The Market Portfolio May Be Mean/Variance Efficient After All

Moshe Levy
Jerusalem School of Business Administration, Hebrew University

Richard Roll
UCLA Anderson

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

Numerous studies have examined the mean/variance efficiency of various market proxies by employing sample parameters and have concluded that these proxies are inefficient. These findings cast doubt about the capital asset pricing model (CAPM), one of the cornerstones of modern finance. This study adopts a reverse-engineering approach: given a particular market proxy, we find the minimal variations in sample parameters required to ensure that the proxy is mean/variance efficient. Surprisingly, slight variations in parameters, well within estimation error bounds, suffice to make the proxy efficient. Thus, many conventional market proxies could be perfectly consistent with the CAPM and useful for estimating expected returns. (*JEL* G11, G12)

Testing the capital asset pricing model (CAPM) of Sharpe (1964) and Lintner (1965) is equivalent to testing the mean/variance efficiency of the market portfolio (see Roll 1977 and Ross 1977). The efficiency of the market portfolio has very important implications regarding the debate over passive versus active investing and regarding the use of betas for pricing risky assets. Many studies that have examined the mean/variance efficiency of various market proxies have found that these proxies are inefficient and typically far from the efficient frontier.[1] Moreover, portfolios on the efficient frontier typically involve many short positions,[2] which implies, of course, that the positive-by-definition market portfolio cannot be efficient. These results hold both with sample parameters and with parameters adjusted by various shrinkage methods.[3] This

We are grateful to Matthew Spiegel, the Editor, and an anonymous referee for their very helpful comments and suggestions. Financial support from the Zagagi Fund is gratefully acknowledged. Send correspondence to Richard Roll, UCLA Anderson, 110 Westwood Plaza, Los Angeles, CA 90095; telephone: (310) 825-6118. E-mail: rroll@anderson.ucla.edu.

[1] See, for example, Gibbons (1982), Jobson and Korkie (1982), Shanken (1985), Kandel and Stambaugh (1987), Gibbons, Ross, and Shanken (1989), Zhou (1991), and MacKinlay and Richardson (1991).

[2] As shown, for example, by Levy (1983), Green and Hollifield (1992), and Jagannathan and Ma (2003).

[3] Jagannathan and Ma (2003) show that constraining the weights of the minimum-variance portfolio to be non-negative is equivalent to modifying the covariance matrix in a way which typically shrinks the large elements of the covariance matrix. When this shrinkage is employed, however, only a small number of assets are held in positive proportions (and the rest have weights of zero). This is, again, not an encouraging result for the hope of finding an efficient market portfolio by employing shrinkage techniques.

© The Author 2010. Published by Oxford University Press on behalf of The Society for Financial Studies. All rights reserved. For Permissions, please e-mail: journals.permissions@oxfordjournals.org.
doi:10.1013/rfs/hhp119 Advance Access publication January 5, 2010

constitutes a very dark cloud hanging over one of the most fundamental models of modern finance. In light of this evidence, should the CAPM be taken seriously or is it just a pedagogical tool for finance classes, grossly inconsistent with the empirical evidence?

This article shows that a small variation of the sample parameters, well within their estimation error bounds, can make a typical market proxy efficient. Thus, the empirically measured return parameters and the market portfolio weights are perfectly consistent with the CAPM using a typical proxy. How is this possible, and how can it be reconciled with the many previous studies that have shown that the market proxy is inefficient? While most studies suggest various variations of the return parameters relative to the sample parameters and check whether these variations lead to an efficient market proxy, we take a reverse approach: We first require that the return parameters ensure that the market proxy is efficient. Given this requirement, we look for parameters that are as "close" as possible to their sample counterparts. Surprisingly, parameters that make the market proxy efficient can be found very close to the sample parameters. Hence, minor changes in estimation error reverse previous negative and disappointing findings for the CAPM.

We hasten to add that the efficiency, or lack thereof, for a market proxy can never be a definitive test of the macro-CAPM, which requires the market portfolio of all assets, including real estate, human capital, etc. Nonetheless, it would be reassuring if typical proxies were less inefficient than previously believed.

This article is organized as follows. The next section introduces the methods employed. Section 2 describes the data and the results. Section 3 provides a detailed comparison of our results with the classical results in the literature. Section 4 discusses implications for asset pricing. Section 5 concludes.

1. Methods

Given a market proxy, m, we look for the "minimal" variation of sample parameters that would make it mean/variance efficient. Denote the vector of market proxy portfolio weights by x_m and denote the vector of sample average returns and the vector of sample standard deviations by μ^{sam} and σ^{sam}, respectively. C^{sam} denotes the sample covariance matrix, and ρ^{sam} denotes the sample correlation matrix.

The objectives being sought are an expected return vector μ and a covariance matrix C that on the one hand make portfolio m mean/variance efficient and on the other hand are as close as possible to their sample counterparts. For simplicity, when considering the covariance matrix C, we allow

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

variation only in the standard deviations, while retaining the same sample correlations:

$$\begin{bmatrix} \\ & C & \\ \\ \end{bmatrix} = \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & & \cdots 0 & \sigma_N \end{bmatrix} \begin{bmatrix} \\ & \rho^{sam} & \\ \\ \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & & \cdots 0 & \sigma_N \end{bmatrix}. \tag{1}$$

Allowing the correlations to vary as well introduces technical difficulties, but can only make the results stronger, as it allows more degrees of freedom in the optimization procedure described below.

In order to obtain the parameters (μ, σ) that are "closest" to their sample counterparts $(\mu^{sam}, \sigma^{sam})$, we define the following distance measure D between any parameter set (μ, σ) and the sample parameter set:

$$D\big((\mu, \sigma), (\mu, \sigma)^{sam}\big) \equiv \sqrt{\alpha \frac{1}{N} \sum_{i=1}^{N} \left(\frac{\mu_i - \mu_i^{sam}}{\sigma_i^{sam}} \right)^2 + (1-\alpha) \frac{1}{N} \sum_{i=1}^{N} \left(\frac{\sigma_i - \sigma_i^{sam}}{\sigma_i^{sam}} \right)^2}, \tag{2}$$

where N is the number of assets, and $0 \le \alpha \le 1$ is a parameter determining the relative weight assigned to deviations of the means relative to deviations of the standard deviations. Recall that the larger the standard deviation of a given asset's returns, the larger the statistical errors involved in estimating this asset's parameters, and the larger the confidence intervals for these parameters. This is the rationale for dividing the deviations in Equation (2) by σ_i^{sam}—the resulting distance measure "punishes" deviations in the parameters of assets with low standard deviations more heavily than similar deviations in assets with higher standard deviations. The ultimate test is whether a set of parameters (μ, σ) can be considered as "reasonably close" to the sample parameters: for example, one can look at the proportion of parameters that deviate from the standard estimation error bounds around their sample counterparts and the size of those deviations. Intuitively, a parameter set can be considered "reasonably close" when 95% or more of the parameters are within the 95% confidence intervals of the sample parameters (below we also employ more formal multivariate tests). The choice of the distance measure D in Equation (2) and its minimization in the optimization problem described below are designed to minimize the statistical significance of the deviations between μ and σ and their sample counterparts, but we should stress that the statistical

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

conclusion regarding the compatibility of the parameters (μ, σ) with the sample parameters is independent of the choice of D.

To find the set of parameters (μ, σ) that make the proxy m mean/variance efficient and are closest to the sample parameters, we solve the following optimization problem:

Optimization Problem 1:

Minimize D $\left((\mu, \sigma), (\mu, \sigma)^{sam}\right)$

Subject to:

$$\text{(i)} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & & \cdots 0 & \sigma_N \end{bmatrix} \begin{bmatrix} & & \\ & \rho^{sam} & \\ & & \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & & \cdots 0 & \sigma_N \end{bmatrix} \begin{bmatrix} x_{m1} \\ x_{m2} \\ \vdots \\ \\ x_{mN} \end{bmatrix} = q \cdot \begin{bmatrix} \mu_1 - r_z \\ \mu_2 - r_z \\ \vdots \\ \\ \mu_N - r_z \end{bmatrix},$$

where $q > 0$ is the constant of proportionality, and r_z is the zero-beta rate. Both q and r_z are free variables in the optimization. Thus, there are $2N + 2$ variables in the optimization: $N\mu$'s, $N\sigma$'s, q, and r_z. Any set of these $2N + 2$ parameters satisfying (i) makes the proxy portfolio mean/variance efficient (see, for example, Roll 1977). We are looking for the set of parameter vectors (μ^*, σ^*) that satisfy this mean/variance efficiency condition and are closest to the sample parameters.[4]

Our approach differs from those employed in previous studies, such as Black, Jensen, and Scholes (1972) and Gibbons, Ross, and Shanken (1989), in two main regards. First, we are not required to assume the existence of a risk-free asset. Second, and more importantly, the standard approach looks at the adjustment to the empirical average returns required to make the market proxy efficient (i.e., the stocks' alphas) and asks whether these adjustments are statistically plausible. In contrast, we are looking at simultaneous adjustments to the average returns and the standard deviations (and could, in principle, include adjustments to the correlations as well). Thus, while the standard approach examines the statistical plausibility of a single vector of alphas, we examine a multitude of vectors of average return and standard deviation adjustments. This allows us many more degrees of freedom relative to the standard approach and explains why we find that only small adjustments are required to make the

[4] This optimization problem is similar in spirit to Sharpe's (2007) "reverse optimization" problem. Levy (2007) employs an analogous technique to find mean/variance efficient portfolios that have all-positive weights. This approach was first used in a very innovative paper by Best and Grauer (1985).

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

market proxy efficient. In Section 3, we discuss the relation of our results to the previous literature in more detail.

In some situations, one may have beliefs about the proxy portfolio's *ex ante* mean and standard deviation and would like to find the set of parameters that are closest to the sample parameters and at the same time ensure that the proxy portfolio is mean/variance efficient with the prespecified mean and standard deviation. Denoting the prespecified mean and standard deviation by μ_0 and σ_0, respectively, the optimization problem solved in this case is:

Optimization Problem 2:

Minimize $D\left((\mu, \sigma), (\mu, \sigma)^{sam}\right)$

Subject to:

$$\text{(i)} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & & \cdots 0 & \sigma_N \end{bmatrix} \begin{bmatrix} & & \\ & \rho^{sam} & \\ & & \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & & \cdots 0 & \sigma_N \end{bmatrix} x_m = q \cdot \begin{bmatrix} \mu_1 - r_z \\ \mu_2 - r_z \\ \vdots \\ \\ \mu_N - r_z \end{bmatrix}$$

$$\text{(ii)} \ x'_m \mu = \mu_0$$

$$\text{(iii)} \ x'_m \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & & \cdots 0 & \sigma_N \end{bmatrix} \begin{bmatrix} & & \\ & \rho^{sam} & \\ & & \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & & \cdots 0 & \sigma_N \end{bmatrix} x_m = \sigma_0^2,$$

where, again, x_m is the vector of a given proxy's portfolio weights.

The next section presents solutions to these optimization problems with empirical equity data in order to ascertain how large the deviations from the sample parameters must be in order to ensure mean/variance efficiency.

2. Data and Results

Our demonstration sample consists of the one hundred largest stocks in the U.S. market (according to December 2006 market capitalizations), which have complete monthly return records over the period January 1997 to December

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

Table 1
The sample parameters and closest parameters ensuring that the market proxy is mean/variance efficient

(1) Stock (i)	(2) μ_i^{sam}	(3) μ_i^*	(4) σ_i^{sam}	(5) σ_i^*	(6) t-Value μ_i^*	(7) $(\sigma_i^*)^2/(\sigma_i^{sam})^2$ (The 95% confidence interval for this value is [0.790–1.319])
1	0.024	0.018	0.165	0.167	−0.423	1.019
2	0.021	0.019	0.115	0.115	−0.170	1.003
3	0.011	0.017	0.106	0.104	0.588	0.963
4	0.029	0.023	0.158	0.160	−0.444	1.028
5	0.039	0.022	0.150	0.156	−1.228	1.077
6	0.005	0.011	0.075	0.073	0.952	0.953
7	0.007	0.013	0.072	0.070	0.938	0.942
8	0.012	0.010	0.051	0.052	−0.433	1.028
9	0.013	0.015	0.070	0.069	0.286	0.978
10	0.016	0.018	0.099	0.098	0.185	0.986
11	0.010	0.013	0.067	0.066	0.344	0.977
12	0.016	0.009	0.092	0.093	−0.819	1.025
13	0.015	0.011	0.071	0.072	−0.627	1.035
14	0.019	0.012	0.100	0.102	−0.702	1.034
15	0.011	0.011	0.061	0.061	−0.029	1.006
16	0.032	0.014	0.159	0.162	−1.215	1.044
17	0.023	0.025	0.158	0.157	0.145	0.990
18	0.024	0.021	0.146	0.147	−0.232	1.016
19	0.011	0.012	0.086	0.085	0.199	0.988
20	0.007	0.010	0.067	0.066	0.477	0.979
21	0.011	0.011	0.065	0.065	0.082	0.996
22	0.018	0.016	0.080	0.081	−0.225	1.018
23	0.012	0.008	0.067	0.068	−0.652	1.023
24	0.013	0.004	0.059	0.059	−1.533	0.995
25	0.017	0.014	0.088	0.088	−0.361	1.021
26	0.014	0.013	0.081	0.082	−0.128	1.007
27	0.006	0.012	0.077	0.075	0.810	0.955
28	0.018	0.011	0.077	0.078	−1.058	1.044
29	0.010	0.012	0.087	0.086	0.276	0.989
30	0.010	0.010	0.065	0.064	0.055	0.999

For the sake of brevity, this table reports only thirty of the one hundred stocks (the complete table is given in the Appendix). The sample parameters are given in the second and fourth columns. The expected returns and standard deviations, which are closest to these parameters and ensure that the market proxy is efficient (i.e., the parameters that solve Optimization Problem 1), are given in columns 3 and 5. The t-values for the expected returns are given in column 6, which shows that none of these values are significant at the 95% level (this is also true for the seventy other stocks not shown in the table). Column 7 reports the ratio between the optimized variances $(\sigma^*)^2$ and the sample variances. The 95% confidence interval for this ratio is [0.790–1.319] (see footnote 5). All of the ratios in the table, as well as the ratios for all other seventy stocks not shown here, fall well within this interval. These results are obtained with a value of $\alpha = 0.75$ in the minimized distance measure D (see Equation (2)). Higher values of α reduce the variation in the expected returns (at the expense of increasing the deviations in the standard deviations).

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

2006 (120 return observations). Columns 2 and 4 in table 1 report the sample average returns and standard deviations for thirty of these stocks (the complete information for all one hundred stocks is given in table A1 in the Appendix). The average sample correlation is 0.24.

Following previous research (e.g., Stambaugh 1982), we examine a market proxy whose weights are market capitalizations, in this case of the one hundred stocks as of December 2006,

$$x_{mi} = \frac{market\ cap\ of\ \ firm\ i}{\sum_{j=1}^{100} market\ cap\ of\ \ firm\ j}.$$

The proxy portfolio and the sample mean/variance frontier are shown in figure 2 by the triangle and thin line, respectively. As the figure illustrates, the proxy portfolio is far from the efficient frontier when the sample parameters are employed. This is consistent with previous studies.

To solve Optimization Problem 1 numerically, we implement Matlab's fmincon function, which is based on the interior-reflective Newton method and the sequential quadratic programming method. The solution (μ^*, σ^*) is given in columns 3 and 5 of table 1.

2.1 Simple tests of significance between sample and adjusted parameters

t-values for the adjusted expected returns $\mu^*$are given in column 6 of table 1. They reveal that the difference between the sample average return, μ_i^{sam}, and μ_i^* is nonsignificant at the 95% level for all stocks (this is true not only for the thirty stocks shown in the table but also for the other seventy stocks as well). Column 7 provides the ratio $\left(\sigma_i^*\right)^2/\left(\sigma_i^{sam}\right)^2$ for each stock. The 95% confidence interval for this ratio is the range [0.790–1.319].[5] The values in column 7 reveal that for all stocks, the ratio $\left(\sigma_i^*\right)^2/\left(\sigma_i^{sam}\right)^2$ is well within this range (and this is also true for the seventy stocks not shown in the table). Thus, the solution (μ^*, σ^*) to the optimization problem is very close to the sample parameter set because no parameters is significantly different from its sample counterpart.

More formally, as we have $2N = 200$ parameters, we are simultaneously testing two hundred hypotheses (each stating that the given parameter is not

[5] The ratio $\frac{(n-1)s^2}{\sigma^2}$ is distributed according to the χ^2_{n-1}distribution, where σ^2 is the population variance, s^2 is the sample variance (or $(\sigma^{sam})^2$ in the notation used in this article), and n is the number of observations. We have 120 monthly return observations, hence n= 120. As we are looking for the 95% confidence interval for s^2/σ^2, we need to find the critical values c_1 and c_2 for which $P\left(\chi^2_{119} > c_1\right) = 0.025$, and $P\left(\chi^2_{119} < c_2\right) = 0.025$. For large n, $\sqrt{2\chi^2_n} - \sqrt{2n-1}$ can be approximated by the standard normal distribution. Thus, the critical values c_1 and c_2 satisfy $\sqrt{2c_1} - \sqrt{2 \cdot 119 - 1} = 1.96$ and $\sqrt{2c_2} - \sqrt{2 \cdot 119 - 1} = -1.96$, which yield $c_1 = 150.6$ and $c_2 = 90.2$. Thus, the 95% confidence interval for $s^2/\sigma^2$is given by $90.2 < 119 \cdot s^2/\sigma^2 < 150.6$ or $0.758 < s^2/\sigma^2 < 1.266$. Alternatively, this range can be also stated as $0.790 < \sigma^2/s^2 < 1.319$.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

different than its sample counterpart at the 5% significance level). The Bonferroni (1935) test states that we should reject the multiple-comparison hypothesis at the 5% level if any one of the parameters is significantly different than its sample counterpart at the (5/200)% level (see also Miller 1991). As none of our parameters are significantly different at the 5% level, of course, none is significant at the much lower (5/200)% level, and we cannot reject the multiple comparison hypothesis.

2.2 Multivariate tests of significance between sample and adjusted parameters

The univariate t-tests reported above and the Bonferroni multiple comparisons test rely on a questionable assumption, viz., that the estimation errors are independent across parameters. Because all sample estimates were obtained with data spanning the same calendar time period, some interdependence in estimation errors would not be all that surprising. To ensure that such a possibility did not seriously affect our inference that the subject portfolio was not statistically significantly off the efficient frontier, we carried out two further tests that take account of possible estimate dependence.

The first test assumes that the individual stock returns are drawn from a multivariate normal distribution. In this case, individual sample mean returns and the sample covariance matrix are jointly distributed as a noncentral Wishart (cf. Johnson and Kotz 1972, p. 175). To test a particular hypothesis about the true population mean returns and covariance matrix, we can employ the likelihood ratio, whose general form is

$$n \cdot \log \left\{ \frac{|S|}{|\Sigma|} - N + trace \left(\Sigma^{-1} \left(S + (\mu - \bar{x}) \, (\mu - \bar{x})\,' \right) \right) \right\},$$

where n is the number of time periods; N is the number of stocks; S and $\bar{x}$ are the sample covariance matrix and vector of sample mean returns, respectively; and Σ and μ are the corresponding hypothesized values.

Because our adjusted means and standard deviations are the hypothesized values, the likelihood ratio above can be calculated after making the following substitutions: $\mu = \mu^*$, $\Sigma = \text{diag}(\sigma^*)\rho^{sam}\text{diag}(\sigma^*)$, $\bar{x} = \mu^{sam}$, and $S = \text{diag}(\sigma^{sam})\rho^{sam}\text{diag}(\sigma^{sam})$, where diag($z$) is a diagonal matrix with the vector z along the diagonal and zeroes off the diagonal. (The sample correlation matrix ρ^{sam} is used in calculating Σ and S, because our optimization adjusted only the standard deviations while holding constant the correlations.)

In the general case (with unrestricted correlations), the likelihood ratio is asymptotically distributed as chi-square with $N + N(N+1)/2$ degrees of freedom. However, in our particular application, with unaltered correlations, there are only $2N$ degrees of freedom (for N means and N standard deviations).

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

The computed value for this likelihood ratio turns out to be 156.8, which is the 0.011 fractile of the chi-square distribution with two hundred degrees of freedom. Hence, one cannot reject the hypothesis that the sample means and standard deviations as a group of possibly correlated parameters are not significantly different from the group of their adjusted counterparts.[6]

Most asset returns, including those used here, exhibit thick tails relative to the normal distribution. Consequently, the sample means and standard deviations may not conform all that well to a noncentral Wishart distribution. We therefore decided to conduct one additional test using the bootstrap, which makes no distributional assumption but merely resamples from the original observations.

To carry out the bootstrap, we first adjust the empirical $T \times N$ return matrix (T monthly returns for N stocks) to create a "true" return matrix with parameters μ^* and σ^*. Then, we resample randomly from this return matrix and calculate the parameters (μ^{BS}, σ^{BS}) obtained in each random draw of T periods. For each draw, a "distance" is calculated between (μ^{BS}, σ^{BS}) and (μ^*, σ^*) and compared with the distance between $(\mu^{sam}, \sigma^{sam})$ and (μ^*, σ^*). If the bootstrap distance exceeds the original sample distance in a large fraction of cases, one can conclude that the sample and adjusted parameters are reasonably close.

Below are the step-by-step details:

1. The sample returns, $r_{i,t}$, are adjusted to create returns with the desired parameters (μ^*, σ^*) by the simple linear transformation $r^*_{i,t} = a_i + b_i r_{i,t}$, with $b_i = \sigma^*/\sigma^{sam}$ and $a_i = \mu^* - b_i \mu^{sam}$. (Obviously, the correlations are unaltered.) The adjusted returns are arranged in a matrix with T columns and N rows.
2. From this $(T \times N)$ matrix, T columns are drawn randomly with replacement, thus maintaining the underlying cross-sectional dependence,[7] and (μ^{BS}, σ^{BS}) are computed for this (re)sample.
3. The "distance" between the sample parameters (μ^{BS}, σ^{BS}) and the true parameters (μ^*, σ^*) is computed as the simple Euclidean distance[8]:

$$d \equiv \sqrt{\sum_{i=1}^{N} \left(\mu_i^{BS} - \mu_i^*\right)^2 + \sum_{i=1}^{N} \left(\sigma_i^{BS} - \sigma_i^*\right)^2}.$$

[6] Since the log likelihood ratio is only asymptotically chi-square, one cannot be certain that the sample size is large enough for a satisfactory convergence, though two hundred degrees of freedom is usually thought to be sufficient.

[7] The returns are assumed to be independent over time.

[8] One could employ various other more sophisticated distance measures (e.g., the distance D in Equation (2)). As will become evident below, the results are very strong, and they are robust to the distance measure employed. Obviously, we employ the same measure D for the distance between $(\mu^{sam}, \sigma^{sam})$ and (μ^*, σ^*) and between (μ^{BS}, σ^{BS}) and (μ^*, σ^*).

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

4. This distance is compared with the corresponding distance between the parameters $(\mu^{sam}, \sigma^{sam})$ and (μ^*, σ^*).

The distance between $(\mu^{sam}, \sigma^{sam})$ and (μ^*, σ^*) is 0.06. Of 10,000 resampled sets of T observations, all had a distance larger than this value. Figure 1 shows the distribution of the distance d obtained with the bootstrap.

It may seem suspicious that none of the bootstrap distances were smaller than the distance between the sample and adjusted values, but remember that the two types of distances are quite different in character. The latter, the distance between (sam) and (*), emerges from a portfolio optimization, whereas the former, the distance between (*) and (BS), is entirely attributable to statistical sampling error. There is no theoretical reason why one cannot be much smaller (or larger) than the other.

To think of it another way, suppose we had the exact same (sam) parameter values and therefore the same (*) values as well, but these were computed from 240 monthly returns rather than 120. In this case, the BS/* distances become smaller but the sam/* distance is unaltered. We actually redid the bootstrap using 240 observations per sample and found that twelve of 10,000 BS/* distances were smaller than the sam/* distance. This is still a very small number, but it is not zero, and it illustrates the fundamental difference between the two procedures.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

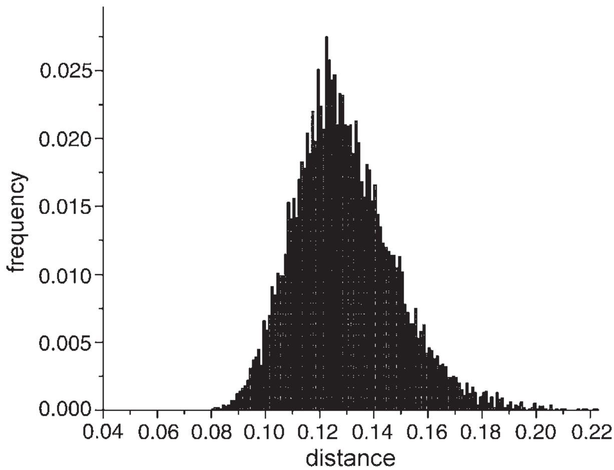

Figure 1
Probability distribution of Euclidean distance between bootstrapped and optimally adjusted parameters
The optimally adjusted parameters (means and standard deviations) are sufficient to make the proxy market portfolio lie on the (adjusted) mean/variance efficient frontier. The Euclidean distance between the adjusted parameters and the original sample parameters is 0.06. Ten thousand resampled sets of returns were drawn, and the Euclidean distance is calculated for each set. As shown above, all resampled distances lie above the sample distance.

Overall, it seems safe to conclude that statistically insignificant parameter adjustments can render our proxy portfolio efficient, even taking account of cross-sectional dependence in the underlying stock returns.

2.3 Interpretation of the results

To confirm that the parameters (μ^*, σ^*) make the proxy portfolio mean/variance efficient, one can examine the efficient frontier and the location of the proxy portfolio in the mean-standard-deviation plane with these parameters. These are illustrated by the bold line and the star in figure 2. The figure shows that with the parameters (μ^*, σ^*), the proxy portfolio lies on the efficient frontier. It is interesting to note that while the modified parameters (μ^*, σ^*) do not have a big impact on the expected return or the standard deviation of the proxy portfolio (the star is located very close to the triangle), they do have a big effect on the shape of the frontier. Why is the modified frontier much flatter than the sample frontier?

The explanation can be found in figure 3, which shows the adjustment to the expected return, $\mu_i^* - \mu_i^{sam}$, as a function of the sample average return, μ_i^{sam}. The figure reveals that high sample returns tend to get negative corrections $\left(\mu_i^* < \mu_i^{sam}\right)$, while the opposite holds for low sample returns. Thus, the cross-sectional variation of μ_i^* is smaller than the cross-sectional variation of μ_i^{sam}, which explains why the frontier is flatter (recall that in the limiting

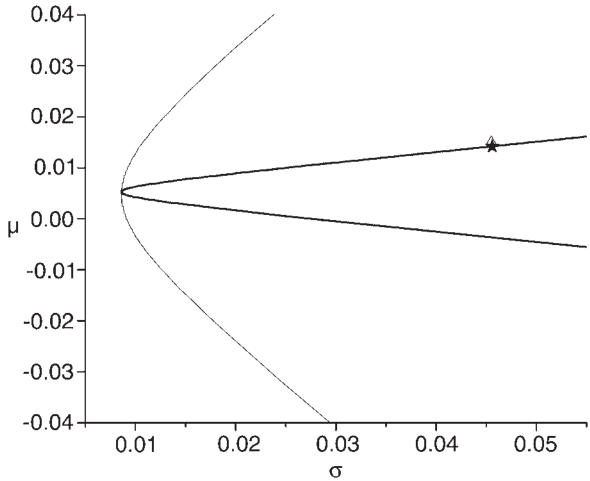

Figure 2
The efficient frontier and market proxy with the sample and the adjusted return parameters
The thin line curve and the triangle (partly hidden behind the star) show the mean/variance frontier and the market proxy with the sample parameters. As is typical of other studies, the market proxy is very far from the efficient frontier when the sample parameters are employed. The bold line and the star show the mean/variance frontier and the market proxy with the adjusted parameters (μ^*, σ^*). With these parameters, the market proxy is mean/variance efficient.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

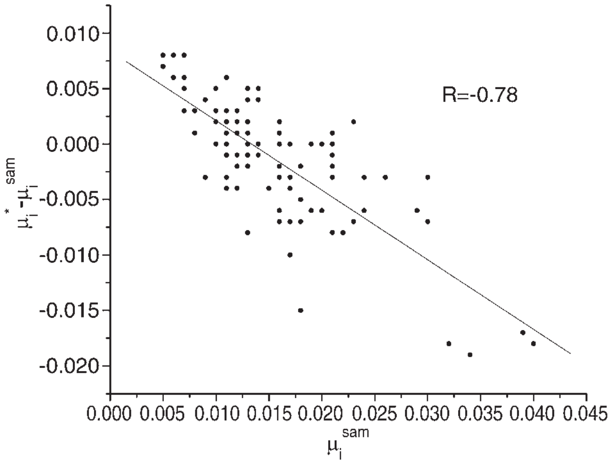

Figure 3
The correction to the estimated expected returns as a function of the sample average return
For stocks with high sample average returns, the correction in the expected return tends to be negative. The opposite holds for stocks with low sample average returns. Thus, the corrections produced by the solution to the optimization problem are reminiscent of statistical shrinkage methods.

case where all expected returns are identical, the frontier becomes completely flat, that is it is a horizontal line). Figure 3 shows that the corrections to the sample means implied by the optimization are reminiscent of standard statistical shrinkage methods. However, unlike the standard shrinkage methods, the method employed here ensures that the proxy is mean/variance efficient.[9]

There is excellent intuition behind such a result when one recalls two facts: (i) The efficient frontier itself is the result of an optimization, giving the minimum variance for each level of mean return; and (ii) sample parameter estimates are equal to true population parameters plus estimation errors. An efficient frontier computed using sample estimates optimizes with respect to sampling errors in addition to true parameters, so assets with overestimated means are likely to be weighted too heavily in frontier portfolios and *vice versa* for assets with underestimated means. This suggests that an efficient frontier computed using population parameters, if they were only known, would fall well inside the frontier computed using sample estimates, at least at most points. The main exception would be near the global minimum variance portfolio, whose weights do not depend on mean returns; indeed, such a relation is exactly what we see depicted in figure 2.

[9] One may wonder whether the adjustment $\mu^* - \mu^{sam}$ is similar for stocks that are relatively highly correlated with one another. In order to check this, we calculate the sample return correlation for each pair of stocks (i, j) and examine the relation across pairs between this sample correlation and the difference between the adjustments of the two stocks, that is $(\mu_i^* - \mu_i^{sam}) - (\mu_j^* - \mu_j^{sam})$. We find no such relation ($R^2 = 0.009$), that is pairs that are more highly correlated are not more likely to have similar adjustments.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

The implication of these results is quite striking. In contrast to "common wisdom," they show that the empirical proxy portfolio parameters are perfectly consistent with the CAPM if one allows for only slight estimation errors in the return moments. The reason that most previous studies have found that the market proxy is inefficient, even when various standard shrinkage methods have been employed, is that the variation of the parameters necessary to make the proxy portfolio efficient is very specific. While this variation is in the spirit of shrinkage, it is specifically designed to ensure the efficiency of the proxy portfolio, and thus it is fundamentally different than the standard statistical shrinkage methods.

With the solution (μ^*, σ^*) to Optimization Problem 1, the proxy portfolio has a monthly expected return of 1.4% and a standard deviation of 4.6% (see figure 2), which are very close to its sample values, 1.5% and 4.6%. These values were produced by the optimization (given the proxy portfolio weights). In some situations, one may have beliefs about the proxy portfolio's *ex ante* return parameters and may wish to look for solutions that are consistent with these beliefs. For example, suppose one would like to find vectors μ and σ such that the proxy portfolio is efficient and has an expected return and a standard deviation of $\mu_0 = 2\%$ and $\sigma_0 = 4\%$, respectively. Are such index values compatible (in a statistical sense) with the sample parameters and with a mean/variance efficient index? To answer this question, Optimization Problem 2 can be solved with $\mu_0 = 2\%$ and $\sigma_0 = 4\%$. We will consider the solution (μ^*, σ^*) compatible with the sample parameters if 95% or more of the parameters are within the 95% confidence intervals of their sample counterparts, and in addition, the adjusted parameters cannot be rejected by the bootstrap test. Of course, $\mu_0 = 2\%$ and $\sigma_0 = 4\%$ are just one example. A more complete picture would scan the mean/variance plane and map the range of proxy portfolios' return parameters, μ_0 and σ_0, that are compatible with the CAPM and the sample returns and market proxy weights.

Figure 4 shows the results of this analysis. For each combination of prespecified proxy portfolio parameters (μ_0, σ_0), we solve Optimization Problem 2. The points scanned are shown by the circles in the mean/variance plane. If the resulting optimal parameter set (μ^*, σ^*) is found to be statistically compatible with the sample parameters $(\mu^{sam}, \sigma^{sam})$ and with the CAPM (mean/variance efficiency of the index), the point is marked as a filled circle; if the parameters are rejected by the univariate test, the point is surrounded by a transparent circle; if the parameters are rejected by the bootstrap test, the point is surrounded by a diamond. For example, the point $(\mu_0 = 2\%,\ \sigma_0 = 4\%)$ (indicated by an up arrow in figure 4) is indeed consistent with the sample parameters and the proxy being efficient. In contrast, the point directly above $(\mu_0 = 2.5\%, \sigma_0 = 4\%)$ is rejected.

We should point out that for a given set of portfolio return parameters (μ_0, σ_0), our procedure produces "the best" portfolio with these parameters, in the sense that this is the portfolio that allows the adjusted parameters to be

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

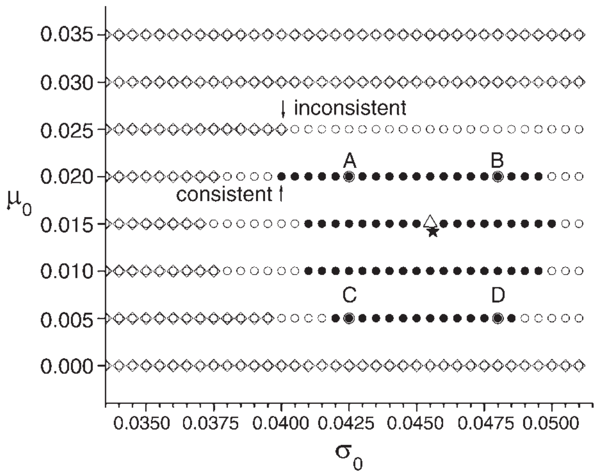

Figure 4
The set of proxy portfolio parameters consistent with mean/variance efficiency and the sample parameters: one hundred stocks
Optimization Problem 2 is solved for a lattice of points on the mean-standard deviation plane (μ_0, σ_0). The resulting parameter set (μ^*, σ^*) is considered consistent with the sample parameters by the univariate test if 95% or more of the parameters are within the 95% confidence intervals of their sample counterparts. Points that are rejected by the univariate test are surrounded by a transparent circle. Points that are rejected by the multivariate bootstrap test are surrounded by a transparent diamond. The (μ_0, σ_0) points that are consistent with the mean/variance efficiency of the proxy portfolio and with the sample parameters (i.e., they are not rejected by either of the tests) are indicated by the filled circles. For example, the proxy portfolio can be made mean/variance efficient with a standard deviation of 4% and a mean return of 2%, but not with a standard deviation of 4% and a mean return of 2.5%. The figure shows that given a set of sample parameters and proxy portfolio weights, the proxy portfolio can be made mean/variance efficient with a large range of possible mean and standard deviation combinations. As in figure 2, the triangle and the star represent the market proxy with the sample parameters and with the parameters solving Optimization Problem 1, respectively.

as close as possible to the sample parameters. Even if a given point (μ_0, σ_0) is in the "dark range" of consistency in the figure, this does not mean that any portfolio with these parameters is efficient (as there are many different portfolios with the same (μ_0, σ_0), and most of them may not be consistent with the sample returns and mean/variance efficiency). On the other hand, if a point (μ_0, σ_0) is in the inconsistent range, there is no portfolio with these parameters that can be consistent. Thus, there is no portfolio with parameters in the inconsistent range that can be moved onto the efficient frontier without violating a statistical p-value.[10]

It would be interesting to redo this analysis using indexes with even more individual assets, but there are technical difficulties. When the number of assets exceeds the number of time series observations, the correlation matrix is

[10] It is interesting to note that the range rejected by the univariate test is larger than the range rejected by the bootstrap test. This is probably because the bootstrap allows for estimation error independence and for departures from normality (because the t-tests are valid for the Gaussian only).

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

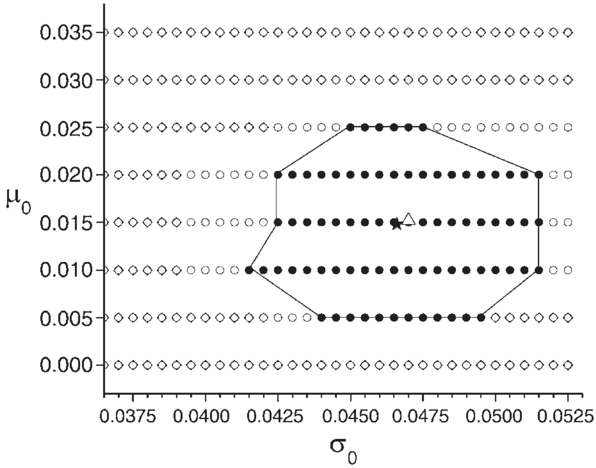

Figure 5
The set of proxy portfolio parameters consistent with mean/variance efficiency and the sample parameters: fifty stocks
This figure is the same as figure 4, but it is constructed with only the fifty largest stocks (rather than one hundred). Again, a wide range of (μ_0, σ_0) is consistent with the efficiency of the market proxy. The area of a polygon drawn through the outer consistent points is an approximation to the range of consistency.

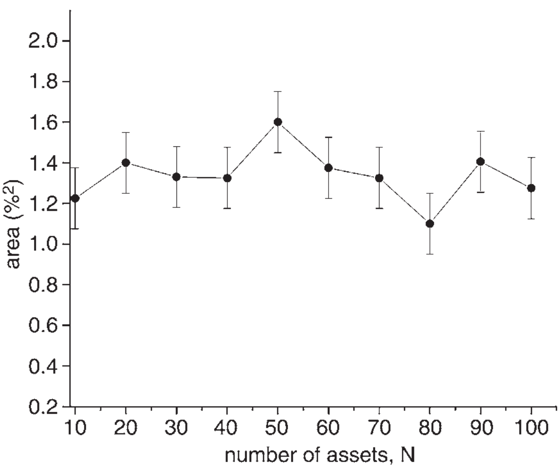

Figure 6
The area of admissible proxy portfolio parameters as a function of the number of assets
For each value of N, starting with the largest ten stocks, the area of a consistency polygon is computed analogous to the one shown in figure 5. This area measures the range of proxy portfolio return parameters consistent with the CAPM and the given proxy portfolio. This is an approximation of the precise area, because it depends on a finite set of parameter points in the mean/variance plane. The error bars reflect this possible estimation error. The figure shows that the area of admissible parameters does not change systematically with the number or the identity of the stocks included in the market index proxy.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

singular, which produces some instabilities in the optimization problem. We can, however, partially investigate this issue by varying the number of assets for $N < 100$ and looking for any trend in the range of proxy portfolio return parameters consistent with the CAPM.

For example, repeating the analysis for the fifty largest stocks (instead of the one hundred) yields the results shown in figure 5. These results are comparable to those obtained with one hundred stocks in the range of proxy portfolio return parameters consistent with the CAPM. To investigate in more detail possible systematic effects of the number of assets, we repeat this analysis for $N = 10$, $20, \ldots, 100$ stocks. For each value of N, we measure the area of admissible proxy portfolio parameters (estimated by the polygon containing the admissible points; see, for example, the polygon in figure 5). The results are shown in figure 6. Although the area is an approximation of the precise area of admissible points, because of the discreteness of the points (and as indicated by the error bars in the figure), figure 6 shows that the area does not seem to change systematically with the number of assets. Thus, the results seem robust to the identity of stocks and to the number of stocks contained in any market index proxy.

3. Detailed Comparison with Previous Results

Our results contradict the prevalent belief that the CAPM is inconsistent with the sample parameters. To understand better why we cannot reject the CAPM, in contrast to many previous studies, it is instructive to perform a detailed comparison with one of the classic studies considered by many as the most definitive rejection of the CAPM—the study by Gibbons, Ross, and Shanken (1989), hereafter GRS.

GRS develop an ingenious multivariate test of the CAPM, which has a very elegant and intuitive graphical interpretation, involving the Sharpe ratios of the proxy portfolio and the *ex post* tangency portfolio. Using monthly returns on twelve industry portfolios during 1926–1982, GRS reject the CAPM at the 1.3% significance level. How can our results be reconciled with this strong rejection?

First, one should note that the GRS test is quite sensitive to the choice of the risk-free rate and the length of the sample period. In order to demonstrate this, we repeat the analysis of GRS for the twelve industry portfolios (the monthly returns for these portfolios over the period 1926–2008 are taken from Ken French's data library[11]). For a monthly risk-free rate of 0.2%, the GRS test rejects the CAPM with a p-value of 0.03, consistent with the findings of GRS. However, the model cannot be rejected for a large range of risk-free rate values. Figure 7 shows the p-values obtained for different risk-free rates; it reveals that the CAPM is rejected at the 5% level only if the risk-free rate is below 0.003

[11] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

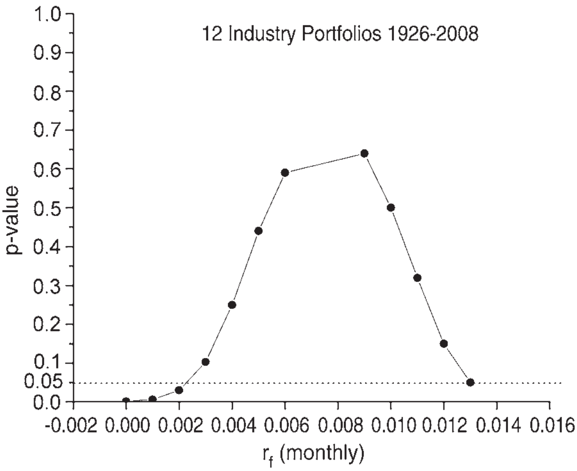

Figure 7
***p*-values in the GRS test for different values of the risk-free rate**
The GRS test is conducted for the twelve industry portfolios and the 990-month sample period of 1926–2008. While the CAPM can be rejected for very low or very high values of the risk-free rate, it cannot be rejected for the wide range of (monthly) interest rate values between 0.3% and 1.3%.

(0.3%) or above 0.013 (even with the very long 990-month sample period). If r_f is anywhere in between these two values, the model cannot be rejected.

Next, in order to examine the length of the time series required to reject the model, we take the value of $r_f = 0.002$ (which yields a p-value of 0.03 for the entire 990-month period of 1926–2008) and check the p-value for different sample sizes. Figure 8 reports the average p-value obtained over all nonoverlapping subperiods of a given length. The figure shows that approximately sixty years (about seven hundred months) are required in order to reject the model.

In light of these observations, it is hardly surprising that we cannot reject the CAPM: We are using only 120 monthly returns, and our test allows us to pick the risk-free rate which is "most favorable" for the model.

Note, however, that there is a key difference between the GRS test and the present approach, a difference that is unrelated to the length of the sample or the value of the risk-free rate. Our method allows for estimation errors not only in the average returns but also in the covariances and consequently in the betas. In order to focus on this effect, which is central to our article, consider a situation where one can reject the CAPM with the GRS test but cannot reject when the errors in covariances are taken into account, as in our test.

Following GRS, we take twelve industry portfolios, the entire 990-month period, and a risk-free rate of 0.002. As mentioned before, in this case, the GRS test rejects with a p-value of 0.03. Our method leads to a different conclusion. We calculate the adjusted parameters by solving Optimization 1 with

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

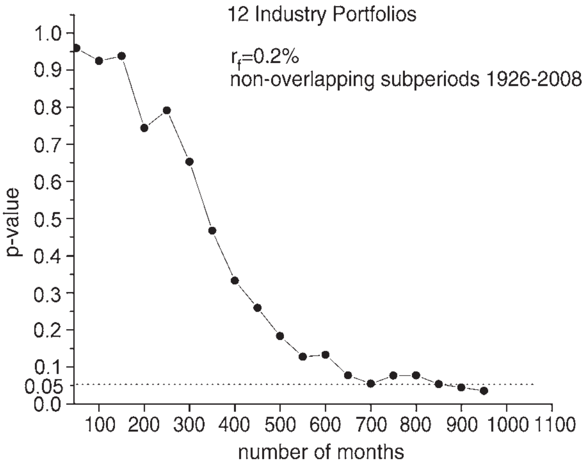

Figure 8
***p*-values in the GRS test for different sample period lengths**
The GRS test is conducted for the twelve industry portfolios taking a monthly risk-free rate of 0.2%. For any sample period length, we report the average *p*-value obtained over all nonoverlapping subperiods of this length. While the CAPM can be rejected for very long sample periods, it cannot be rejected with less than seven hundred months of data.

$r_z = 0.002$, and the market proxy is taken as the value weighted portfolio of the twelve industries at December 2008.[12] We find that all of the parameters fall within the 95% confidence intervals of their sample counterparts. As GRS correctly claim, the appropriate test is a multivariate test, but even with this in mind, the fact that all the parameters are well within the 95% confidence intervals is indicative.

Employing the multivariate likelihood ratio test, we obtain a statistic of 16.03, which is in the 0.11 fractile of the chi-square distribution with twenty-four degrees of freedom. Hence, one cannot reject the hypothesis that the sample means and standard deviations as a group of possibly correlated parameters are not significantly different from the group of their adjusted counterparts.

The bootstrap results also support this conclusion. The distance between $(\mu^{sam}, \sigma^{sam})$ and (μ^*, σ^*) is 0.0058. Out of 10,000 resampled sets of 990 observations, 96.3% had a distance larger than this value. Thus, we clearly cannot reject with the bootstrap.

To illustrate graphically the difference between our method and GRS, figure 9 shows the beta–expected return relationship for the twelve industry portfolios, with the sample parameters and with the adjusted parameters. The numbers represent the twelve industry portfolios with the sample parameters,

12 Very similar results are obtained if instead of the December 2008 weights, we calculate the value weighted portfolio for each month and then take the proxy portfolio weights as the value weighted portfolio weights averaged over the entire period.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

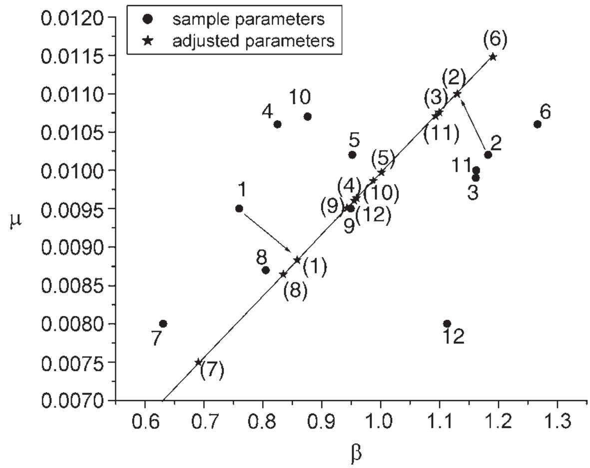

Figure 9
The SML and the twelve industry portfolios with sample and adjusted parameters
The numbers represent the location of the twelve industry portfolios as calculated with the sample parameters, while the numbers in parentheses represent the location of these same twelve portfolios with the adjusted parameters. Note that while the standard tests allow for corrections in only the expected returns, our method also allows for corrections in the standard deviation and therefore betas. Thus, we may have both vertical and horizontal corrections.

and the numbers in parentheses represent the corresponding portfolios with the adjusted parameters. While the average return corrections, as employed in GRS and others, allow for only vertical corrections, our method also allows for corrections in the covariances (hence the betas) and thus for horizontal corrections.[13] Thus, even with a very long sample period (990 months) and a rather low risk-free rate value (0.2%), the CAPM cannot be rejected when estimation errors in the standard deviations are taken into account.[14]

4. Implications for Asset Pricing and Practical Use of the CAPM

The security market line (SML) formula is probably the most widespread method for estimating the cost of capital and for pricing risky assets. Using beta and the SML formula for estimating the expected return, rather than employing the sample average return directly, is usually justified on the basis that

[13] This figure was drawn for a value of $\alpha = 0.98$ in the distance measure (2), in order to emphasize the horizontal corrections. With lower values of α and a larger number of assets, the horizontal corrections are typically smaller (see Section 4), but may still be important.

[14] Very similar results are obtained when the analysis is performed on portfolios formed based on book-to-market ratios instead of the industry portfolios. These results are available from the authors upon request.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

the statistical estimation of beta is more stable than that of the average return. However, when there are questions about how well the SML relationship holds empirically, there are serious doubts about employing betas for pricing.[15] While we cannot prove that the SML relationship holds empirically with the *ex ante* parameters, our analysis does provide another reason for employing betas for estimating the cost of capital.

Suppose that the CAPM holds with the true *ex ante* parameters (μ^*, σ^*) and that the empirically measured parameters are $(\mu^{sam}, \sigma^{sam})$. The true and sample betas of stock i are given respectively by:

$$\beta_i^* = \frac{\sum_{j=1}^{N} x_{mj}\sigma_i^*\sigma_j^*\rho_{ij}}{x_m^{'} C x_m} \tag{3a}$$

$$\beta_i^{sam} = \frac{\sum_{j=1}^{N} x_{mj}\sigma_i^{sam}\sigma_j^{sam}\rho_{ij}}{x_m^{'} C^{sam} x_m}, \tag{3b}$$

where x_m denotes the market portfolio weights. The true cost of equity of firm i is μ_i^*. If one employs the observable β_i^{sam} in the SML formula instead of the correct β_i^*, how accurate will the resulting cost of capital estimate be? In other words, how close are β_i^{sam} and β_i^*? We show that when the number of assets is large, the difference will typically be small. Figure 10 shows β_i^{sam} and β_i^*, where the parameter set (μ^*, σ^*) employed is the solution to Optimization Problem 1. The figure reveals that the difference between β_i^{sam} and β_i^* is very small. The reason is that both the denominators and the numerators of Equations (3a) and (3b) are very similar. The variance of the market proxy is quite close, whether the optimized parameters or the sample parameters are employed (compare the horizontal location of star and the triangle in figures 3 and 4). As for the covariances in the numerator, note that $\sigma_j^* \approx \sigma_j^{sam}$, and in addition, when the number of assets is large, the deviations tend to cancel each other out in the summation, as in some cases $\sigma_j^* > \sigma_j^{sam}$, while in others $\sigma_j^* < \sigma_j^{sam}$ (see column 7 in table 1).[16]

[15] This is, of course, one of the major debates in finance. See, for example, Reinganum (1981), Levy (1981), Lakonishok and Shapiro (1986), Fama and French (1992), and Roll and Ross (1994).

[16] Figure 10 shows the relation between the β_i^{sam}'s and the β_i^*'s when we use a value of $\alpha = 0.75$ in the distance measure D (see Equation (2)). When a higher value of α is employed, the μ_i^*'s are closer to their sample counterparts, and the σ_i^*'s are more distant from their sample counterparts. As a result, the differences between the β_i^{sam}'s and the β_i^*'s also increase. Yet, even with a very high value of $\alpha = 0.97$, with one hundred assets, the β_i^*'s are still very close to the β_i^{sam}'s, with a correlation of 0.96. When the number of assets is smaller, as in the twelve-asset analysis in Section 3, the differences between the β_i^{sam}'s and the β_i^*'s will typically be larger.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

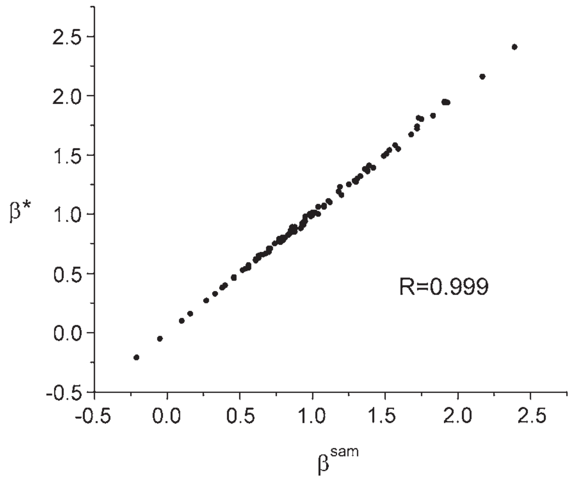

Figure 10
The relation between sample betas and the "true" betas
The "true" parameters are those that solve Optimization Problem 1 and satisfy the CAPM: (μ^*, σ^*). The sample parameters are $(\mu^{sam}, \sigma^{sam})$. The true and sample betas are given by Equations (3a) and (3b). The figure shows that the sample betas are very close to the true betas and thus yield excellent estimates of the expected returns.

Because the market proxy is efficient with the true parameters (μ^*, σ^*), the following relationship holds exactly:

$$\mu_i^* = r_z + \beta_i^*(\mu_m - r_z), \tag{4}$$

where r_z is the expected return on the zero-beta portfolio for index m. Common practice substitutes a "riskless" rate, r_f, for r_z, but this is appropriate only when f and z have the same mean return. Since $\beta_i^{sam} \approx \beta_i^*$, employing the SML with the sample beta, as is commonly done in practice, provides an excellent estimate for the true expected return (assuming $r_f = r_z$):

$$\mu_i^* - \left[r_f + \beta_i^{sam}(\mu_m^{sam} - r_f)\right] = \beta_i^*(\mu_m^* - r_f) - \beta_i^{sam}(\mu_m^{sam} - r_f) \approx 0. \tag{5}$$

The above argument is based on taking the true *ex ante* parameters as the (μ^*, σ^*) vectors solving Optimization Problem 1, that is the parameters ensuring the CAPM that are closest to the sample parameters. What if, instead, we take another set of parameters that ensures the efficiency of the proxy and is consistent with the sample parameters? For example, suppose that we take as the true parameters those that solve Optimization Problem 2 with $\mu_0 = 2\%$ and $\sigma_0 = 4.25\%$ (see point A in figure 4). It turns out that with these parameters,

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

the β_i^*'s and theβ_i^{sam}'s are still very close—see panel A in figure 11. This is also true for other points with very different proxy portfolio expected returns and standard deviations—see panels B, C, and D in figure 11, corresponding to the points B, C, and D in figure 4.

This is a strong result: if the CAPM holds in a way that is consistent with the sample parameters, the differences between sample betas and true betas are going to be small. Thus, if one employs the SML formula for pricing, which implies that the CAPM holds with the *ex ante* parameters, one can be confident about using the sample betas and should not worry about estimation errors in the betas. This conclusion is reached because we are not just looking at the statistical estimation error of a single asset's beta in isolation, as is typically

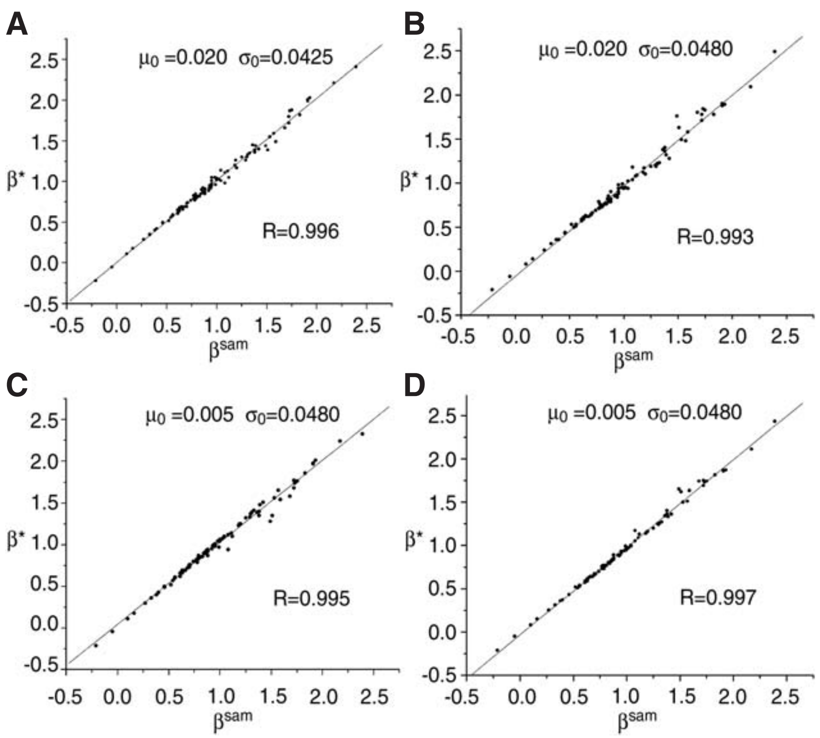

Figure 11
The relation between sample betas and "true" betas for varying values of the market proxy's expected return and volatility
The "true" parameters are those that satisfy the CAPM and solve Optimization Problem 2. Each panel corresponds to a different combination of values of the prespecified expected return and standard deviation of the proxy portfolio, μ_0 and σ_0. (The points corresponding to these four panels are indicated by A, B, C, and D, respectively, in figure 4.) The true and sample betas are given by Equations (3a) and (3b). The figure shows that with one hundred assets, the sample betas are very close to the true betas and thus yield excellent estimates of the true expected returns, even when μ_0 and σ_0 are not close to the values obtained with the sample parameters.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

done, but rather at the error in beta given that the CAPM holds in a way that is consistent with the sample parameters $(\mu^{sam}, \sigma^{sam})$.

From a practical perspective, because sample betas are quite close to betas that have been adjusted to render the market proxy mean/variance efficient,

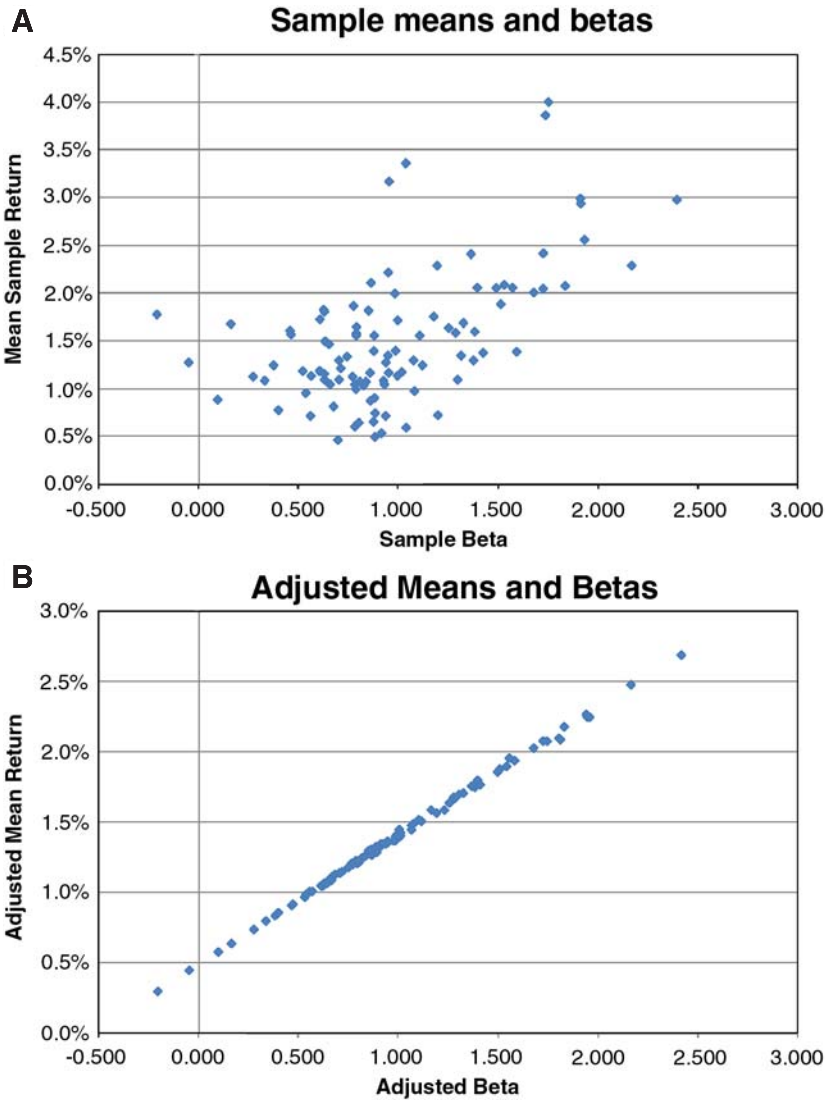

Figure 12
The SML scatter for sample versus adjusted means and betas
Sample estimates of means and betas for our one hundred stocks are plotted against each other in panel A. Panel B plots the corresponding adjusted means and betas that are obtained from Optimization Problem 1.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

improved estimates of expected returns can be obtained from sample betas alone. Sample mean returns should be ignored! To illustrate, figure 12, panel A, shows the cross-sectional relation between sample mean returns and sample betas for our one hundred stocks, while figure 12, panel B, shows the analogous relation for adjusted means and betas. Clearly, the sample means in panel A are not closely related at all to sample betas, but the adjusted means in panel B are perfectly related to adjusted betas.[17]

Consequently, to obtain an improved expected return estimate for any stock, first calculate the adjusted mean return for the market index proxy and for its corresponding zero-beta portfolio.[18] Plugging these numbers along with the sample beta (because it is close to the adjusted beta) into the usual CAPM formula delivers the improved estimate of expected return. Making the market index proxy mean/variance efficient produces useful betas for many practical purposes such as estimation of the cost of equity capital for a firm or of the discount rate for a risky project.

5. Conclusion

Market proxy portfolios are typically very far from the sample efficient frontier. Many studies have tried various adjustments to the sample parameters to make the market proxy mean/variance efficient, without success. Thus, the "common wisdom" is that the empirical return parameters and market portfolio weights are incompatible with the CAPM theory.

In this article, we hope to change that perception. We show that small variations of the sample parameters, well within the range of estimation error, can make a typical market proxy mean/variance efficient. While such parameter variations are reminiscent of "shrinkage," they differ from those obtained with the standard statistical shrinkage methods: They are the result of "reverse optimization." In this reverse optimization, return parameters are derived to make the market proxy mean/variance efficient while being "close" to their sample counterparts.

The fact that we find many such parameter sets and the fact that many previous attempts to vary the return parameters in order to obtain an efficient proxy were unsuccessful seem to indicate that such parameter sets may be very rare in parameter space—it is very unlikely to "stumble onto one of them" by coincidence. Yet, the reverse optimization problem delivers them simply and directly.

These findings suggest that the CAPM (i.e., *ex ante* mean/variance efficiency of the market index proxy) is consistent with the empirically observed return parameters and the market proxy portfolio weights. Of course, this does not constitute a proof of the empirical validity of the model, but it shows that the

[17] The slight deviations from linearity in figure 12, panel B, are caused by rounding error.

[18] For most proxies, the sample means will be close to the adjusted means.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

model cannot be rejected, in contrast to the widespread belief in our profession. The intuitive idea that shrinkage corrections should increase the empirical validity of the CAPM is shown to be valid—with the right corrections, which are small, the index proxy is perfectly efficient. The analysis also shows that in this framework, employing the sample betas typically provides an excellent estimate of the true expected returns.

Appendix

Table A1
The sample parameters and closest parameters ensuring that the market proxy is mean/variance efficient

(1) Stock (i)	(2) μ_i^{sam}	(3) μ_i^*	(4) σ_i^{sam}	(5) σ_i^*	(6) t-Value μ_i^*	(7) $(\sigma_i^*)^2/(\sigma_i^{sam})^2$ (The 95% confidence interval for this value is [0.790–1.319])
1	0.024	0.018	0.165	0.167	−0.423	1.019
2	0.021	0.019	0.115	0.115	−0.170	1.003
3	0.011	0.017	0.106	0.104	0.588	0.963
4	0.029	0.023	0.158	0.160	−0.444	1.028
5	0.039	0.022	0.150	0.156	−1.228	1.077
6	0.005	0.011	0.075	0.073	0.952	0.953
7	0.007	0.013	0.072	0.070	0.938	0.942
8	0.012	0.010	0.051	0.052	−0.433	1.028
9	0.013	0.015	0.070	0.069	0.286	0.978
10	0.016	0.018	0.099	0.098	0.185	0.986
11	0.010	0.013	0.067	0.066	0.344	0.977
12	0.016	0.009	0.092	0.093	−0.819	1.025
13	0.015	0.011	0.071	0.072	−0.627	1.035
14	0.019	0.012	0.100	0.102	−0.702	1.034
15	0.011	0.011	0.061	0.061	−0.029	1.006
16	0.032	0.014	0.159	0.162	−1.215	1.044
17	0.023	0.025	0.158	0.157	0.145	0.990
18	0.024	0.021	0.146	0.147	−0.232	1.016
19	0.011	0.012	0.086	0.085	0.199	0.988
20	0.007	0.010	0.067	0.066	0.477	0.979
21	0.011	0.011	0.065	0.065	0.082	0.996
22	0.018	0.016	0.080	0.081	−0.225	1.018
23	0.012	0.008	0.067	0.068	−0.652	1.023
24	0.013	0.004	0.059	0.059	−1.533	0.995
25	0.017	0.014	0.088	0.088	−0.361	1.021
26	0.014	0.013	0.081	0.082	−0.128	1.007
27	0.006	0.012	0.077	0.075	0.810	0.955
28	0.018	0.011	0.077	0.078	−1.058	1.044
29	0.010	0.012	0.087	0.086	0.276	0.989
30	0.010	0.010	0.065	0.064	0.055	0.999
31	0.012	0.013	0.086	0.085	0.147	0.991
32	0.009	0.006	0.082	0.082	−0.406	1.004
33	0.016	0.009	0.082	0.083	−0.862	1.026
34	0.017	0.006	0.077	0.078	−1.461	1.018
35	0.011	0.012	0.072	0.072	0.243	0.984
36	0.009	0.013	0.064	0.062	0.658	0.954
37	0.012	0.011	0.064	0.064	−0.228	1.012
38	0.026	0.023	0.203	0.204	−0.142	1.006
39	0.011	0.010	0.065	0.065	−0.195	1.009
40	0.006	0.012	0.087	0.085	0.749	0.960

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

Table A1
(continued)

(1) Stock (i)	(2) μ_i^{sam}	(3) μ_i^*	(4) σ_i^{sam}	(5) σ_i^*	(6) t-Value μ_i^*	(7) $(\sigma_i^*)^2/(\sigma_i^{sam})^2$ (The 95% confidence interval for this value is [0.790–1.319])
41	0.010	0.015	0.115	0.114	0.480	0.978
42	0.016	0.017	0.119	0.119	0.011	1.001
43	0.018	0.003	0.100	0.099	−1.615	0.986
44	0.013	0.017	0.105	0.104	0.364	0.976
45	0.009	0.013	0.088	0.087	0.499	0.974
46	0.006	0.014	0.085	0.082	1.067	0.932
47	0.013	0.018	0.124	0.122	0.409	0.978
48	0.011	0.011	0.084	0.083	0.057	0.997
49	0.008	0.009	0.077	0.077	0.112	0.998
50	0.017	0.011	0.082	0.084	−0.884	1.036
51	0.012	0.014	0.081	0.081	0.265	0.984
52	0.021	0.018	0.105	0.106	−0.277	1.019
53	0.016	0.012	0.072	0.073	−0.517	1.030
54	0.011	0.014	0.106	0.105	0.281	0.984
55	0.011	0.012	0.074	0.074	0.118	0.993
56	0.007	0.013	0.076	0.074	0.889	0.945
57	0.011	0.013	0.072	0.071	0.379	0.975
58	0.014	0.019	0.102	0.100	0.581	0.952
59	0.023	0.016	0.089	0.091	−0.807	1.056
60	0.014	0.018	0.090	0.088	0.489	0.961
61	0.012	0.012	0.070	0.070	−0.095	1.005
62	0.012	0.011	0.093	0.093	−0.095	1.000
63	0.008	0.011	0.075	0.074	0.436	0.979
64	0.021	0.019	0.106	0.107	−0.172	1.012
65	0.016	0.013	0.077	0.078	−0.336	1.018
66	0.013	0.014	0.074	0.074	0.110	0.993
67	0.016	0.017	0.076	0.075	0.130	0.988
68	0.011	0.008	0.052	0.052	−0.610	1.020
69	0.020	0.020	0.134	0.133	0.029	0.994
70	0.014	0.014	0.076	0.076	0.009	0.997
71	0.010	0.013	0.094	0.094	0.346	0.983
72	0.015	0.011	0.070	0.071	−0.560	1.028
73	0.018	0.013	0.088	0.089	−0.658	1.033
74	0.022	0.014	0.096	0.098	−0.934	1.049
75	0.011	0.007	0.059	0.059	−0.705	1.018
76	0.005	0.013	0.083	0.081	1.013	0.937
77	0.007	0.013	0.083	0.081	0.718	0.957
78	0.005	0.013	0.083	0.081	1.032	0.938
79	0.013	0.014	0.086	0.086	0.028	0.997
80	0.016	0.015	0.090	0.090	−0.046	1.006
81	0.012	0.015	0.074	0.072	0.392	0.964
82	0.011	0.013	0.070	0.069	0.290	0.983
83	0.021	0.022	0.117	0.116	0.099	0.992
84	0.019	0.019	0.089	0.088	−0.004	0.993
85	0.018	0.011	0.098	0.100	−0.800	1.029
86	0.013	0.012	0.073	0.073	−0.228	1.012
87	0.021	0.021	0.130	0.130	0.031	0.996
88	0.007	0.016	0.095	0.092	0.968	0.939
89	0.021	0.020	0.100	0.100	−0.109	1.009
90	0.040	0.022	0.193	0.199	−1.035	1.052
91	0.034	0.015	0.161	0.164	−1.274	1.046
92	0.030	0.027	0.170	0.171	−0.163	1.014
93	0.012	0.014	0.086	0.086	0.310	0.982
94	0.013	0.011	0.080	0.080	−0.204	1.009

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

Table A1
(continued)

(1) Stock (i)	(2) μ_i^{sam}	(3) μ_i^*	(4) σ_i^{sam}	(5) σ_i^*	(6) t-Value μ_i^*	(7) $(\sigma_i^*)^2 / (\sigma_i^{sam})^2$ (The 95% confidence interval for this value is [0.790–1.319])
95	0.030	0.023	0.130	0.133	−0.579	1.045
96	0.016	0.012	0.147	0.147	−0.245	1.009
97	0.017	0.012	0.087	0.088	−0.523	1.024
98	0.017	0.017	0.102	0.102	0.035	0.997
99	0.020	0.014	0.089	0.090	−0.704	1.041
100	0.021	0.013	0.087	0.089	−0.997	1.057

This is the complete version of table 1 given in the text, where here the data are provided for all one hundred stocks. The sample parameters are given in the second and fourth columns. The expected returns and standard deviations, which are closest to these parameters and ensure that the market proxy is efficient (i.e., the parameters that solve Optimization Problem 1), are given in columns 3 and 5. The t-values for the expected returns are given in column 6, which shows that none of these values are significant at the 95% level. Column 7 reports the ratio between the variances $(\sigma^*)^2$ and the sample variances. The 95% confidence interval for this ratio is [0.790–1.319] (see footnote 5). All of the ratios in the table fall well within this interval.

References

Best, M. J., and R. R. Grauer. 1985. Capital Asset Pricing Compatible with Observed Market Value Weights. *Journal of Finance* 40:85–103.

Black, F., M. C. Jensen, and M. Scholes. 1972. The Capital Asset Pricing Model: Some Empirical Tests. In M.C. Jensen (ed.), *Studies in the Theory of Capital Market*, pp. 78–121. New York: Praeger.

Bonferroni, C. E. 1935. Il Calcolo delle Assicurazioni su Gruppi di Teste. In *Studi in Onore del Professore Salvatore Ortu Carboni*, pp. 13–60. Rome. n.p.

Fama, E. F., and K. R. French. 1992. The Cross-Section of Expected Stock Returns. *Journal of Finance* 47:425–65.

Gibbons, M. R. 1982. Multivariate Tests of Financial Models: A New Approach. *Journal of Financial Economics* 10:3–27.

Gibbons, M. R., S. A. Ross, and J. Shanken. 1989. A Test of the Efficiency of a Given Portfolio. *Econometrica* 57:1121–52.

Green, R. C., and B. Hollifield. 1992. When Will Mean/Variance Efficient Portfolios Be Well Diversified? *Journal of Finance* 47:1785–809.

Jagannathan, R., and T. Ma. 2003. Risk Reduction in Large Portfolios: A Role for Portfolio Weight Constraints. *Journal of Finance* 58:1651–84.

Jobson, J. D., and B. Korkie. 1982. Potential Performance and Tests of Portfolio Efficiency. *Journal of Financial Economics* 10:433–66.

Johnson, N. L., and S. Kotz. 1972. *Distributions in Statistics: Continuous Multivariate Distributions*. New York: Wiley.

Kandel, S., and R. F. Stambaugh. 1987. On Correlations and Inferences about Mean/Variance Efficiency. *Journal of Financial Economics* 18:61–90.

Lakonishok, J., and A. C. Shapiro. 1986. Systematic Risk, Total Risk and Size as Determinants of Stock Market Returns. *Journal of Banking and Finance* 10:115–32.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012

Levy, H. 1981. A Test of the CAPM via a Confidence Level Approach. *Journal of Portfolio Management* 9:56–61.

———. 1983. The Capital Asset Pricing Model: Theory and Empiricism. *Economic Journal* 93:145–65.

Levy, M. 2007. Positive Portfolios Are All Around. Working Paper, Hebrew University.

Lintner, J. 1965. Security Prices, Risk, and the Maximal Gains from Diversification. *Journal of Finance* 20:587–615.

MacKinlay, A. C., and M. P. Richardson. 1991. Using Generalized Method of Moments to Test Mean/Variance Efficiency. *Journal of Finance* 46:511–27.

Miller, R. G. 1991. *Simultaneous Statistical Inference.* New York: Springer.

Reinganum, M. R. 1981. A New Empirical Perspective on the CAPM. *Journal of Financial and Quantitative Analysis* 16:439–62.

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* 4:129–76.

Roll, R., and S. Ross. 1994. On the Cross-sectional Relation Between Expected Returns and Betas. *Journal of Finance* 49:101–21.

Ross, S. 1977. The Capital Asset Pricing Model (CAPM), Short-Sale Restrictions and Related Issues. *Journal of Finance* 32:177–83.

Shanken, J. 1985. Multivariate Tests of the Zero-Beta CAPM. *Journal of Financial Economics* 14:327–48.

Sharpe, W. F. 1964. Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *Journal of Finance* 19:425–42.

———.2007. Expected Utility Asset Allocation. *Financial Analysts Journal* 63:18–30.

Stambaugh, R. F. 1982. On the Exclusion of Assets from Tests of the Two-Parameter Model: A Sensitivity Analysis. *Journal of Financial Economics* 10:237–68.

Zhou, G. 1991. Small Sample Tests of Portfolio Efficiency. *Journal of Financial Economics* 30:165–91.

Downloaded from http://rfs.oxfordjournals.org/ at University of Toronto Library on October 25, 2012