Journal of Financial Economics 14 (1985) 349–357. North-Holland

A NOTE ON THE GEOMETRY OF SHANKEN'S CSR T^2 TEST FOR MEAN/VARIANCE EFFICIENCY

Richard ROLL*

University of California at Los Angeles, Los Angeles, CA 90024, USA

Received October 1984, final version received November 1984

Shanken (1985) derives a test for the zero-beta capital asset pricing model (CAPM) which, as he points out, is equivalent to a test of the mean/variance efficiency of the market portfolio. This note illustrates the geometry of Shanken's test in the mean/variance space.

1. Introduction

In an interesting contribution to the econometric literature of finance, Shanken (1985) derives a cross-sectional regression (CSR) test for the mean/variance efficiency of a particular market index. The test statistic is based on a quadratic form involving the sample covariance matrix and, as such, it is related to the classic Hotelling T^2 statistic; thus, the name CSR T^2 test. As Shanken points out, testing the mean/variance efficiency of a given portfolio is equivalent to testing the validity of the zero-beta capital asset pricing model with a given index employed as the 'market portfolio'. This note is really just an appendix to Shanken's paper. It derives the geometry of his test in mean/variance space. This geometry illustrates the intuition of Shanken's test.

2. The geometry of Shanken's test

The test statistic is a residual sum of squares from a generalized least squares (GLS) cross-sectional regression. Using Shanken's notation, the statistic is expressed as

$$Q = Te'\hat{V}^{-1}e, \tag{1}$$

*The comments and suggestions of Chi-Cheng Hsia, Bob Korkie, and Jay Shanken are gratefully acknowledged.

where

T = time series sample size,
e = vector of residuals from a generalized least squares (GLS) cross-sectional regression of mean returns of 'betas' for $N+1$ assets,
$\hat{V}$ = sample covariance matrix $[(N+1)\times(N+1)]$.

The residual vector e is computed as

$$e = \bar{R} - \hat{\gamma}_0\mathbf{1} - \hat{\gamma}_1\beta, \tag{2}$$

where $\hat{\gamma}_0$ and $\hat{\gamma}_1$ are GLS coefficients, and

$\bar{R}$ = vector of individual mean returns,
$\mathbf{1}$ = vector of ones,
β = vector of beta coefficients.

All three of these vectors are $[(N+1)\times 1]$.

The cross-sectional coefficients $\hat{\Gamma} = (\hat{\gamma}_0\ \hat{\gamma}_1)'$ are given by

$$\hat{\Gamma} = (X'\hat{V}^{-1}X)^{-1}(X'\hat{V}^{-1}\bar{R}), \tag{3}$$

(see Shanken), where

$$X = (\mathbf{1}\ \beta) = \left[\mathbf{1}\ (q'Vq)^{-1}Vq\right], \tag{4}$$

and where q is an $[(N+1)\times 1]$ column vector of the investment proportions of the particular market index being tested for efficiency.

Shanken analyzes this statistic in an 'errors-in-variables' framework, noting that the regressors in the cross-sectional regression (3) should be the true betas; the β in (4) is supposed to be computed from the true covariance matrix V, not from its estimate $\hat{V}$.

Operationally, the sample estimate of β is used in the cross-sectional regression. The sample estimate of beta is a function of $\hat{V}$, i.e., $\hat{X} \equiv (\mathbf{1}\ \hat{\beta}) = [\mathbf{1}\ (q'\hat{V}q)^{-1}\hat{V}q]$. This suggests that there is really no errors-in-variables problem in the usual econometric sense of the phrase. If we substitute for $\hat{X}$ in (3), betas disappear completely and Q can be expressed in terms of $\hat{V}$ and $\bar{R}$ alone. The sampling distribution of Q depends on the joint sampling distribution of the mean vector $\bar{R}$ and the covariance matrix $\hat{V}$. The intermediate step (3) can be finessed and, in doing so, the errors-in-variables problem is eliminated.[1] This is not to say that the problem is simple; indeed, the sampling

[1] This is true, however only when the index contains no individual assets other than those used in computing $\bar{R}$ and $\hat{V}$. The index need not contain all of these $N+1$ assets, but it must not contain others.

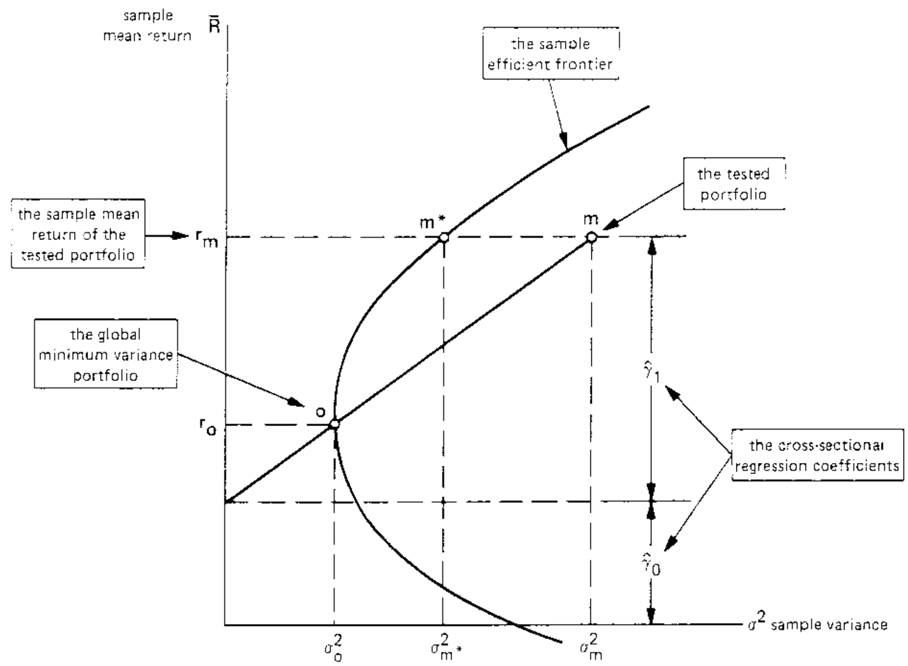

Fig. 1. The geometry of Shanken's test. *Note:* His test statistic is $Q = T\Omega(\sigma_m^2 - \sigma_{m^*}^2)/(\sigma_m^2 - \sigma_0^2)$, where T is time series sample size and the value of Ω is determined by the sample efficient frontier (see fig. 2). The CSR regression coefficients $\hat{\gamma}_0$ and $\hat{\gamma}_1$ have the values indicated.

distribution of the test is extremely complex. It is just not an errors-in-variables problem. Using $\hat{X}$ in (3) and thereby eliminating $\hat{\beta}$ also permits us to derive the geometric properties of the test statistic. The necessary algebra is given in the appendix.

Fig. 1 illustrates the geometry of Shanken's test. First note that Shanken's GLS regression coefficients have a natural interpretation. The coefficient $\hat{\gamma}_1$ is the 'excess return' on the market index m over its 'zero-beta' portfolio return. Since m is not exactly (sample) efficient, there are zero-beta portfolios at all return levels [Roll (1980)], and the particular one selected by the GLS regression is positioned as indicated by $\hat{\gamma}_0$.[2]

The test statistic Q can be expressed as

$$Q = T\Omega\left[\left(\sigma_m^2 - \sigma_{m^*}^2\right)/\left(\sigma_m^2 - \sigma_0^2\right)\right], \tag{5}$$

[2] The zero-beta portfolio whose return is $\hat{\gamma}_0$ is not itself on the mean/variance frontier. It is, however, the global minimum variance zero-beta portfolio for m. See Roll (1980, p. 1008).

where σ_m^2 is the index' sample variance, $\sigma_{m^*}^2$ is the sample variance of a portfolio that is (sample) mean/variance efficient and has the same mean return as m, and σ_0^2 is the sample global minimum variance. The scalar quantity Ω, a constant given the sample, will be discussed momentarily. As the illustration shows, the ratio $(\sigma_m^2 - \sigma_{m^*}^2)/(\sigma_m^2 - \sigma_0^2)$ lies between zero and one and it measures the horizontal distance of m from the sample efficient frontier, relative to a normalization factor, $\sigma_m^2 - \sigma_0^2$, which is the horizontal distance of m from the global minimum variance portfolio.

The reader can readily verify that this ratio makes a lot of intuitive sense. When m is far away from the sample frontier, Q is large (and this rejects the hypothesis that m is ex ante efficient). The normalization is important. For example, market indices near the global minimum variance portfolio must be very close indeed to the frontier to avoid being judged inefficient while indexes with large variance can be quite far from the frontier without being so penalized. This too is sensible because the sampling variability of a portfolio far away from the global minimum/variance is considerably larger.

The sample statistic Ω has its own interesting features. The value of Ω is determined entirely by the sample efficient frontier and is unrelated to the particular market index being tested. As shown in the appendix,

$$\Omega = b(r_1 - r_0), \qquad r_0 \neq 0,^3 \tag{6}$$

where r_1 is the return on a sample efficient portfolio positioned on a ray from the origin through the position of the sample global minimum variance portfolio (whose return is r_0). The slope of this ray is b. (See fig. 2.)

Fig. 2 illustrates the effect of the components of Ω on the overall test statistic Q. The larger b, the larger Shanken's Q. As the figure shows, b is larger when the global minimum variance portfolio is closer to zero, i.e., when there is less correlation among returns and less variability in general. But given b, Shanken's statistic is larger the further is the special portfolio 1 from the minimum variance portfolio. In other words, the more acute the curvature of the sample frontier (dashed curve), the lower Q and the less likely the index is to be judged inefficient. The curvature is influenced by correlations among individual returns. The relative value of time series and cross-sectional sample sizes (T and $N+1$, respectively) partially determine the values of $r_1 - r_0$ and b. As $(N+1) \to T$, $r_1 \to r_0$ and $b \to \infty$; naturally, the test has no power when there are as many assets as time series observations and the statistic cannot then be computed.

[3] In the special case $r_0 = 0$, $\Omega = \bar{R}'\hat{V}^{-1}\bar{R}$.

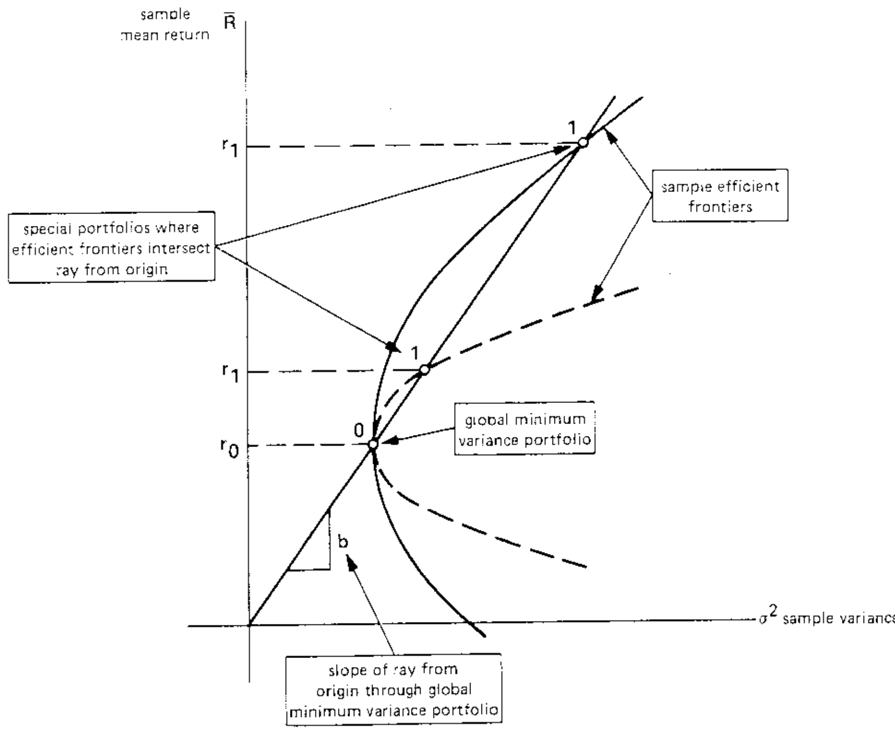

Fig. 2. The effect of the position of the sample efficient frontier on Shanken's test statistic. *Note:* Dashed curve indicates a particular sample frontier and solid curve a different sample frontier (both with the same global minimum variance point). In Shanken's test, the parameter Ω is given by $b(r_1 - r_0)$.

3. Remarks

The geometry illustrates why small sample exact distributions for tests of portfolio efficiency are difficult to obtain. The ratio $(\sigma_m^2 - \sigma_{m^*}^2)/(\sigma_m^2 - \sigma_0^2)$ characterizes m's position relative to the sample efficient frontier; but the sample frontier is itself a random variable. In the decomposition of Q given by (5), the first term Ω is entirely attributable to the sample efficient frontier and is unrelated to the index being tested. We can see in figs. 1 and 2 that the Q statistic is determined not only by the index' (m's) position in mean/variance space, but also by the position of the sample frontier. The latter constitutes the main problem in determining the finite sampling distribution of Q. The sampling distribution of the position of m is relatively easy to derive.

If the asymptotic (chi-square) distribution of Q is used to test the index' efficiency, one is likely to reject efficiency too often, as Shanken emphasizes. The geometry helps elucidate why this happens. The positively-sloped segment of the sample efficient frontier is likely to be positioned to the left and above the positively-sloped segment of the true frontier. This is due to the frontier being the solution to a minimization problem. It traces out minimum variance portfolios *in the sample*. But the sample contains random errors so, roughly speaking, the sample frontier minimizes over both population values and sampling errors.

To help alleviate this problem, Shanken presents an adjusted Q statistic which does not reject too often asymptotically, i.e., as $(T-N)\to\infty$. However, the adjusted statistics still reject too often in small samples.[4] Jobson and Korkie (1980) illustrate this small sample problem by showing that the expected value of the sample slope, $\hat{b}$, of the ray shown in fig. 2 exceeds the population slope, $E(\hat{b}) = [(T-1)/(T-N-1)]b$, and that the 'Sharpe ratio', the slope of the tangent line to the frontier in mean/standard deviation space, has a similar upward bias. See also Jobson and Korkie (1982).

A test related to Shanken's[5] was derived by Kandel (1984). [See also Gibbons (1982).] This test involves a likelihood ratio computed from the sample position of the index in mean/variance space relative to the sample efficient frontier. The test statistic is T times the logarithm of the ratio Q_K where

$$\frac{\sigma_m^2}{\sigma_{m^*}^2} > Q_k \equiv \frac{(r_0-\gamma^*)/\sigma_0^2}{(r_m-\gamma^*)/\sigma_m^2} > 1. \tag{7}$$

This is illustrated in fig. 3. Kandel proves that γ^* is less than Shanken's CSR regression coefficient $\hat{\gamma}_0$ and greater than a lower bound γ_2 which, as fig. 3 indicates, is positioned where a line from m^* through 0 intersects the return axis. The geometry can be used to prove that Q_K lies between the bounds given by (7). The exact value of γ^* can be computed from the sample but there does not seem to be a convenient analytic representation.

Appendix: Derivation of the geometric properties of Shanken's estimator Q

Shanken's T^2 estimator is

$$Q = Te'\hat{V}^{-1}e,$$

where

$$e \equiv \overline{R} - X\hat{\Gamma}, \qquad \hat{X} \equiv [\mathbf{1}\ \hat{\beta}] = [\mathbf{1}\ k\hat{V}q],$$

[4] Shanken derives a small sample approximation to the adjusted Q statistic.

[5] Shanken's test is a transformation of the *one-step* version of the likelihood ratio test. See Gibbons (1982, p. 10).

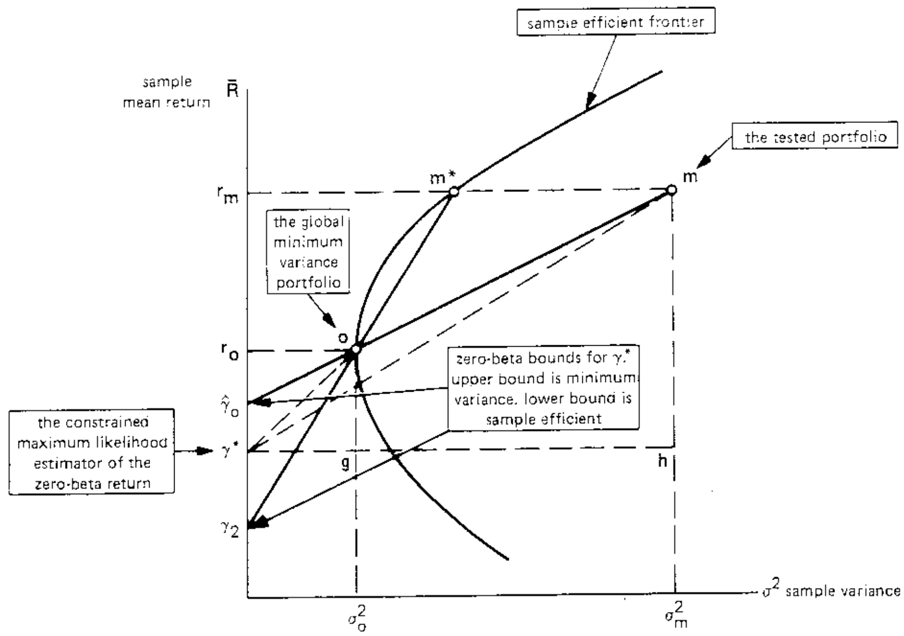

Fig. 3. The Gibbons/Kandel likelihood ratio test. *Note:* The test statistic involves the similarity of the two triangles outlined by dashed lines ($\gamma^* 0 g$ and $\gamma^* mh$).

and

k = $(q'\hat{V}q)^{-1} = 1/\sigma_m^2$,
q = investment proportions vector $[(N+1)\times 1]$ of market index m, where $N+1$ is number of assets,
T = time series sample size,
$\mathbf{1}$ = $[(N+1)\times 1]$ vector of ones,
$\hat{V}$ = sample covariance matrix,
$\hat{\Gamma}$ = $(\hat{\gamma}_0\hat{\gamma}_1)$ = Shanken's estimators for the zero-beta return ($\hat{\gamma}_0$) and the market's excess return ($\hat{\gamma}_1$), $\hat{\Gamma} = (\hat{X}'\hat{V}^{-1}\hat{X})^{-1}(\hat{X}'\hat{V}^{-1}\overline{R})$,
$\overline{R}$ = $[(N+1)\times 1]$ vector of individual asset mean returns,
r_m = $q'\overline{R}$ is the (scalar) mean return on the index.

For simplicity of notation, drop the hats on V and Γ.

Thus,

$$Q/T = (\overline{R} - X\Gamma)'V^{-1}(\overline{R} - X\Gamma)$$
$$= \overline{R}'V^{-1}\overline{R} + \Gamma'X'V^{-1}X\Gamma - 2\overline{R}'V^{-1}X\Gamma$$
$$= a - \overline{R}'V^{-1}X(X'V^{-1}X)^{-1}X'V^{-1}\overline{R}.$$

where

$$a = \bar{R}'V^{-1}\bar{R}.$$

Note that a is one of the efficient set constants [Roll (1977, app.)].

Since $X = [\mathbf{1}\ kVq]$, we can simplify the expression above to

$$Q/T = a - [b\ kr_m]\begin{bmatrix} c & k \\ k & k \end{bmatrix}^{-1}\begin{bmatrix} b \\ kr_m \end{bmatrix},$$

where $b = \mathbf{1}'V^{-1}\bar{R}$ and $c = \mathbf{1}'V^{-1}\mathbf{1}$ are the other two 'efficient set constants'. Simplifying again,

$$Q/T = a - \frac{b^2 - 2bkr_m + ckr_m^2}{c-k}$$

$$= \frac{ac-b^2}{c-k}\left[1 - \frac{a - 2br_m + cr_m^2}{ac-b^2}k\right].$$

The expression $\sigma_{m*}^2 = (a - 2br_m + cr_m^2)/(ac-b^2)$ can be recognized as the variance of the efficient portfolio whose mean return is the same as r_m. Simplifying again,

$$Q = T\sigma_0^2(ac-b^2)\frac{\sigma_m^2 - \sigma_{m*}^2}{\sigma_m^2 - \sigma_0^2},$$

where we have used the fact that $\sigma_0^2 = 1/c$ is the variance of the (sample) global minimum variance portfolio. It can be easily shown [Roll (1977, app.)] that

$$\Omega \equiv \sigma_0^2(ac-b^2) = b(r_1 - r_0),$$

where b is the slope of a ray from the origin through the global minimum variance portfolio and r_1 is the return of the other efficient portfolio (besides 0) which lies along this same ray.

The geometric interpretation of Shanken's estimators $\hat{\gamma}_0$ and $\hat{\gamma}_1$ in the cross-sectional regression of mean returns on betas can be derived directly from his formulae,

$$\hat{\Gamma} = (X'V^{-1}X)^{-1}(X'V^{-1}\bar{R}).$$

Substituting for X, we have

$$\hat{\Gamma} = \frac{1}{k(c-k)}\begin{bmatrix} k & -k \\ -k & c \end{bmatrix}\begin{bmatrix} b \\ kr_m \end{bmatrix}.$$

or

$$\hat{\Gamma} = \begin{pmatrix} \hat{\gamma}_0 \\ \hat{\gamma}_1 \end{pmatrix} = \begin{pmatrix} b - kr_m \\ cr_m - b \end{pmatrix} \Big/ (c-k).$$

Simplifying,

$$\hat{\Gamma} = \begin{bmatrix} \dfrac{\bar{r}_0\sigma_m^2 - \bar{r}_m\sigma_0^2}{\sigma_m^2 - \sigma_0^2} \\ \sigma_m^2\left(\dfrac{\bar{r}_m - \bar{r}_0}{\sigma_m^2 - \sigma_0^2}\right) \end{bmatrix}.$$

It is easily shown that these estimates are positioned as indicated in fig. 1 of the text.

References

Gibbons, Michael, 1982, Multivariate tests of financial models: A new approach, Journal of Financial Economics 10, 3–27.

Jobson, J.D. and B. Korkie, 1980, Estimation for Markowitz efficient portfolios, Journal of the American Statistical Association 75, 544–554.

Jobson, J.D. and B. Korkie, 1982, Potential performance and tests of portfolio efficiency, Journal of Financial Economics 10, 433–466.

Kandel, Shmuel, 1984, The likelihood ratio test statistic of mean-variance efficiency without a riskless asset, Journal of Financial Economics 13, 575–592.

Roll, Richard, 1977, A critique of the asset pricing theory's tests, Journal of Financial Economics 4, 129–176.

Roll, Richard, 1980, Orthogonal portfolios, Journal of Financial and Quantitative Analysis 15, 1005–1023.

Shanken, Jay, 1985, Multivariate tests of the zero-beta CAPM, Journal of Financial Economics 14, this issue.