European Journal of Operational Research 265 (2018) 571–579

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Stochastics and Statistics

Comparing large-sample maximum Sharpe ratios and incremental variable testing

Michael Hanke[a,*], Spiridon Penev[b]

[a] Institute for Financial Services, University of Liechtenstein, Fürst-Franz-Josef-Strasse, 9490 Vaduz, Liechtenstein

[b] School of Mathematics and Statistics, The University of New South Wales, Sydney NSW 2052, Australia

ARTICLE INFO

Article history:
Received 2 October 2016
Accepted 9 August 2017
Available online 12 August 2017

Keywords:
Multivariate statistics
Finance
Maximum Sharpe ratio
Incremental variables
Mean-variance spanning

ABSTRACT

Most existing results on the distribution of the maximum Sharpe ratio depend on the assumption of multivariate normal return distributions. We use recent results from the literature to provide an analytical representation of the distribution of the difference between two maximum Sharpe ratios for much less restrictive distributional assumptions, both with and without short sales. Knowing the distribution of the difference enables us to test ex ante whether or not the inclusion of additional variables leads to a significant improvement in the maximum Sharpe ratio. In addition, we characterize the optimal long-only solution and provide conditions for global optimality.

1. Introduction

Particularly in times of low interest rates and decreasing risk premia, investors look for potential improvements to their portfolios by including new asset classes. Also within asset classes, the performance might be improved by considering variables such as book-to-market or momentum as potential return drivers. For some of these variables, their impact on average returns is generally accepted, since there exists both a sound theoretical foundation and solid empirical evidence. For others, it is less clear whether their inclusion significantly enhances the portfolio opportunity set. Fama and French (2015) discuss the related literature and criticize the standard line of argument used there, which employs Fama and MacBeth (1973) cross-sectional regressions to sort stocks into portfolios and illustrates the (seeming) importance of a new candidate variable on the basis of univariate return spreads. Fama and French (2015) demonstrate that judging a variable's importance by the univariate spread it produces in average returns overestimates its positive impact compared to the true impact observed when adding this variable in a portfolio context.

When starting from an existing portfolio, only the incremental improvement from new additions is relevant. If a candidate component improves the portfolio opportunity set, it should be included, and disregarded otherwise. In classical portfolio theory,[1] if a riskless asset is available, this may be recast into an equivalent statement formulated in terms of the Sharpe ratio: A candidate component should be included if and only if it significantly increases the maximum attainable Sharpe ratio. Gibbons, Ross, and Shanken (1989) note that their approach to testing portfolio efficiency, which is formulated in a regression framework, can also be interpreted as a test for the significance in the Sharpe ratio improvement caused by the improvement in the opportunity set.

The Sharpe ratio has some well-known shortcomings: It is based on the standard deviation, which is not a coherent risk measure in the sense of Artzner, Delbaen, Eber, and Heath (1999). Moreover, real-world return distributions are not necessarily multivariate normal, but frequently exhibit skewness and excess kurtosis. To account for these deviations from multivariate normality, extensions to the Sharpe ratio have been proposed (see, e.g., Biglova, Ortobelli, Rachev, & Stoyanov, 2004). Despite these shortcomings and the existence of alternatives that are theoretically more appealing, the Sharpe ratio is still the most widely used reward-to-risk ratio in practice, and it may even be the main relevant criterion for many practitioners. A portfolio manager considering a new variable may hence be interested in testing whether or not its inclusion would lead to a significant increase in the maximum attainable Sharpe ratio. Fama and French (2015) discuss the use of Gibbons et al. (1989) (GRS) tests for this purpose and illustrate a major disadvantage of these tests from a practical perspective: Since the tests assume that unlimited short sales are allowed,

* Corresponding author.
E-mail address: Michael.Hanke@uni.li (M. Hanke).

[1] For an overview of alternative theoretical foundations of mean-variance analysis, see Liu (2004). Regarding its robustness to alternative model assumptions, Simaan (2014) shows that mean-variance based portfolio optimization also leads to good results in a more general expected utility maximizing framework.

http://dx.doi.org/10.1016/j.ejor.2017.08.018

they may indicate a significant increase in the maximum attainable Sharpe ratio in cases where the portfolios leading to this increase require unrealistic leverage. Fama and French (2015) illustrate numerically that despite highly significant GRS statistics, Sharpe ratio improvements for a long-only investor from adding a third variable to an initial set of two are typically very small. A further disadvantage of the GRS statistic is its assumption of multivariate normal asset returns.

The reason why Fama and French (2015) resort to numerical illustrations for the case of the long-only investor is the lack of a suitable test for the difference between the maximum Sharpe ratios that are attainable with and without an incremental variable when short sales are not allowed. In this paper, we derive such a test for an ex ante comparison of optimal portfolios from different investment opportunity sets, which generalizes their GRS-based test not only for the long-only case, but also for deviations from multivariate normality. Furthermore, whereas the GRS-based test is formulated such that one of the investment opportunity sets is a proper subset of the other, our test allows for comparing any two investment opportunity sets. To derive our test, we will use recent results by Maller, Roberts, and Tourky (2016) on the distribution of the maximum Sharpe ratio under quite general distributional assumptions. The resulting formulae form the basis for comparing maximum Sharpe ratios for two investment opportunity sets. This includes, but is not limited to, testing whether increasing a given investment opportunity set significantly increases the maximum attainable Sharpe ratio.

One potential technical issue that arises with maximization of the Sharpe ratio in the long-only case is in the nature of the optimization problem. The long-only constraints make the problem quasi-concave, which in principle may give rise to multiple (local) optima. However, as noted in Stoyanov, Rachev, and Fabozzi (2007), since the function $g(x)$ (see e.g., Eq. (1) in Section 3) is a ratio of a linear and a *differentiable* convex function, the ratio is pseudo-concave in its argument. It is known (Stoyanov et al., 2007) that local extrema of pseudo-concave function are also global. For quasi-concave functions defined in a non-empty convex set, the optimal solution is an extreme point of the convex set. As an additional contribution of this paper, we provide conditions that allow us to check whether any such extreme point is indeed the global optimum we are looking for.

The difference in maximum attainable Sharpe ratios is closely related to mean-variance spanning tests, where the object of interest is usually the entire efficient frontier constructed from the available risky assets. In one sense, the problem studied in this paper appears to be simpler, since it involves only the tangency portfolio (as opposed to the entire efficient frontier). However, most existing mean-variance spanning tests suffer from the assumption of unlimited short sales, and many of these tests assume (conditional) multivariate normality of returns. The relation of this paper to the literature on mean-variance spanning tests will be discussed in more detail in Section 2.

The remainder of the paper is organized as follows: Section 3 reviews recent results on the distribution of the maximum Sharpe ratio when asset returns are not assumed to be multivariate normal. Section 4 provides a general result for the distribution of the difference between two maximum Sharpe ratios, both with and without short sales. Section 5 briefly discusses potential applications. Section 6 illustrates our results using a numerical example, and Section 7 concludes.

2. Tests on the difference between maximum Sharpe ratios and their relation to mean-variance spanning

One of the questions in this paper is whether the maximum Sharpe ratios attainable by forming portfolios of assets from two investment opportunity sets differ significantly. This question arises in classical portfolio theory if a riskless asset is available. In this case, the Sharpe ratio of the tangency portfolio determines the risk-return tradeoff that is attainable from investing in risky assets. If a riskless asset is not available, this efficient frontier is completely described by two portfolios on the frontier (two-fund separation). Frequently, these two portfolios are chosen as the tangency portfolio and the global minimum-variance portfolio. In the absence of a riskless asset, the relevant question is whether the efficient frontier of one set of assets (the benchmark or reference set) is identical to the efficient frontier of another set of assets (the test set). The finite sample distribution and moments of the sample minimum variance frontier have been analyzed by Kan and Smith (2008) for the case of multivariate normal asset returns. Corresponding tests for the difference between two such frontiers are known as mean-variance spanning tests in the literature (for an overview, see DeRoon & Nijman, 2001; Kan & Zhou, 2012). Since these tests analyze the difference between two efficient frontiers, they are not usually based on the distribution of the Sharpe ratio of the tangency portfolio, which would be an insufficient statistic for their purpose.

In the framework of regression-based spanning tests introduced by Huberman and Kandel (1987), Kan and Zhou (2012) discuss and compare various such tests. They point out that tests of spanning can be viewed as joint tests on the differences between the minimum variance portfolios and the tangency portfolios constructed from the reference set and the test set. The components are weighted implicitly according to the statistical accuracy of the estimates. Since the estimate for the minimum variance portfolio does not depend on the estimated means, it will be more accurate than the estimate for the tangency portfolio. For this reason, mean-variance spanning tests place more weight on the difference between the minimum variance portfolios than on the difference between the tangency portfolios, although the difference between the tangency portfolios may be economically much more important. Whereas standard tests may be able to distinguish clearly between largely similar efficient frontiers with a small difference in minimum variance portfolios, they find it much more difficult to classify a large difference in the tangency portfolios as statistically significant.

Because of the problems of joint testing and the difficulty in distinguishing between Sharpe ratios of tangency portfolios due to the estimation uncertainty in means, Kan and Zhou (2012, Section 4) suggest a step-down procedure for spanning tests, i.e., a sequential procedure that tests first for any difference in the tangency portfolios and then, in a second step, for the difference in minimum variance portfolios (conditional on no difference in tangency portfolios). They also discuss using different significance levels for the two parts, and an adjustment of these to the economic significance of the two components. Instead of using the traditional significance level of 5% for both components, this may mean a lower value of, e.g., 1% for the minimum-variance part, because the difference in minimum-variance portfolios is easier to detect and less important economically. In contrast, for the tangency portfolio part, a higher value of, e.g., 10% might be used to account for both higher economic importance and higher estimation uncertainty in means.

3. The maximum Sharpe ratio and its large-sample distribution

In this section, we introduce the notation and review existing results on the distribution of the maximum Sharpe ratio. We assume the existence of a risk-free asset. Let $\mu \in \mathbb{R}^p$ represent a vector of expected excess returns for p risky assets ($p \geq 1$). The efficient frontier in the presence of the risk-free asset is given by the tangent from the origin to the “risky-assets-only” efficient frontier,