CEBRIG
CENTRE EMILE BERNHEIM
DE RECHERCHE INTERDISCIPLINAIRE EN GESTION

ULB Solvay Brussels School Economics & Management

Sensitivity to measurement errors of the distance to the efficient frontier

Marie Briere, Léopold Simar, Ariane Szafarz & Anne Vanhems

This paper builds confidence intervals for the distance in the mean-variance plan between any portfolio and the Markowitz efficient frontier. The distance can be calculated in any risk-return direction chosen by the investor. To do so, we introduce random variations of inputs and outputs and estimate the frontier. We then use subsampling approximations to derive confidence intervals around the distance of portfolios to the efficient frontier. This methodology offers a novel statistical approach to mean-variance portfolio choice, which is key for asset management. We apply this approach to show that the distance between the S&P 500 index and the efficient frontier made up of all the shares in the index is significantly different from zero in all testable directions. This result adds robustness to the still controversial Roll critique of the Capital Asset Pricing Model (CAPM). In the general setup of production theory, our paper addresses the sensitivity of the estimated efficiency scores to random variations in the original inputs-outputs.

Keywords Portfolio choice, Random inputs-outputs, Directional distance, FDH estimator, Efficient frontier.

JEL Classifications C44, C12, C67, G11, G14.

CEBRIG Working Paper N°23-004
May 2023

Université Libre de Bruxelles - Solvay Brussels School of Economics and Management
Centre Emile Bernheim de Recherche Interdisciplinaire en Gestion
ULB CP114/03 50, avenue F.D. Roosevelt 1050 Brussels BELGIUM
cebrig@ulb.be - Tel : +32 (0)2/650.48.64

Sensitivity to measurement errors of the distance to the efficient frontier

Marie Brière[*]
marie.briere@amundi.com

Léopold Simar[‡]
leopold.simar@uclouvain.be

Ariane Szafarz[§]
ariane.Szafarz@ulb.be

Anne Vanhems[¶]
a.vanhems@tbs-education.fr

May 5, 2023

Abstract

This paper builds confidence intervals for the distance in the mean-variance plan between any portfolio and the Markowitz efficient frontier. The distance can be calculated in any risk-return direction chosen by the investor. To do so, we introduce random variations of inputs and outputs and estimate the frontier. We then use subsampling approximations to derive confidence intervals around the distance of portfolios to the efficient frontier. This methodology offers a novel statistical approach to mean-variance portfolio choice, which is key for asset management. We apply this approach to show that the distance between the S&P 500 index and the efficient frontier made up of all the shares in the index is significantly different from zero in all testable directions. This result adds robustness to the still controversial Roll critique of the Capital Asset Pricing Model (CAPM). In the general setup of production theory, our paper addresses the sensitivity of the estimated efficiency scores to random variations in the original inputs-outputs.

Keywords: Portfolio choice, Random inputs-outputs, Directional distance, FDH estimator, Efficient frontier

JEL Classifications: **C44, C12, C67, G11, G14**

[*]Amundi Asset Management; Paris Dauphine University; Université Libre de Bruxelles

[‡]Institut de Statistique, Biostatistique et Sciences Actuarielles (ISBA), LIDAM, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.

[§]Université Libre de Bruxelles (ULB), Solvay Brussels School of Economics and Management, CEBRIG and CERMi, Belgium.

[¶]TBS Business School, France.

1 Introduction

In modern portfolio theory (Markowitz 1952), the investor selects an optimal portfolio based on an expected utility function that depends positively on the expected return and negatively on the risk. The efficient frontier is the curve in the mean-variance plan that includes all non-dominated portfolios. Equivalently, in the mean-variance framework, any portfolio located at a significantly non-zero distance to the efficient frontier cannot be optimal for any investor. This fundamental property highlights the need for developing tests with the null hypothesis that a given portfolio belongs (i.e., is at zero distance) to the efficient frontier. Even though previous work has addressed this testing problem for specific—vertical and horizontal—distances (Basak et al. 2002, Brière et al. 2013), this paper is the first, to our knowledge, to offer a general solution applicable to any risk-return direction chosen by the investor depending on her level of risk aversion.

We start by embedding the testing problem in the larger setting of input-output spaces and taking advantage of the algorithm developed by Nalpas et al. (2017). For any given set of observed inputs and outputs, this efficient and numerically robust algorithm evaluates the financial performances of risky portfolios with nonparametric efficiency measurements. Inputs can include variance and kurtosis, while outputs can include mean, return, and skewness. In practice, portfolio characteristics are empirical estimates derived from a sample of historical returns observed over a long period, which introduces randomness in the original data. Our paper develops a convenient method to analyze variations, with confidence intervals, in performance measures relative to this randomness. The methodology is based on subsampling approximations, adapted by Politis et al. (1997) to heteroskedastic time series, having inherent time dependence.

Next, we apply the new algorithm to the portfolio selection problem and the Markowitz mean-variance efficiency frontier, where we only have empirical (random) estimators of the mean and the variance. Mean-variance efficiency is the key criterion used for testing whether a portfolio is optimal in a given investment universe. In this setting, optimality means that the distance to the efficiency frontier is not significantly different from zero, suggesting that one cannot reject that the portfolio of interest belongs to the frontier. Running tests along this line can be particularly tricky for investment universes with no risk-free asset, so that the efficient frontier is nonlinear. This situation is arguably the most realistic one, given the liquidity risks associated with cash investments during crises (Acharya et al. 2011). Several approaches have been proposed for testing mean-variance efficiency in a universe with risky assets only (Kandel 1984, Levy and Roll 2010) by adapting the tests based on a risk-free asset. Alternatively, Basak et al. (2002) suggest using the (horizontal) distance between any portfolio and its same-return counterpart on the efficient frontier, while Brière et al. (2013) exploit the Kandel and Stambaugh (1995) measure of the portfolio inefficiency to develop a test for the (vertical) distance between any portfolio and its same-volatility counterpart on the frontier. While the two tests provide complementary views on directional efficiency, their conclusions are contingent on specific (horizontal and vertical) distances. Results obtained with other directional distances might therefore challenge the generality of these conclusions. Moreover, not all portfolios of interest possess a horizontal counterpart on the efficient frontier.

Testing the mean-variance efficiency of actual portfolios in any risk-return direction is of considerable use for portfolio managers willing to a) assess the financial performance of their portfolio, and b) check their alignment with the investor's risk aversion. Our novel approach addresses both concerns. Yet there is a special portfolio, namely the market portfolio, for which the test has deep implications for finance theory. The mean-variance efficiency of the market portfolio lies indeed at the heart of the Capital asset pricing model (CAPM), the foundational pricing model of financial theory. In the CAPM framework, asset pricing crucially relies on the efficiency of the market portfolio. Despite the centrality of this theoretical result, the efficiency of the real-life proxies of the market portfolio, such as large stock market indices, is still controversial. Diverging views have been voiced since at least the Roll (1977) critique, and no consensus has emerged yet. For this reason, the application of our theoretical framework to the market portfolio goes beyond a simple numerical example as it can yield important financial implications. Testing whether the S&P 500 index corresponds to an efficient portfolio is a fundamental question in finance that touches both asset pricing and portfolio management.

Beyond theoretical concerns, the lack of options to assess the statistical significance of directional distances from a given portfolio to the efficient frontier represents a serious limitation for asset-management professionals. The new tool we propose fills the gap by informing investment advisors and fund managers on how to fine-tune portfolio mean-variance efficiency in directions corresponding to the risk-return trade-offs defined by investors.

Section 2 presents the problem and describes the statistical setting. Section 3 shows how we use subsampling to build confidence intervals for the distance to the efficient frontier. Section 4 applies the method to assess the distance between the S&P 500 stock index and the efficient frontier derived from the corresponding investment universe. Section 5 concludes.

2 Estimating the distance to the efficient frontier

2.1 Portfolio selection and production theory

Consider the problem of an investor selecting a portfolio made up of n risky assets. Under the assumption that no risk-free asset is available and short sales are forbidden, the investment opportunity set consists of all linear combinations of the n initial assets:

$$\mathcal{F} = \{w \in \mathbb{R}^n_+ | w' i_n = 1\}, \tag{2.1}$$

where i_n is a vector $(n \times 1)$ of ones. In the classical mean-variance setting, the objectives of the investor can be split into $x \in \mathbb{R}$ and $y \in \mathbb{R}$, which relate in the production theory to the input to be minimized (the variance of the portfolio return) and the output to the maximized (the mean of the portfolio return), respectively. We then define the set $\mathcal{X}_n = \{(x_i, y_i) \mid i = 1, \ldots, n\}$ representing the input/output of the original data set (the n risky assets). For a given portfolio, $w \in \mathcal{F}$, we compute its input and output, (x_w, y_w) from the previous set. In this setting, the objective of any investor is characterized by a vector of weights.

The input/output representation of the investment opportunity set is given by:

$$\mathcal{N} = \{(x_w, y_w) \in \mathbb{R}^2 | w \in \mathcal{F}\}. \tag{2.2}$$

It refers to all the portfolios generated by linear combinations in $\mathcal{F}$. Note that $\mathcal{X}_n \subset \mathcal{N}$, which corresponds to $\{(x_{e_i}, y_{e_i}) \mid i = 1, \ldots, n\}$, where e_i is the ith column of I_n, the identity matrix of order n.

Under the free-disposability assumption stating that it is always possible to achieve lower outputs with more inputs, we can define the Free Disposal Hull (FDH) of $\mathcal{N}$ by:

$$\Psi = \bigcup_{(x_w, y_w) \in \mathcal{N}} \left\{(x, y) \in \mathbb{R}^2 \mid x \geq x_w, y \leq y_w\right\}, \tag{2.3}$$

and the efficient frontier is then characterized as:

$$\Psi^\partial = \{(x, y) \in \Psi \mid \text{ for any } (\tilde{x}, \tilde{y}) \text{ such that } \tilde{x} < x, \tilde{y} > y, (\tilde{x}, \tilde{y}) \notin \Psi\}. \tag{2.4}$$

Any efficient portfolios belongs to the efficient frontier and is identified by a specific value of (x, y).

Next, let us consider portfolio $(x, y) \in \Psi$. We characterize the efficient frontier using a flexible approach based on directional distance functions introduced by Chambers et al. (1998). These functions generalize the traditional radial measures provided by both input and output distance functions. Given a direction vector $(-g_x, g_y)$ where $(g_x, g_y) \in \mathbb{R}^2_+$, the directional distance function projects portfolio (x, y) onto the efficient frontier in the chosen direction:

$$\delta(x, y; g_x, g_y) = \sup\{\beta \mid (x - \beta g_x, y + \beta g_y) \in \Psi\}. \tag{2.5}$$

By definition $\delta(x, y; g_x, g_y) \geq 0$ if and only if $(x, y) \in \Psi$ and the points belonging to Ψ^∂ are characterized by $\delta(x, y; g_x, g_y) = 0$. Since the measure is additive, it allows to handle jointly any positive or negative values of the input and output. The measure is also translation-invariant and independent from the measurement units provided that the units of the direction vector and of the input and output are the same.

Nalpas et al. (2017) propose a simple and fast methodology to evaluate the frontier of Ψ, as described in (2.4). In a nutshell, their algorithm involves Monte-Carlo simulations for covering the set of possible portfolios in (2.3) in an iterative and efficient way using only enumerative techniques (in computing at each step FDH estimators). Remarkably, this method can be extended to more complex settings, including additional inputs/outputs, like skewness and kurtosis, without increasing the level of computational complexity. This possibility confers on users a significant advantage because the theoretical optimization problem hidden in searching the characterization of Ψ^∂ is typically much harder, as in Briec et al. (2007), Jurczenko et al. (2012), and the discussion in Nalpas et al. (2017).

2.2 Statistical framework

The FDH estimator of Ψ was introduced in production efficiency analysis by Deprins et al. (1984) allowing the estimation of non-convex attainable sets. Suppose we consider a set of N random weights w_j, each of size n, $j = 1, \ldots, N$, so that we can build a sample of size N of inputs and outputs (x_{w_j}, y_{w_j}). The asymptotic properties of the FDH estimator derived in Korostelëv et al. (1995) and Park et al. (2000), indicate that, under mild regularity conditions, the rate of convergence of the resulting estimator is $N^{1/(p+q)}$ where p is the number of inputs and q the number of outputs.

In our case, since $p = q = 1$, we achieve the rate $N^{1/2}$ and the corresponding estimator of the distance between a given asset (x, y) and the efficient frontier converges at this rate to a limiting Weibull distribution. So, in theory, if N is large enough, we can estimate, at any desired level of statistical accuracy, the distance from a given portfolio to the efficient frontier. In practice however, memory limitations make the direct generation of the N random weights numerically intractable for huge values of N (say, above 10^6).

Drawing on this, the algorithm of Nalpas et al. (2017), hereafter the NSV algorithm, avoids direct evaluation of a huge number of weights. It generates iteratively and selectively a limited number of random weights, which are numerically easy to handle at each step using FDM estimators. The iteration process is shown to converge to the true value. Its simple stopping rule is comparable to any stopping rule or tolerance level used in nonlinear optimization routines (for details, see Nalpas et al. 2017). Ultimately, the algorithm produces several thousands of weights, and the error associated with the FDH estimator is negligible.

In the general setup of production theory, our methodology addresses the sensitivity of estimated efficiency scores of a given production unit (x, y) to input and output variations from the original sample $\mathcal{X}_n$. Precisely, we use the NSV algorithm to evaluate the distance function for a (x, y) in (2.5) with a negligible error, under the assumption that the (x_i, y_i)'s in the original dataset $\mathcal{X}_n$ are measured without statistical error. Transposed to the financial context, our objective is to produce a sensitivity analysis in the performance of a portfolio (x, y), measured by $\delta(x, y; g_x, g_y)$, given that the original data (x_i, y_i) in $\mathcal{X}_n$ are estimated from time series of asset returns. The problem is however complicated by the fact that the original dataset, made up of financial returns, includes time series with time dependence.

Consider a sample of n historical returns, R_{it}, $i = 1, ..., n$, observed over a given period of time $t = 1, ..., T$. From this sample, we derive the empirical counterparts $\widehat{\mathbb{E}}_T$ and $\widehat{\mathbb{V}}_T$ of the unknown $(n \times 1)$ vector of means $\mathbb{E}$ and the unknown $(n \times n)$ variance-covariance matrix $\mathbb{V}$. Element by element, for $i, j = 1, \ldots, n$, we have:

$$\widehat{\mathbb{E}}_{i,T} = \frac{1}{T} \sum_{t=1}^{T} R_{it},$$
$$\widehat{\mathbb{V}}_{ij,T} = \frac{1}{T} \sum_{t=1}^{T} (R_{it} - \widehat{\mathbb{E}}_{iT})(R_{jt} - \widehat{\mathbb{E}}_{jT}). \tag{2.6}$$

For any portfolio $w \in \mathcal{F}$, we have the corresponding moments:

$$\widehat{\mathbb{E}}_T(w) := \widehat{y}_{wT} = w'\widehat{\mathbb{E}}_T,$$
$$\widehat{\mathbb{V}}_T(w) := \widehat{x}_{wT} = w'\widehat{\mathbb{V}}_T w. \tag{2.7}$$

These equations provide the input/output representation of the opportunity set $\mathcal{N}$ defined in (2.2) given the information on the mean $\widehat{\mathbb{E}}_T$ and the covariance matrix $\widehat{\mathbb{V}}_T$ of the n original assets.[1]

In practice, the original dataset $\mathcal{X}_n$ we use to apply the NSV algorithm is:

$$\mathcal{X}_{n,T} := \{(\widehat{x}_{iT}, \widehat{y}_{iT}) \mid i = 1, \ldots, n\}, \tag{2.8}$$

where $\widehat{x}_{iT} := \widehat{x}_{e_iT} = \widehat{\mathbb{V}}_{ii,T}$ and $\widehat{y}_{iT} := \widehat{y}_{e_iT} = \widehat{\mathbb{E}}_{i,T}$, depend on the realized time series R_{it}, $i = 1, ..., n$, and $t = 1, ..., T$. As a consequence, the evaluation of the efficient frontier and

[1] Nalpas et al. (2017) give similar relations for skewness and kurtosis.

the distance measures depend on the empirical estimates we use. For example, changing the observation period affects $\mathcal{X}_{n,T}$, and so the efficient frontier. As any procedure solving (2.5)), the NSV algorithm produces point estimates of $\delta(\cdot)$, denoted $\widehat{\delta}_T(\cdot)$.

In sum, our objective is to use the NSV algorithm to analyze, for any given portfolio (x, y), the sensitivity of directional distances to the efficient frontier $\widehat{\delta}_T(x, y; g_x, g_y)$ to random changes in $\mathcal{X}_{n,T}$. From there, we will derive confidence intervals for $\delta(x, y; g_x, g_y)$. In the Markowitz mean-variance setting, these confidence intervals will allow us to test directional efficiencies of selected portfolios of interest.

3 Sensitivity analysis through subsampling

A natural way to meet our objective is to simulate random variations of the initial dataset $\mathcal{X}_{n,T}$ in an appropriate way and analyze how the efficient frontier, and the derived distance measures, are sensitive to these random variations. Resampling and bootstrap techniques are often appropriate but the returns are time series data having potential time dependence. Two main techniques have been adapted to time series, the *moving block bootstrap* developed and analyzed, e.g., by Hall et al. (1995), Lahiri (2002), Lahiri (2003), Lahiri et al. (2007), and the *subsampling methodology* proposed by Politis et al. (1997), Politis et al. (1999) and Politis et al. (2001). We adopt the subsampling approach for its simplicity and because it is consistent under very mild conditions (Politis et al. 2001). In addition, this approach is relevant for time series with heteroskedasticity (see Politis et al. 1997, and Chapter 4 in Politis et al. 1999).

Consider the infinite sequence of random returns of n risky assets $\{R_t\}_{t=-\infty}^{+\infty}$ where $R_t \in \mathbb{R}^n$ and denote by P its joint probability distribution. We impose on $\{R_t\}$ a weak-dependence condition. Precisely, we assume that the sequence is α-mixing, which intuitively means that the future of the variables of interest tends to be independent from its past when the time lag between the two periods goes to infinity (see e.g. Politis et al. (1997) for a more formal presentation).

3.1 Variations in expected return and volatility

The subsampling proposed by Politis et al. (1997) can be summarized as follows. We are interested in a real-valued or vector-valued parameter $\theta(P) \in \mathbb{R}^k$ estimated from the sample $\{R_1, \ldots, R_T\}$. Let $\widehat{\theta}_T = \widehat{\theta}_T(R_1, \ldots, R_T)$ be the estimator of $\theta(P)$. Our goal is to build confidence regions for $\theta(P)$. In our setup, the initial ingredients for the NSV algorithm are the vector of means and the covariance matrix of the returns. We thus have: $\theta(P) =$

$(\mathbb{E}, \mathbb{V}) \in \mathbb{R}^n \times \mathbb{R}^{n(n+1)/2}$, where the variances and covariances are stacked in a vector of length $n(n+1)/2$. So, we have $\widehat{\theta}_T = (\widehat{\mathbb{E}}_T, \widehat{\mathbb{V}}_T)$ and $k = n(n+3)/2$.

From the full series $(R_1, \ldots, R_T)$, we consider all the subsamples of size b starting successively at date a, where $a = 1, \ldots, T-b+1$. The resulting series are denoted $R_a, \ldots, R_{a+b-1}$. Next, we define $\widehat{\theta}_{b,a} = \widehat{\theta}_b(R_a, \ldots, R_{a+b-1})$, the estimator of $\theta(P)$ based on these subsamples of size b. We denote $J_{b,a}(P)$ the sampling distribution of $\sqrt{b}(\widehat{\theta}_{b,a} - \theta(P))$ (we use the rate $\sqrt{b}$ since our estimators are empirical means and covariances). Hence, for any Borel set A defined on $\mathbb{R}^k$ we may define:

$$J_{b,a}(A, P) = \text{Prob}_P\big[\sqrt{b}(\widehat{\theta}_{b,a} - \theta(P)) \in A\big] \tag{3.1}$$

Our object of interest is the sampling distribution of the estimator over the full set of data, $\widehat{\theta}_T = \widehat{\theta}_{T,1}$. It can be characterized by the asymptotic distribution of $\sqrt{T}(\widehat{\theta}_T - \theta(P))$, i.e., for any Borel set A

$$J_{T,1}(A, P) = \text{Prob}_P\big[\sqrt{T}(\widehat{\theta}_{T,1} - \theta(P)) \in A\big]. \tag{3.2}$$

The idea is that if, under workable assumptions, $b/T \to 0$ and $b \to \infty$ as $T \to \infty$, then the empirical distribution of the $T - b + 1$ values of $\sqrt{b}(\widehat{\theta}_{b,a} - \theta(P))$ can serve as a good approximation of our object of interest $J_{T,1}(P)$. Moreover, to get a feasible estimator, we can replace the unknown $\theta(P)$ by $\widehat{\theta}_T$ because $\sqrt{b}(\widehat{\theta}_T - \theta(P))$ is of order $\sqrt{b/T}$ in probability and we have $b/T \to 0$.

Formally, Politis et al. (1997) define the following approximation of $J_{T,1}(P)$:

$$L_T(A) = \frac{1}{T-b+1} \sum_{a=1}^{T-b+1} \mathbb{1}\big[\sqrt{b}(\widehat{\theta}_{b,a} - \widehat{\theta}_T) \in A\big] \tag{3.3}$$

The regularity conditions state that (i) $J_{T,1}(P)$ converges weakly to some limiting law $J(P)$, as $T \to \infty$ and (ii) the distribution functions of the normalized estimator based on the subsamples are close on average to the normalized estimator of the full sample, i.e., $(T - b+1)^{-1} \sum_{a=1}^{T-b+1} J_{b,a}(A, P) \to J(A, P)$, if $b/T \to 0$ and $b \to \infty$ and $T \to \infty$. Theorem 2.1 in Politis et al. (2001) establishes the validity of the approximation and explains how to build asymptotic confidence regions for $\theta(P) \in \mathbb{R}^k$.

In practice, the confidence regions are obtained as follows. Consider the Euclidean norm $||\cdot||$ on $\mathbb{R}^k$ and build, as an analog to (3.3), the univariate empirical distribution of $||\sqrt{b}(\widehat{\theta}_{b,a} - \widehat{\theta}_T)||$:

$$F_{T,b}(w) = \frac{1}{T-b+1} \sum_{a=1}^{T-b+1} \mathbb{1}\big[||\sqrt{b}(\widehat{\theta}_{b,a} - \widehat{\theta}_T)|| \le x\big]. \tag{3.4}$$

This is an approximation of the sampling distribution of the univariate $||\sqrt{T}(\widehat{\theta}_T - \theta(P))||$. With the $(1-\alpha)$-quantile of the distribution $F_{T,b}$, denoted by:

$$w_{T,b,\alpha} = \inf\{w \mid F_{T,b}(w) \geq (1-\alpha)\}, \tag{3.5}$$

the region $\{\theta \in \mathbb{R}^k \mid ||\sqrt{T}(\theta - \widehat{\theta}_T)|| \leq w_{T,b,\alpha}\}$ has the asymptotic coverage level of $(1-\alpha)$.

Selecting an appropriate subsample size b

Theorem 2.1 in Politis et al. (2001) requires that as $T \to \infty$, $b \to \infty$ with $b/T \to 0$, but finding an optimal subsample size can be delicate in practice. For the moving block bootstrap, prevailing calibration methods involve numerical burdens (with double bootstrap), as discussed in Section 9.3.1 of Politis et al. (1999). A simpler approach to select a suitable subsample size is provided by the minimum volatility method presented in Politis et al. (2001). It is based on the fact that the asymptotic validity of the subsampling approximations holds for a broad range of choice of the subsample size b. In this range, the confidence regions considered as a functions of the subsample size b should be stable. The minimum volatility method exploits this stability by using a volatility measure described below.

The algorithm goes as follows (see Algorithm 6.1 in Politis et al. 2001). Define $b_{\min}$ and $b_{\max}$ as the range where we will seek a value for b. We may choose these numbers as small (close to 0) and large (close to 1) powers of T.

[1] For each value of b in a grid a values in $[b_{\min}, b_{\max}]$, compute the subsampling confidence region for $\theta(P)$ at a given level $(1-\alpha)$, having the endpoints $I_{b,\text{low}}$ and $I_{b,\text{up}}$;

[2] For a small integer m, let the volatility index VI_b be measured by the standard deviation of the endpoints $\{I_{b-m,\text{low}}, \ldots, I_{b+m,\text{low}}\}$ plus the standard deviation of the endpoints $\{I_{b-m,\text{up}}, \ldots, I_{b+m,\text{up}}\}$;

[3] The selected value of b is the one having the smallest volatility index VI_b.

Politis et al. (1999) mention that other measures can be used for the volatility index, like the volume of the confidence region. In any case, the idea is to obtain a stable behavior around the selected value of b.

3.2 Confidence interval for the distance

We use the subsampling algorithm to analyze for any given (x, y) the sensitivity of $\widehat{\delta}_T(x, y; g_x, g_y)$ to random changes in $\mathcal{X}_{n,T}$, which in turn depend on variations of $\widehat{\mathbb{E}}_T$ and $\widehat{\mathbb{V}}_T$. To avoid the numerical burden of running the NSV algorithm for many values of b, we select a value of b

that stabilizes the variations in the components of $\mathcal{X}_{n,T}$. To do so, we apply the methodology described in the previous section to $\theta_1(P) = (\mathbb{E}, \text{diag}(\mathbb{V}))$, representing the $2n$-vector of the n means and of the n variances. Together with the selected value of b, we obtain the $T-b+1$ subsamples of size b of the n returns: $R_a, \ldots, R_{a+b-1}$ for $a = 1, \ldots, T-b+1$, providing the estimators $\widehat{\mathbb{E}}_{a,b}$ and $\widehat{\mathbb{V}}_{a,b}$.

In the next step, thanks to the NSV algorithm, $\widehat{\mathbb{E}}_{a,b}$ and $\widehat{\mathbb{V}}_{a,b}$ lead to the estimator $\widehat{\delta}_{a,b}(x, y; g_x, g_y)$ of $\delta(x, y; g_x, g_y)$. As in (3.4), the sampling distribution of the normalized error $|\sqrt{T}(\widehat{\delta}_T(x, y; g_x, g_y) - \delta(x, y; g_x, g_y))|$, is then approximated by

$$F_{T,b}^{\delta_{x,y}}(w) = \frac{1}{T-b+1} \sum_{a=1}^{T-b+1} \mathbb{1}\Big\{ |\sqrt{b}\big[\widehat{\delta}_{a,b}(x, y; g_x, g_y) - \widehat{\delta}_T(x, y; g_x, g_y)\big]| \leq x \Big\}, \tag{3.6}$$

where we use simple absolute value because δ is a scalar. Therefore, the confidence interval for $\delta(x, y; g_x, g_y)$ with asymptotic coverage level $(1-\alpha)$ is given by:

$$I_{\delta_{x,y}}(1-\alpha) = \{\delta \text{ such that } |\sqrt{T}(\delta - \widehat{\delta}_T(x, y; g_x, g_y))| \leq w_{T,b,\alpha}^{\delta_{x,y}}\}, \tag{3.7}$$

where $w_{T,b,\alpha}^{\delta_{x,y}}$ is the $(1-\alpha)$ quantile of $F_{T,b}^{\delta_{x,y}}(\cdot)$.

Finally, this results in a symmetric two-sided interval:

$$\delta(x, y; g_x, g_y) \in \left[\widehat{\delta}_T(x, y; g_x, g_y) - \frac{w_{T,b,\alpha}^{\delta_{x,y}}}{\sqrt{T}}, \widehat{\delta}_T(x, y; g_x, g_y) + \frac{w_{T,b,\alpha}^{\delta_{x,y}}}{\sqrt{T}} \right]. \tag{3.8}$$

Depending on the context, we may also derive one-sided confidence intervals and two-sided equal tailed intervals.

4 Application: Is the S&P 500 index efficient?

This section uses our new method to revisit a long-standing controversy located at the heart of portfolio theory and having significant consequences for portfolio management. The debate about the efficiency of index investing – i.e., holding portfolios mimicking large and diversified financial indexes – is almost as old as the Markowitz approach to portfolio selection. Under the assumptions of the CAPM developed in the 1960s (see Perold (2004) for a historical perspective), the market portfolio, defined as the portfolio composed of all the capitalization-weighted assets in the market under consideration, is proven mathematically to be efficient. Shortly after this theoretical result was established, finance practitioners and scholars started raising critical issues about its practical implications for optimal diversification and passive investment strategies. The controversy culminated with the Roll (1977) critique stating that market indices, arguably the best proxies of true market portfolios, fail to reach optimal diversification and can therefore be located away from the efficient frontier. As a consequence, the CAPM appeared to some as an essentially untestable theoretical construct, distant from the everyday functioning of actual financial markets.

As a matter of facts, both the Markowitz portfolio selection theory and the CAPM survived the controversies. Textbooks still introduce these contributions as core building blocks of market finance. Even more so, passive index investing has gained traction, notably with the development of the Exchange Traded Funds (ETFs) favoring easy access to index investment. Besides diversification benefits, index investing is associated with reduced management fees and transaction costs (Gârleanu and Pedersen 2022). In sum, the question of whether (the proxy of) the market portfolio is efficient is still topical.

This section contributes to the conversation by examining whether 1) the distance from the S&P 500 index to its efficient frontier depends on the direction used to reach the frontier, and 2) the corresponding directional distances are significantly different from zero. It shows that the directional distances slightly vary according to the direction chosen, but all are significant, suggesting that Roll's critique is robust.

4.1 Tuning the parameters of the FDH estimation

In this part, we focus on the practical aspects of estimating the efficiency frontier: 1) defining the relevant direction vectors (g_x, g_y); 2) estimating the directional distance between from given portfolio to the frontier with the FDH algorithm.

Direction vectors. Let us consider a portfolio characterised by its mean return y_M and its variance x_M, which are its coordinates in the mean-variance plan. Segments linking this point to a point of the efficient frontier have a large spectrum of possible directions (g_{x_M}, g_{y_M}). However, some directions are unachievable because they do not correspond to any point of the efficient frontier.

In practice, the decision of choosing the direction belongs to the investor holding the portfolio of interest. Therefore, our approach applies to any achievable direction. The only constraint we impose is is that g_{x_M} and g_{y_M} have to be positive. In the mean-volatility plan, we define direction vectors of norm 1, which have the same unit as input-output vectors. Our efficiency measure is the Euclidean distance δ from the portfolio to the frontier along the chosen direction.

In this section, we consider several direction vectors, all of norm 1. In polar coordinates, each vector can be characterized by its angle with the horizontal axis. In our application, we rely on visual intuition by mentioning the angle rather than the vector.

FDH approximation of the efficient frontier. We use the NSV algorithm proposed by Nalpas et al. (2017) to generate an approximation of the efficient frontier. At each step, we iterate portfolios made up of a linear combination of initial assets and estimate their FDH efficiency performances. The last iteration produces reference portfolios for building the efficient frontier. The FDH approximation of the efficient frontier is thus computed as a step function where each step corresponds to a reference portfolio.

To estimate the distance from a given portfolio to the frontier in a given direction, we can compute both the distance to the approximated FDH frontier in the chosen direction (with the exactly corresponding angle) and the distance to the closest achievable efficient reference portfolio (with a slightly different angle). Both distances (and angles) are very close to each other for a sufficiently large number of iterations. In the application below, we use 100,000 iterations ($k_{\max}$ in Nalpas et al. 2017) and 100 portfolios generated at each iteration (N_c in Nalpas et al. 2017).

4.2 Results

We collected the daily returns of the S&P 500 stock index and of all the stocks covering the index, from January 2015 to May 2022. The FDH approximation of the efficient frontier represented in Figure 1 is obtained after 100,000 iterations of the NSV algorithm. In the mean-volatility plan, the figure shows the frontier, the S&P 500 index, and the directions of interest. The achievable directions from the index to the efficient frontier correspond to efficient portfolios with volatility taking values between the minimal volatility on the frontier (utmost left point) and the same volatility as the index (on the vertical segment). Only this portion of the frontier is composed of portfolios that can dominate the S&P 500 portfolio in the mean-variance setting. Angles of more than 90 degrees (measured from the horizontal axis on the left) lead to efficient portfolios riskier than the index, which is incompatible with these portfolios dominating the index in the mean-variance setting.

The efficient portfolios in Figure 1 are presented with angles varying from 0 to 90 degrees. Actually, the efficient portfolios obtained with the FDH algorithm correspond to angles from around 20 degrees up to 90 degrees. Below 20 degrees, the FDH method does not identify efficient portfolios anymore and conventionally terminates the frontier with a vertical line. Therefore, we restrict our directions of interest between 20 degrees and 90 degrees, with steps of 10 degrees. Table 1 provides descriptive statistics for the returns of the 8 efficient portfolios located at the intersections of the efficient frontier and the directional segments from the S&P 500 index. As Figure 1 suggests, Table 1 shows that the S&P 500 index is dominated by the every reference efficient portfolios (i.e., smaller annualized return and larger or equal volatility). The main question is thus whether the observed dominance is associated with directional distances significantly different from zero.

Note that the graph in Figure 1 might lead to believe that the distance to the frontier increases systematically with portfolio volatility. This is not necessarily the case in general. With the same frontier, the distance would not be monotonic for some portfolios with an expected return (and/or a volatility) larger than that of the S&P 500 index.

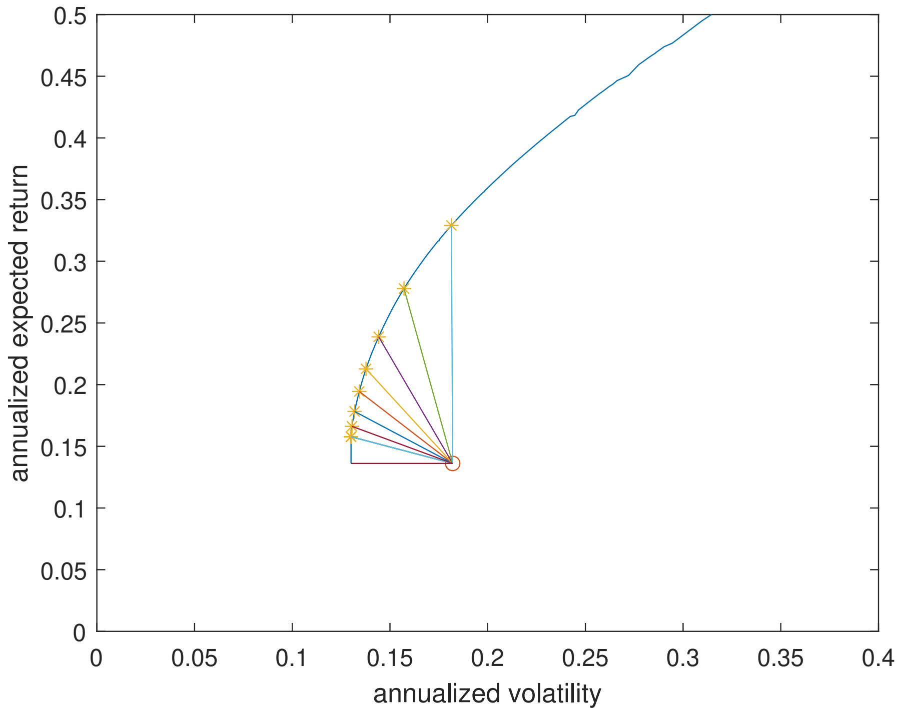

Figure 1: Efficient frontier estimation.
The horizontal axis gives the volatility and the vertical axis the annualized return. The symbol o represents the S&P500 market index, and * the reference efficient portfolios for various angles.

	S&P 500	Ptf 1	Ptf 2	Ptf 3	Ptf 4	Ptf 5	Ptf 6	Ptf 7	Ptf 8
Chosen angle (degree)		90	80	70	60	50	40	30	20
Annualized return	0.136	0.329	0.278	0.239	0.213	0.194	0.178	0.166	0.158
Volatility	0.181	0.181	0.157	0.144	0.138	0.134	0.132	0.131	0.131

Table 1: **Mean-variance characteristics of the S&P 500 index and reference portfolios on the efficient frontier**

	Chosen angle	Approx. angle	Distance (chosen angle)	Distance (approx. angle)	Confidence interval ($\alpha = 5\%$)	
Pft1	90	89.812	0.193	0.193	0.1382	0.2478
Pft2	80	80.043	0.143	0.144	0.1174	0.1686
Pft3	70	69.755	0.109	0.109	0.0968	0.1212
Pft4	60	59.948	0.088	0.088	0.0826	0.0934
Pft5	50	50.681	0.074	0.075	0.0717	0.0763
Pft6	40	40.015	0.065	0.065	0.0649	0.0651
Pft7	30	30.265	0.059	0.060	0.0588	0.0592
Pft8	20	22.595	0.055	0.056	0.0549	0.0551

Table 2: **Directional distances between the S&P 500 index and the efficient frontier, with confidence intervals (at 5%)**
Columns 1 and 3 give the angle chosen (in degrees) and the corresponding distance of the S&P 500 index to the frontier, respectively. Columns 2 and 4 show their respective proxies obtained after 100,000 iterations of the NSV algorithm. The last two columns delineate the confidence intervals.

In Table 2, we run the tests for directional efficiency of the market portfolio. Comparing columns 1 and 2 in the table shows that the proxies for the angles delivered by our algorithm after 100,000 iterations are, except for Portfolio 8, in a range of less than one degree away from the original angle. Likewise, the directional distance to the frontier with the approximate angle is very close to the corresponding distance with the chosen angle. The proxies we compute allow building confidence intervals. The last two columns of Table 2 show that none of the confidence intervals obtained at the 5% level includes zero, leading us to reject the null that the proxy for the market portfolio belongs to the efficient frontier along the chosen direction. Our testing method is the first to bring such a uniform rejection in all achievable directions, thereby substantiating the doubts expressed in the profession about the optimality of index investing. Using our directional approach therefore reinforces the validity of Roll's critique stating that even large and well-diversified financial indices, meant to proxy the theoretical market portfolio, mostly fail being efficient portfolios.

5 Conclusions

To our knowledge, this paper is the first to address the impact of random variations in inputs and outputs on efficiency scores. The originality of our approach stems from combining the NSV algorithm and the subsampling method to build confidence intervals around the scores of interest. We target key applications in portfolio management since our approach offers an original way to assess the distance between a given portfolio and the efficient frontier, while controlling for the direction used to reach the frontier.

The statistical tool to assess any directional distance between a given portfolio and the efficient frontier overcome the restrictions of existing tests for mean-variance efficiency to the horizontal and vertical directions, so opening new perspectives for asset management. Identifying and parametrizing the direction of the straight line between the portfolio of interest and the efficient frontier may provide additional flexibility to portfolio managers willing to increase the efficiency of their portfolio while maintaining a given risk-return trade-off. It can also help them address investors' changes in risk/return sensitivity in portfolio updates.

The empirical exercise in this paper contributes to a long-standing debate taking place in finance about the optimality of index investing. It is however fair to say that our approach is based on in-sample optimal portfolios (out of sample, the gains from optimal diversification can be offset by estimation errors, as shown for example by DeMiguel et al. 2009) and excludes any transaction costs, which are notoriously difficult to assess. Our findings emphasize the robustness of previous evidence suggesting that the empirical counterpart of the market portfolio - here, the S&P 500 stock index - is not efficient. This, still controversial, result is typically used to question the empirical relevance of the CAPM's original formulation.

A promising avenue for future research stems from determining the smallest number of changes in a portfolio composition needed to transform an inefficient portfolio into an efficient one along a given direction (at a fixed confidence level). Our innovative method can also help design portfolios constrained by social and environmental criteria with minimal loss of financial efficiency along the risk/return direction chosen by the investor.

References

Acharya, V. V., Gale, D., and Yorulmazer, T. (2011). Rollover risk and market freezes. *The Journal of Finance*, 66(4):1177–1209.

Basak, G., Jagannathan, R., and Sun, G. (2002). A direct test for the mean variance efficiency of a portfolio. *Journal of Economic Dynamics and Control*, 26(7-8):1195–1215.

Briec, W., Kerstens, K., and Jokung, O. (2007). Mean-variance-skewness portfolio performance gauging: a general shortage function and dual approach. *Management Science*, 53(1):135–149.

Brière, M., Drut, B., Mignon, V., Oosterlinck, K., and Szafarz, A. (2013). Is the market portfolio efficient? a new test of mean-variance efficiency when all assets are risky. *Finance*, 34(1):7–41.

Chambers, R. G., Chung, Y., and Färe, R. (1998). Profit, directional distance functions, and nerlovian efficiency. *Journal of Optimization Theory and Applications*, 98:351–364.

DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The review of Financial studies*, 22(5):1915–1953.

Deprins, D., Simar, L., and Tulkens, H. (1984). Measuring labour-efficiency in post offices. *The Performance of Public Enterprises: Concepts and Measurements*.

Gârleanu, N. and Pedersen, L. H. (2022). Active and passive investing: Understanding samuelson's dictum. *The Review of Asset Pricing Studies*, 12(2):389–446.

Hall, P., Horowitz, J. L., and Jing, B.-Y. (1995). On blocking rules for the bootstrap with dependent data. *Biometrika*, 82(3):561–574.

Jurczenko, E., Maillet, B., and Merlin, P. (2012). Hedge fund portfolio selection with higher-order moments: A nonparametric mean–variance–skewness–kurtosis efficient frontier. *Multi-moment Asset Allocation and Pricing Models*, pages 51–66.

Kandel, S. (1984). The likelihood ratio test statistic of mean-variance efficiency without a riskless asset. *Journal of Financial Economics*, 13(4):575–592.

Kandel, S. and Stambaugh, R. F. (1995). Portfolio inefficiency and the cross-section of expected returns. *The Journal of Finance*, 50(1):157–184.

Korostelëv, A. P., Simar, L., and Tsybakov, A. B. (1995). Efficient estimation of monotone boundaries. *The Annals of Statistics*, pages 476–489.

Lahiri, S. (2002). On the jackknife-after-bootstrap method for dependent data and its consistency properties. *Econometric Theory*, 18(1):79–98.

Lahiri, S. (2003). Selecting optimal block lengths for block bootstrap methods. *Department of Statistics Iowa State University. Retrieved from https://www. interfacesymposia. org I*, 3.

Lahiri, S., Furukawa, K., and Lee, Y.-D. (2007). A nonparametric plug-in rule for selecting optimal block lengths for block bootstrap methods. *Statistical Methodology*, 4(3):292–321.

Levy, M. and Roll, R. (2010). The market portfolio may be mean/variance efficient after all: The market portfolio. *The Review of Financial Studies*, 23(6):2464–2491.

Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy*, 60(2):151–158.

Nalpas, N., Simar, L., and Vanhems, A. (2017). Portfolio selection in a multi-moment setting: A simple monte-carlo-fdh algorithm. *European Journal of Operational Research*, 263(1):308–320.

Park, B. U., Simar, L., and Weiner, C. (2000). The fdh estimator for productivity efficiency scores: asymptotic properties. *Econometric Theory*, 16(6):855–877.

Perold, A. F. (2004). The capital asset pricing model. *Journal of Economic Perspectives*, 18(3):3–24.

Politis, D., Romano, J., and Wolf, M. (1999). Subsampling. *Springer-Verlag, New-York*.

Politis, D. N., Romano, J. P., and Wolf, M. (1997). Subsampling for heteroskedastic time series. *Journal of Econometrics*, 81(2):281–317.

Politis, D. N., Romano, J. P., and Wolf, M. (2001). On the asymptotic theory of subsampling. *Statistica Sinica*, pages 1105–1124.

Roll, R. (1977). A critique of the asset pricing theory's tests part i: On past and potential testability of the theory. *Journal of Financial Economics*, 4(2):129–176.