation for the mapping,

$$\nu(X,Y) = \mu(X - EX, Y - EY).$$

Relation to probability metrics

- The expected return of the portfolio and some other characteristics can be incorporated into the constraint set of the benchmark-tracking problem (3).
- For example, the expected alpha constraint imposes a lower bound on the expected alpha, or the expected outperformance relative to the benchmark.

Relation to probability metrics

Let's finally define the r.d. metrics.

Any functional μ which is weakly regular, location invariant, positively homogeneous of degree s, and satisfies Property 1 and Property 3.

The structural classification of probability metrics holds for r.d. metrics as well. We distinguish between *compound*, *simple*, and *primary* r.d. metrics depending on the degree of sameness implied by the r.d. metric.

Relation to probability metrics

Now let us revisit the classical tracking error and try to classify it.

- *First*, it is a special case of the average compound metric with $p = 2$ and therefore it satisfies Property 1 - Property 3. Of course, this also means it is a compound metric, hence it implies the strongest form of sameness — in almost sure sense.
- *Second*, concerning the group of the additional axioms, it is positively homogeneous of degree 1, translation invariant, and satisfies the location invariance property.

Further, our goal is to give other examples of r.d. metrics, which are substantially different from classical tracking error, and to see their properties in a practical setting.

Examples of r.d. metrics

- A deviation measure $D(X)$ can generate a functional which is a reasonable candidate for a measure of distance in the optimization problem (3).
- For example,

$$\mu_D(X, Y) = D(X - Y) \tag{6}$$

is a translation invariant probability semimetric, homogeneous of degree 1 on condition that D is a symmetric deviation measure.

- A converse relation holds as well. That is,

$$D_\mu(X) = \mu(X - EX, 0) \tag{7}$$

is a symmetric deviation measure, where μ is a translation invariant probability metric, homogeneous of degree 1.

Examples of r.d. metrics

- If D is general deviation measure, then μ_D is a r.d. metric. In a similar way, if μ is a r.d. metric, then D_μ is a general deviation measure[3].

$\Rightarrow$ All deviation measures turn out to be spawned from the class of translation invariant r.d. metrics with degree of homogeneity $s = 1$.

- This relationship already almost completely classifies the functional μ_D arising from the deviation measure D.
- Note that μ_D is a compound metric and therefore it implies the strongest, almost sure sense of similarity. This can be seen by considering an example in which X and Y are independent and identically distributed. Then the difference $X - Y$ is a r.v. with non-zero uncertainty, hence $D(X - Y) > 0$.

[3] See the appendix of this lecture for the details

Examples of r.d. metrics

- We proceed by providing two r.d. metrics belonging to a completely different category - they both are simple and therefore the sense of similarity they imply is only up to equality of distribution functions.
- These functionals are defined through the equations

$$\theta_p^*(X,Y) = \left[\int_{-\infty}^{\infty} (\max(F_X(t) - F_Y(t), 0))^p dt\right]^{1/p}, \; p \geq 1 \qquad (8)$$

and

$$\ell_p^*(X,Y) = \left[\int_0^1 (\max(F_Y^{-1}(t) - F_X^{-1}(t), 0))^p dt\right]^{1/p}, \; p \geq 1 \qquad (9)$$

where

X and Y are zero-mean random variables,
$F_X(t) = P(X < t)$ is the distribution function of X and
$F_X^{-1}(t) = \inf\{x : F_X(x) \geq t\}$ is the generalized inverse of the distribution function.

The intuition behind (9) and (8) is the following.

- Suppose that X and Y represent the centered random return of two portfolios and that their distribution functions are as shown in *Figure 2*.
- Both functionals measure the relative deviation of X and Y using only the part of the distribution functions, or the inverse distribution functions, which describes losses.

Examples of r.d. metrics

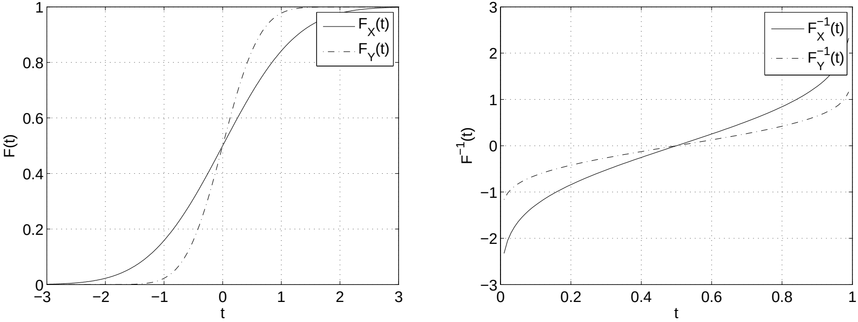

Figure 2. The distribution functions (left) and the inverse distribution functions (right) of X and Y.

Examples of r.d. metrics

- A closer look at the left plot in the figure reveals that the difference $F_X(t) - F_Y(t)$ is non-negative only for negative values of t and therefore $\theta_p^*(X, Y)$ essentially uses the information about losses contained in the distribution function. The same holds for the other functional.
- In the case where $p = 1$, then $\theta_p^*(X, Y)$ calculates the area between the two distribution functions to the left of the origin, which is exactly the same area between the inverse distribution functions to the left of $t = 1/2$.
- It is easy to notice that $\theta_1^*(X, Y) = \ell_1^*(X, Y)$ but this is, generally, not true if $p \neq 1$.

Examples of r.d. metrics

A remark for (8) and (9) on Property 1 follows as there is a subtle nuance which has to be explained.

- If the two distribution functions coincide, then both (9) and (8) become equal to zero.
- The converse statement holds as well. Suppose that the two r.d. metrics are equal to zero. Then, the distribution functions of the random variables may diverge but only in a very special way,

$$\theta_p^*(X, Y) = 0 \quad \Longrightarrow \quad F_X(t) \leq F_Y(t), \ \forall t \in \mathbb{R}.$$

See the illustration on the next slide.

Examples of r.d. metrics

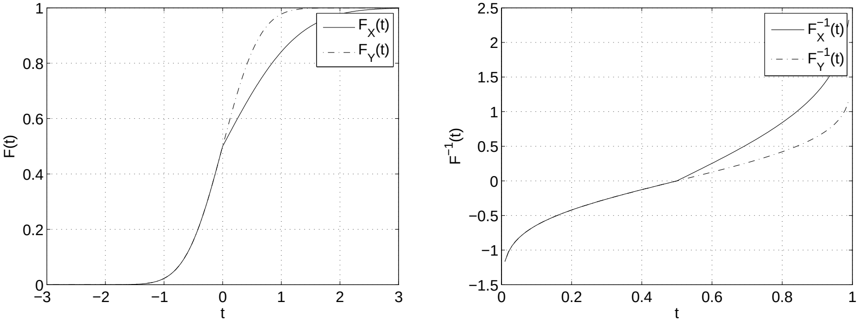

Figure 3. The distribution functions (left) and the inverse distribution functions (right) of X and Y.

Examples of r.d. metrics

- However, the inequality is impossible to hold for r.d. metrics because of the location invariance property; that is, we consider only zero-mean random variables and the inequality between the distribution functions above implies that $EX \leq EY$, hence one of the random variables may have a non-zero mean.
- As a result, if $\theta_p^*(X, Y) = 0$, then it follows that the c.d.f.s coincide for all values of the argument, $F_Y(t) = F_X(t)$, $\forall t \in \mathbb{R}$ and, therefore, the two random variables have identical probabilistic properties.
- Exactly the same argument applies to $\ell_p^*(X, Y)$.

Examples of r.d. metrics

- In summary, we showed that Property 1 holds in a stronger sense for (9) and (8).
- Not only does the distance between X and X equal zero, $\mu(X, X) = 0$, but if $\mu(X, Y) = 0$, then for these two cases, because of the location invariance property, it follows that X is equivalent to Y to the extent that their distribution functions coincide.
- However, there are examples of r.d. metrics for which the location invariance property is insufficient to guarantee this stronger identity property.

Examples of r.d. metrics

Is is possible to calculate explicitly the r.d. metrics (9) and (8)?

- The answer is negative but in some special cases, this can be done.
- Suppose that $p = 1$ and that both random variables have the normal distribution, $X \in N(0, \sigma_X^2)$ and $Y \in N(0, \sigma_Y^2)$.
- Due to the equality $\boldsymbol{\theta}_1^*(X, Y) = \ell_1^*(X, Y)$, it makes no difference which r.d. metric we choose for this calculation. Then, under these assumptions,

$$\begin{aligned} \ell_1(X, Y) &= \int_0^1 (\max(F_Y^{-1}(t) - F_X^{-1}(t), 0))dt \\ &= \frac{1}{\sqrt{2\pi}} \left|\sigma_X - \sigma_Y\right|. \end{aligned} \tag{10}$$

- In this special case, (10) is actually a primary metric because it measures the distance between the standard deviations of the portfolio return and the benchmark return. It is because we have restricted our reasoning to the normal distribution only that the simple metric ℓ_1 takes this special form. Otherwise, it is a simple r.d. metric.

Examples of r.d. metrics

- Looking more carefully at (10), we notice that the symmetry property holds due to the absolute value; that is, in this special case, $\ell_1(X, Y) = \ell_1(Y, X)$.
- This may appear striking because equation (9) is asymmetric by construction, or so it may seem.
- The symmetry property appears, again, because of the normality assumption — the left and the right tails of the distributions disagree symmetrically in this case.
- In other words, the particular form of (9) *allows for* asymmetry if the corresponding distributions are skewed. If X and Y are symmetric, then this is a fundamental limitation and the potential for asymmetry, granted by the r.d. metric, cannot be exploited.

- We can use equation (10) to illustrate a point about the relationship between the compound r.d. metrics and the minimal r.d. metrics corresponding to them.
- Using very general arguments, only the triangle inequality, we can show that equation (10) is related to the tracking error through the inequality

$$|\sigma_X - \sigma_Y| \leq \sigma(X - Y). \tag{11}$$

  It is true not only when X and Y are normal but in general.
- Equation (11) shows that if the tracking error is zero, then $|\sigma_X - \sigma_Y| = 0$ which, in the normal distribution case, means that $\ell_1(X, Y) = 0$.
- Conversely, in the normal distribution case one can find two random variables X and Y with $\sigma_X = \sigma_Y$ and, yet, the tracking error can be non-zero, $\sigma(X - Y) \neq 0$, because of the dependence between the two random variables.

Examples of r.d. metrics

Example

- If X and Y are independent, then $X - Y$ is a r.v. with the normal distribution and its standard deviation is strictly positive.
- $\ell_1(X, Y) = 0 = |\sigma_X - \sigma_Y|$ means that X and Y have the same probabilistic properties and, nevertheless, the tracking error may be strictly positive.

- Similar conclusion holds in general, when we consider compound versus simple metrics but an inequality such as (11) is guaranteed to hold between compound metrics and the minimal metrics corresponding to them.
- It is in the normal distribution case that the left side of (11) coincides with the minimal metric of the tracking error.

Numerical example

- We showed that both functionals (9) and (8) are meaningful objectives in the benchmark-tracking problem. They are very different from the classical tracking error as they are simple metrics and imply a weaker form of sameness.
- Even if they are both simple, the optimal solutions corresponding to (9) and (8) will, generally, not be the same if $p \neq 1$. This is understandable as the functionals are not identical.
- There is one important difference between them concerning Property 4. The functional $\theta_p^*(X, Y)$ is positively homogeneous of degree $1/p$ while $\ell_p^*(X, Y)$ is positively homogeneous of degree 1 irrespective of the value of p.

Numerical example

We'll provide a numerical example. Our goals are:

1. Observe the difference between the optimal solutions of the classical tracking error on the one hand and (9) and (8) on the other. In this example, the optimal solution is represented by a portfolio, the empirical c.d.f. of which is closest to the empirical c.d.f. of the benchmark as measured by the corresponding r.d. metric.
2. Examine the effect of the degree of homogeneity in the case of $\theta_p^*(X, Y)$, our expectation being that the higher the degree of homogeneity, the more sensitive $\theta_p^*(X, Y)$ is.

Our dataset includes 10 randomly selected equities from the S&P 500 universe and the benchmark is the S&P 500 index. The data cover the one-year period from December 31, 2002 to December 31, 2003.

Numerical example

- The optimization problem we solve is the benchmark-tracking problem given by (3) in which $R_* = 0$ and the objective function $\mu(r_p, r_b)$ is represented by the corresponding empirical counterpart,

$$\hat{\sigma}(r_p - r_b) \quad (12)$$

$$\hat{\boldsymbol{\theta}}^*_p(r_{p0}, r_{b0}) \quad (13)$$

$$\hat{\ell}^*_p(r_{p0}, r_{b0}) \quad (14)$$

  where the index 0 signifies that the corresponding returns are centered.

- In the case of tracking error, the sample counterpart $\hat{\sigma}$ is the sample standard deviation.

Numerical example

- In the three problems, we start from an equally weighted portfolio and then solve the optimization problems.
- Therefore, in all cases, our initial portfolio is an equally weighted portfolio of the 10 randomly selected stocks.
- The constraint set guarantees that the expected return of the optimal portfolio will not be worse than that of the benchmark.
- We compare the inverse distribution functions of the centered returns of the optimal portfolios in order to assess which problem better tracks the benchmark in terms of the distribution function.
- Note that it makes no difference whether we compare the distribution functions or the inverse distribution functions, the conclusion will not change.

Numerical example

Figure below compares the inverse distribution functions of the centered returns of the initial portfolio, the optimal portfolio of the classical tracking error problem and the optimal portfolio obtained with objective (14) in which $p = 1$.

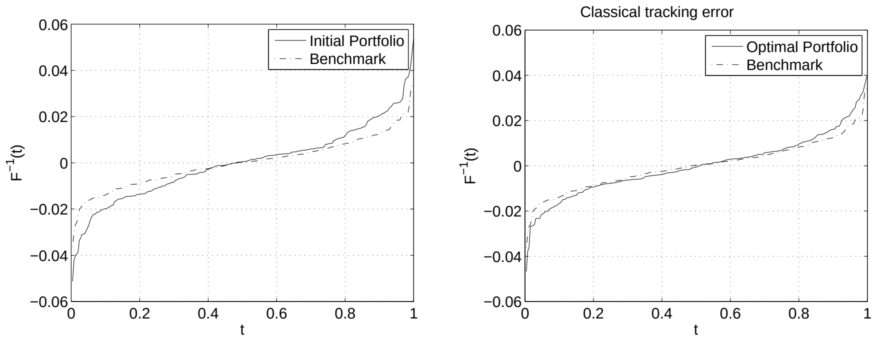

Figure 4. The inverse distribution functions of the S&P 500 index (the benchmark), equally weighted portfolio (initial portfolio) and the two optimal portfolios.

Numerical example

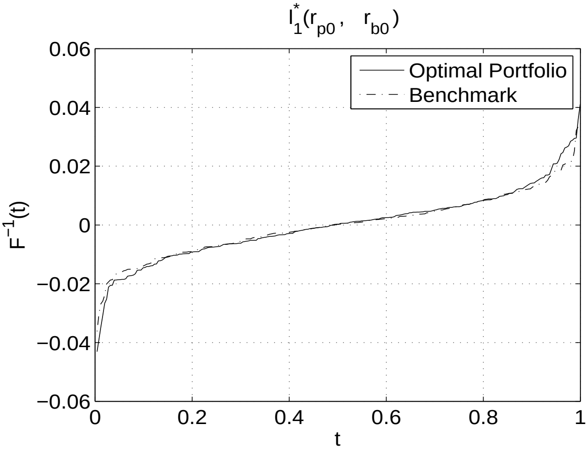

Figure 4 (cont.). The inverse distribution functions of the S&P 500 index (the benchmark), equally weighted portfolio (initial portfolio) and the two optimal portfolios.

Numerical example

- It is obvious that both optimization problems provide solutions that better track the benchmark than the trivial strategy of holding an equally weighted portfolio.
- The functional (14) does a better job at approximating the distribution function of the benchmark returns, allowing for asymmetries in the loss versus the profit part.
- The part of the inverse distribution function of the optimal solution describing losses, the one closer to $t = 0$ is closer to the corresponding part of the inverse distribution function of the benchmark returns, while this is not true for the profit part closer to $t = 1$.
- Actually, the fact that the inverse distribution function of the optimal solution is above the inverse distribution function of the benchmark returns close to $t = 1$ means that the probability of a large positive return of the optimal portfolio is larger than that of the benchmark.

Numerical example

- In order to explore the question of how the degree of homogeneity might influence the solution, we solve the tracking problem with objective (13) in which we choose $p = 1$ and $p = 10$.
- The degree of homogeneity is equal to 1 and $1/10$ respectively.
- We noted already that $\theta_1^*(r_{p0}, r_{b0}) = \ell_1^*(r_{p0}, r_{b0})$ and therefore the optimal solutions will coincide.
- The inverse distribution functions of the returns of the optimal portfolios are shown on *Figure 5*.

Numerical example

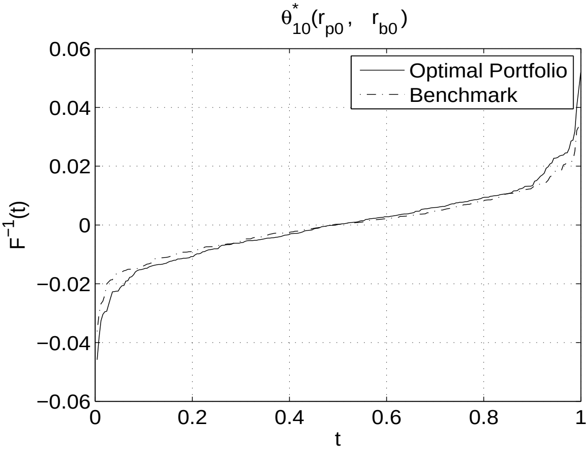

Figure 5. The inverse distribution functions of the SP500 index (the benchmark), and the optimal portfolios obtained with (13) as objective with $p = 10$.

Numerical example

- Apparently, the degree of sensitivity of the objective is directly influenced by the parameter p in line with our expectations.
- The integrand in the functional (13) measures the differences of the two distribution functions and therefore its functional values are small numbers, converging to zero in the tails.
- Holding other things equal, raising the integrand to a higher power deteriorates the sensitivity of the functional with respect to deviations in the tails of the two distribution functions.
- This observation becomes obvious when we compare the bottom plot of *Figure 4* to *Figure 5*.

Numerical example

- Generally, the optimization problems involving simple r.d. metrics may not belong to the family of convex problems because the simple r.d. metric may not appear to be a convex function of portfolio weights.
- Stoyanov, Rachev and Fabozzi (2007) show that, in particular, this holds for the minimal r.d. metrics.

Summary

- We considered the problem of benchmark-tracking.
- The classical problem relies on the tracking error to measure the degree of similarity between the portfolio and the benchmark.
- Making use of the approach of the theory of probability metrics, we extended significantly the framework by introducing axiomatically relative deviation metrics replacing the tracking error in the objective function of the optimization problem.
- We provided two examples of relative deviation metrics and a numerical illustration of the corresponding optimization problems.

Svetlozar T. Rachev, Stoyan Stoyanov, and Frank J. Fabozzi
Advanced Stochastic Models, Risk Assessment, and Portfolio Optimization: The Ideal Risk, Uncertainty, and Performance Measures
John Wiley, Finance, 2007.

Chapter 9.