European Journal of Operational Research 234 (2014) 382–391

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

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

The opportunity cost of mean–variance choice under estimation risk

Yusif Simaan *

Graduate School of Business, Fordham University, 113 West 60th Street, New York, NY 10023, United States

ARTICLE INFO

Article history:
Available online 24 January 2013

Keywords:
Mean–variance choice
Estimation risk
Expected utility maximization
Opportunity cost of sub-optimal portfolios

ABSTRACT

Mean–variance portfolio choice is often criticized as sub-optimal in the more general expected utility framework. It is argued that the expected utility framework takes into consideration higher moments ignored by mean variance analysis. A body of research suggests that mean–variance choice, though arguably sub-optimal, provides very close-to-expected utility maximizing portfolios and their expected utilities, basing its evaluation on in-sample analysis where mean–variance choice is sub-optimal by definition. In order to clarify this existing research, this study provides a framework that allows comparing in-sample and out-of-sample performance of the mean variance portfolios against expected utility maximizing portfolios. Our in-sample results confirm the results of earlier studies. On the other hand, our out-of-sample results show that the expected utility model performs worse. The out-of-sample inferiority of the expected utility model is more pronounced for preferences and constraints under which in-sample mean variance approximations are weakest. We argue that, in addition to its elegance and simplicity, the mean–variance model extracts more information from sample data because it uses the covariance matrix of returns. The expected utility model may reach its optimal solution without using information from the covariance matrix.

1. Introduction

Mean variance portfolio choice is often considered sub-optimal when evaluated in the subjective expected utility framework of Savage et al. (1954). Savage et al. (1954) indicates that a set of rationality axioms on the behavior of a decision maker is both necessary and sufficient for maximizing expected subjective expected utility.[1] Seminal work implies that investors rank investment prospects based on their mean–variance pairs. The two frameworks are consistent when either the utility function is restricted or when the joint distribution of asset returns is restricted. Hanoch and Levy (1970) show that Quadratic utility is sufficient for this consistency Chamberlain (1983) shows that joint elliptical distribution with existing covariance matrix is sufficient for the consistency. Gaussian returns are elliptical. The presence of skewness is inconsistent with elliptical returns. Empirical stock returns exhibit skewness and thus highlights the inconsistency for the application of portfolio selection derived from either model. For people who accept the behavioral axioms of Savage et al. (1954), mean–variance portfolio choice is sub-optimal. A body of literature examined whether mean–variance portfolio choice approximates expected utility maximization. Markowitz et al. (1959), Levy and Markowitz (1979), Dexter et al. (1980), Pulley (1983), and Kroll et al. (1984), Simaan (1993a,b), Hlawitschka (1994) use different methodologies under different assumptions, preferences and asset returns, and show that a utility-tailored Mean–Variance (MV) efficient portfolio provides a good approximation to both the expected-utility-maximizing-portfolio and to its expected utility. Under both paradigms, uncertainty is introduced in the form of probability distribution that depends on parameters known to the decision maker. However, in reality, a joint distribution of asset returns has to be specified and its parameters have to be estimated before a portfolio choice is made. In implementation, the decision maker risks both a specification error as she specifies the probability model and an estimation error as she estimates the parameters of the specified model. Complex probability models provide flexibility and thus reduce specification error. However, complexity requires estimating a larger number of parameters and thus is likely to cause more estimation errors, as relevant sample sizes limit degrees of freedom to estimate larger number of parameters.

This study compares mean variance choice to its subjective expected utility *ideal* under estimation risk. Estimation risk forces us to distinguish between in-sample and out-of-sample quality of a portfolio choice. Under full information of the joint distribution parameters, the in-sample opportunity cost of the mean–variance choice is positive, because the expected utility model is taken as

* Tel.: +1 646 220 0652; fax: +1 212 214 0747.
E-mail address: simaan@fordham.edu

[1] The expected utility model dates back to Daniel Bernoulli in the 18th century and was formally developed by John von Neumann et al. (1944) using objective probabilities. Savage proved the representation theorem of choice behavior for subjective uncertainty. See Read (2012) for a thorough review of the expected utility model and the mean–variance model.

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.01.025

the ideal. To put aside specification error, we assume no parametric structure on the joint asset-returns-distribution. We replace expectations and other moments with their sample counterparts. This approach is taken by Kroll et al. (1984). This study examines portfolio choices based on sample estimates of expected utility and sample-means and sample covariance matrix estimates of asset returns. Out-of-sample evaluation of in-sample portfolio choices (based on sample estimates) reveals that the expected utility choice is likely to be worse than their mean–variance counterparts. This result is more pronounced in cases where in-sample expected utility choice strongly dominates its mean–variance alternative.

We follow Kroll et al. (1984) in (i) using empirical returns; (ii) considering similar collection of utility functions; (iii) solving for the average-utility-maximizing portfolio; (iv) generating a dense MV efficient frontier; and (v) finding a mean–variance alternative to the portfolio in (iii) on the mean–variance-efficient-frontier. However, we differ in three issues. First we use a different metric to compare the average-utility-maximizing portfolio with its mean–variance-efficient alternative. We use the opportunity cost metric of Simaan (1993a,b), to be defined later, in order to place a dollar value in comparing any two distinct portfolio choices. Second, we distinguish in-sample performance from out-of-sample performance of the two strategies: We take ten years of monthly returns on the thirty stocks of the Dow Jones Industrial index stocks. We draw one thousand samples each consisting of 60-months from of the 120 months. Each sample is used to make investment decisions and derive in-sample portfolio choices. The remaining 60-months of each sample are used only to evaluate the out-of-sample performance of the portfolio decisions made based on the chosen sample. For each utility function, we compute its in-sample optimal portfolio, its best MV efficient portfolio, and the opportunity cost of the latter relative to the former. The opportunity cost is calculated once based on in-sample performance and another based on an out-of-sample performance. We repeat this process 1000 times and examine the statistical evidence of this experiment. Third, we solve each portfolio problem with two sets of constraints: (i) The first set imposes no-short sales constraints as is the case in Kroll et al. (1984); (ii) The second set allows short-sales but limit long and short position to 20% or less, and imposes the portfolio constraint. Expanding the set of portfolio constraints in the latter case is likely to result in a higher in-sample expected utility and consequently higher in-sample opportunity cost for the mean–variance strategy.

Ederington et al. (1995) questions using the empirical returns distribution employed in Levy and Markowitz (1979) by arguing that such returns constitute a small sample from a distribution whose non-Gaussian properties may fail to appear in a relatively small sample size. He simulates 4 quarterly returns from 40 quarterly returns on 130 mutual funds over a period of ten years, 1970–1979. Each 4 returns are used to build a single annual return. This process is repeated 10,000 times to construct 10,000 annual returns. The study evaluates the correlation of the expected utility to three expected utility approximations: (i) On a two-terms mean–variance Taylor expansion; (ii) A four-term mean–variance Taylor expansion done by forcing skewness to be zero and forcing Kurtosis to be a function of variance, as is the case with Gaussian returns, and (iii) a four terms Taylor expansion that does not constrain the first four moments. The Ederington simulated returns pose a stronger challenge to the MV approximations in Levy and Markowitz (1979), especially for exponential utility functions with an absolute risk aversion of more than 3. For the latter utilities, the two-terms-Taylor-quadratic approximations deteriorate. The near-normal second one, which constitutes an alternative mean–variance approximation, provides similar results to the four-term expansion which is a four-moment approximation. This Study is similar to Ederington et al. (1995) in its use of a bootstrapping methodology. It differs in at least three aspects. Like Kroll et al. (1984) it solves for optimal portfolios, their MV alternatives, and compares expected utilities of optimal portfolios to their MV alternatives. Like Simaan (1993a,b) it compares utilities using the opportunity cost of MV choice portfolios. Unlike Ederington et al. (1995) and Simaan (1993a,b), this study makes a distinction between in-sample expected performance and actual out-of-sample performance. It is the only one in the literature that makes such a distinction.[2]

The in-sample performance of MV choice in this experiment confirms the results reported in earlier studies like Levy and Markowitz (1979), Kroll et al. (1984) and Simaan (1993a,b). When short sales are not allowed we accept the hypothesis that the opportunity cost for each utility function is zero for log utility function, power utility functions and exponential utility functions with non-extreme measures of absolute risk aversion. On the other hand, we reject the hypothesis that the opportunity cost of mean–variance choice is zero for exponential utility functions with extreme measures of absolute risk aversion. Moreover, the opportunity cost of the latter utility functions is positive, and both statistically significant and economically significant. Simaan (1993a,b) points out similar results for these utilities. We argue that exponential utility investors with extreme measures of absolute risk aversion exhibit pathological focus on wealth preservation: No increase in expected return compensates for even infinitesimal increase in risk.

The study finds the out-of-sample opportunity cost differs sharply from the in-sample opportunity cost. Its sign is negative for almost all utility functions, unlike the in-sample case. Nevertheless, we still accept the hypothesis that the opportunity costs of MV choices are zero for log, power, and exponential utility functions with non-extreme measures of risk aversion. On the other hand, the opportunity costs are consistently negative, statistically and economically significant, and similar in magnitude to the in-sample ones for the exponential utility functions with extreme measures of absolute risk aversion. Hence, the dominance of expected-utility-maximization-strategies for investors obsessed in wealth preservation is deceptive. When estimation risk is considered such investors are better served with mean–variance-strategies.

When we allow short sales, the in-sample opportunity cost of MV choice increases and becomes significantly positive and statistically significant for each utility function. However, out-of-sample, expected-utility-maximization-strategies are consistently negative and statistically significant for all utility functions. Thus for all utility functions out-of-sample performance of the expected utility strategies are worse than mean–variance-strategies. In Section 2 we set up the evaluation framework, describe the data, outline our experiment and define the opportunity cost concept. In Section 3 we discuss our empirical findings. In Section 4 we provide an explanation to our findings in Section 3.

2. The evaluation framework

We use the framework of Kroll et al. (1984) with one extension and a modification. The framework compares expected utility choice and mean–variance choice using joint empirical returns-

[2] Another distinction is the difference in investment horizon. Ederington uses annual returns and we use monthly returns. Arditti and Levy (1975) show that the skewness of return distributions depend on the investment horizon. Specifically, compounding creates skewness in distributions that do not exhibit skewness. Levy and Duchin (2004) find that distributions that best fit stock returns depend on the investment horizon. Given this evidence, our result may be sensitive to the investment horizon.