Journal of Empirical Finance 16 (2009) 816–829

Contents lists available at ScienceDirect

Journal of Empirical Finance

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

Exact distribution-free tests of mean-variance efficiency

Sermin Gungor, Richard Luger *

Department of Economics, Emory University, Atlanta, GA 30322-2240, USA

ARTICLE INFO

Article history:
Received 6 November 2008
Received in revised form 28 April 2009
Accepted 11 June 2009
Available online 21 June 2009

JEL classification:
C12
C14
C22
G11
G12

Keywords:
CAPM
Conditional heteroskedasticity
Non-parametric tests
Robust inference

ABSTRACT

This paper develops exact distribution-free tests of unconditional mean-variance efficiency. These new tests allow for unknown forms of non-normalities, conditional heteroskedasticity, and other non-linear temporal dependencies among the absolute values of the error terms in the asset pricing model. Exactness here rests on the assumption that the joint temporal error density is symmetric around zero. This still leaves open the possibility of return distribution asymmetry via coskewness with the benchmark portfolio. A simulation study shows that the new tests have very good power relative to that of many commonly used tests. The inference procedures developed are further illustrated by tests of the mean-variance efficiency of a market index using a 42-year sample of monthly returns on ten U.S. equity portfolios.

1. Introduction

The celebrated capital asset pricing model (CAPM) of Sharpe (1964) and Lintner (1965) extends the notion of a mean-variance efficient portfolio to the portfolio of all invested wealth—the market portfolio. A given portfolio is mean-variance efficient if it has the smallest possible variance of return given its expected return, or, more appropriately, if it has the largest expected return given its variance. This theory implies that expected excess returns on assets are linearly related to the slope, or beta, of their regression on the expected excess return of the benchmark portfolio. Here excess returns are those in excess of the riskless rate of return. Under mean-variance efficiency, the risk premium of an asset is a linear function of the asset's beta.

Empirical tests of the mean-variance efficiency hypothesis are usually conducted within the context of a multivariate linear regression. The application of tests based on asymptotic theory can lead to misleading conclusions as the approximation to the finite-sample distribution of test statistics can be quite poor, especially as the number of equations included in the system increases; see Shanken (1996), Campbell et al. (1997), Dufour and Khalaf (2002), and the numerical evidence presented here. The findings show that many standard parametric tests are unreliable, rejecting the null hypothesis of mean-variance efficiency far too often.

Gibbons et al. (1989) (GRS) propose a truly finite-sample test. The exact distribution theory for their multivariate *F*-test rests on the assumption that the regression error terms are independent over time and jointly normally distributed each period, conditional on the returns of the benchmark portfolio under test. These assumptions are at odds with some well-known facts about financial asset returns. Indeed, it has long been recognized that financial returns depart from Gaussian conditions; see Fama (1965), Blattberg and Gonedes (1974), and Hsu (1982). In particular, the distribution of asset returns appears to have fatter tails than those of a normal distribution.

* Corresponding author. Tel.: +1 404 727 0328; fax: +1 404 727 4639.
E-mail address: rluger@emory.edu (R. Luger).

0927-5398/$ – see front matter
doi:10.1016/j.jempfin.2009.06.003

Beaulieu et al. (2007) (BDK) propose an exact randomized likelihood-based test procedure that relaxes normality. Their framework assumes that the error distribution is known, or at least specified up to some unknown nuisance parameters. If normality is assumed, the BDK test becomes the simulation-based equivalent of the GRS test. More generally, the BDK test procedure can be thought of as an exact parametric bootstrap. When nuisance parameters are present, the computational cost of the BDK test procedure grows with the number of nuisance parameters since it then involves finding maximal p-values over a confidence region for the intervening nuisance parameters. The first-step confidence region is established by inverting a simulation-based goodness-of-fit test (of the assumed distribution), which involves a grid search over the nuisance parameter space.

Without such parametric assumptions it would seem difficult to derive an exact finite-sample distribution theory. Despite this apparent difficulty, we propose in this paper new non-randomized tests that are exact in finite samples without any parametric assumptions about the distribution of the error terms in the simple multivariate linear regression model. The methods exploit results derived by Luger (2003) in the context of testing for a random walk. Here we propose three testing approaches for joint inference on several parameters that differ mainly by what is assumed about the covariance of the errors across equations.

The first approach is an induced test procedure that allows for arbitrary covariances in the cross-section of error terms. The price to pay for this extra flexibility is that those tests can be conservative and lead to power losses if the number of test assets is large. The second approach assumes that the errors are independent across equations, conditional on the returns of the benchmark portfolio. This delivers tests with the correct size and more power than the induced tests. The third approach is based on a simple linear combination of the test assets (or portfolios of test assets) and, like the second approach, provides a test procedure with the correct size no matter the number of included assets. These single-portfolio tests allow some forms of covariation in the cross-section of error terms. The number of assets in the cross-section may even exceed the number of time-series observations, making these tests particularly attractive when testing mean-variance efficiency with many test assets or when the portfolios have relatively short histories. This stands in contrast to extant approaches based on estimates of the covariance matrix of the regression errors. In order to avoid singularities, those approaches require the size of the cross-section be less than that of the time series. The distribution-free approaches to inference developed here do not require the error covariance matrix.

The proposed distribution-free (or non-parametric) tests have several appealing features, since they are built on the mere assumption that the joint temporal error density is symmetric around zero. This means that no restrictions are placed on the degree of non-normality or the degree of heterogeneity across marginal distributions. In fact, the existence of moments need not be assumed for the validity of the new tests. It is important to note that this framework still leaves open the possibility of asymmetries in the distribution of test asset returns via coskewness with the benchmark portfolio.

Asset returns typically display clear patterns of volatility clustering for which generalized autoregressive conditional heteroskedasticity (GARCH) models are often used; see Bollerslev (1986). The tests proposed here allow not only for non-normalities, but also for *unknown* forms of conditional heteroskedasticity and other intertemporal dependencies among the absolute values of the error terms in the asset pricing model. Such forms of intertemporal dependencies invalidate the exact statistical theory underlying the parametric GRS and BDK test procedures. Since it is well known that asset returns depart from homogeneous conditions, the new tests of the mean-variance efficiency hypothesis developed here offer a valid and useful distribution-free testing alternative to potentially misleading parametric procedures.

Section 2 of the paper presents the framework, the hypothesis of interest, and the extant GRS and BDK test procedures that are exact under parametric distributional assumptions. Section 3 describes the building blocks of our three approaches to distribution-free inference. Section 4 is divided into three subsections, each one describing a proposed test procedure. Section 5 begins by presenting some additional (asymptotic) extant tests of mean-variance efficiency and then presents the results of some simulation examples to illustrate the behavior of the proposed tests relative to the commonly used procedures. Section 6 provides an empirical illustration of the new tests in the context of the Sharpe–Lintner version of the CAPM. Section 7 concludes.

2. Exact parametric tests

A benchmark portfolio with excess returns r_p is said to be mean-variance efficient with respect to a given set of N test assets with excess returns r_i, $i=1,\ldots,N$, if it is not possible to form another portfolio of those N assets and the benchmark portfolio with the same expected return as r_p but a lower variance, or equivalently, with the same variance but a higher expected return. More formally, portfolio p is mean-variance efficient if the following first-order condition is satisfied for the N test assets:

$$E[r_{it}] = \beta_i E[r_{pt}], \quad i = 1, \ldots, N, \tag{1}$$

where r_{it} and r_{pt} are the time-t returns on asset i and portfolio p, respectively, in excess of the riskless rate of return. The term β_i captures the degree of association between the expected excess return on the individual asset i and the expected excess return for portfolio p. Accordingly, assets with higher betas should offer in equilibrium higher expected returns. Consider the excess-return system of equations

$$r_{it} = a_i + \beta_i r_{pt} + \varepsilon_{it}, \quad t = 1, \ldots, T, \quad i = 1, \ldots, N, \tag{2}$$

where ε_{it} is a random error term for asset i in period t with the property that $E[\varepsilon_{it}]=0$. The specification in Eq. (2) is a seemingly unrelated equations model. The mean-variance efficiency condition in Eq. (1) can then be assessed by testing

$$H_0 : a_i = 0, \quad i = 1, \ldots, N, \tag{3}$$