JOIM
www.joim.com

JOURNAL OF INVESTMENT MANAGEMENT, Vol. 10, No. 3, (2012), pp. 9–20

A NEW PERSPECTIVE ON THE VALIDITY OF THE CAPM: STILL ALIVE AND WELL

Moshe Levy[a] *and Richard Roll*[b]

The Capital Asset Pricing Model (CAPM) has far-reaching practical implications for both investors and corporate managers. The model implies that the market portfolio is mean variance efficient, and thus advocates passive investment. It also provides the most widely used measure of risk, beta, which is used to calculate the cost of capital and excess return (alpha). Most academic studies empirically reject the CAPM, leaving the lack of a better alternative as the only uneasy justification for using the model. Here we take a reverse-engineering approach for testing the model and show that with slight variations in the empirically estimated parameters, well within their estimation-error bounds, the CAPM perfectly holds. Thus, in contrast to the widely held belief, the CAPM cannot be empirically rejected.

1 Introduction

The Capital Asset Pricing Model (CAPM) has fundamental implications for the debate about active versus passive investment. If the model holds, the market portfolio is mean variance efficient, implying that stock picking is futile. The model also implies that the risk priced is the investment's beta, and that the cost-of-capital is given by the well-known Security Market Line (SML) equation. While the CAPM is very widely used in practice, most academic studies reject the empirical validity of the model,[1] and the typical view in the financial community is that the model is simply inconsistent with the empirical evidence. This leaves practitioners and academics employing the model at an uneasy situation, where the only justification for using the empirically rejected model is the lack of a better alternative.

The typical approach for testing the CAPM involves empirically estimating the stock return parameters (average returns and covariances), and

[a]Professor of Finance, Jerusalem School of Business, Administration at the Hebrew University, Jerusalem 91905, Israel. Tel.: 972-2588-3219, E-mail: mslm@huji.ac.il
[b]Joel Fried Professor of Finance, Anderson School at UCLA, 110 Westwood Plaza, Los Angeles, CA 90095. Tel.: 310-825-6118, E-mail: rroll@anderson.ucla.edu

examining whether these parameters satisfy the relations implied by the model, i.e., whether the market proxy is mean variance efficient or whether the stocks lie on the SML (note that if the market proxy is mean variance efficient, the SML risk-return linear relation automatically also holds, see Roll (1977)). The findings are typically negative: the sample parameters are inconsistent with the market proxy being mean variance efficient. In fact, the market proxy is typically very far from the sample efficient frontier. Obviously, empirical estimation involves estimation errors. However, the market proxy remains inefficient even when various shrinkage adjustment methods are applied to the sample estimates (see Levy, 1983; Green and Hollifield, 1992; and Jagannathan and Ma, 2003). These findings are the basis for the "common wisdom" about the inconsistency of the model with the empirical evidence.

In this paper we suggest a new approach for empirically testing the CAPM. We take a reverse-engineering 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, well within their estimation error bounds. Hence, minor changes in return parameters reverse previous negative and disappointing finding for the CAPM, and it is shown that the model cannot be empirically rejected.

2 Methodology

Given a market proxy, m, we look for the "minimal" variation of sample parameters that would make this proxy 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 (i.e., located on the Markowitz efficient frontier), and on the other hand are as close as possible to their sample counterparts. For simplicity, when considering the covariance matrix C we allow variation only in the standard deviations, while retaining the same sample correlations:

$$[C] = \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & \cdots & 0 & \sigma_N \end{bmatrix} [\rho^{\text{sam}}] \times \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots \\ & & & 0 \\ 0 & \cdots & 0 & \sigma_N \end{bmatrix} \quad (1)$$

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, $(\mu^{\text{sam}}, \sigma^{\text{sam}})$, we define the following distance measure D between any parameter set (μ, σ) and the sample parameter set:

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

where N is the number of assets and $0 \leq \alpha \leq 1$ 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 of whether a set of parameters (μ,σ) can be considered as "reasonably close" to the sample parameters is 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 (in Levy and Roll (2010) 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 conclusion regarding the compatibility of the parameters (μ,σ) with the sample parameters is independent of the choice of the distance measure 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:

(i) *Optimization Problem*:

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

Subject to:

$$\begin{bmatrix}\sigma_1 & 0 & \cdots & 0\\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots\\ & & & 0\\ 0 & \cdots & 0 & \sigma_N\end{bmatrix}[\rho^{\text{sam}}] \times \begin{bmatrix}\sigma_1 & 0 & \cdots & 0\\ \vdots & \sigma_2 & & \\ & & \ddots & \vdots\\ & & & 0\\ 0 & \cdots & 0 & \sigma_N\end{bmatrix}\begin{bmatrix}x_{m1}\\ x_{m2}\\ \vdots\\ x_{mN}\end{bmatrix} = q\cdot\begin{bmatrix}\mu_1-r_z\\ \mu_2-r_z\\ \vdots\\ \mu_N-r_z\end{bmatrix}, \quad (3)$$

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 (i.e., they ensure that the proxy portfolio is the optimal tangency portfolio from the point of view of the mean variance investor) and are closest to the sample parameters.[2]

Our approach is different from the approach employed in previous studies, such as Black *et al.* (1972) and Gibbons *et al.* (1989), for example, 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 market proxy efficient. For a detailed comparison of our results with the pervious literature, see Levy and Roll (2010).

3 Data and results

Our demonstration sample consists of the 100 largest stocks in the U.S. market (according to December 2006 market capitalizations), which have a complete monthly return records over the period January 1997–December 2006 (120 return observations). Columns (2) and (4) in Table 1 report the sample average returns and standard deviations for 30 of these stocks (the complete information for all 100 stocks is given in Levy and Roll (2010)). 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 100 stocks as of December 2006,

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

The proxy portfolio and the sample mean variance frontier are shown in Figure 1 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 the optimization problem 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.

t-Values for the adjusted expected returns μ^* 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 30 stocks shown in the table, but for the other 70 stocks as well). Column (7) provides the ratio $(\sigma_i^*)^2/(\sigma_i^{\text{sam}})^2$ for each stock. The 95% confidence interval for this ratio is the range 0.790–1.319.[3] The values in Column (7) reveal that for all stocks the ratio $(\sigma_i^*)^2/(\sigma_i^{\text{sam}})^2$ is well within this range (and this is also true for the 70 stocks not shown in the table). Thus, the solution (μ^*, σ^*) to the optimization problem is very close to the sample parameter set in the sense that none of the parameters is significantly different from its sample counterpart.

The t-tests reported above rely on what might be considered a problematic assumption, *viz.* that the estimation errors are independent across parameters. Since 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

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 for μ_i^*	(7) $(\sigma_i^*)^2/(\sigma_i^{\text{sam}})^2$ (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 30 of the 100 stocks (the complete table is given in the appendix of Levy and Roll (2010)). 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 given by Eq. (3)) 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 70 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 3). All of the ratios in the table, as well as the ratios for all other 70 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).

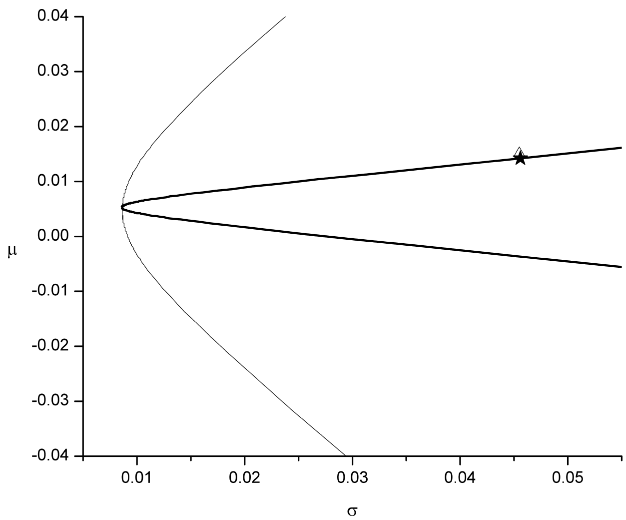

Figure 1 The efficient frontier and market proxy with the sample and the adjusted return parameters.

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 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.

assure that such a possibility did not seriously affect our inference that the proxy 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 and employs the likelihood ratio test. With this test we again find that the hypothesis that the proxy portfolio is mean variance efficient cannot be rejected (for more details, see Levy and Roll (2010)).

The second test does not assume normal return distributions. 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 non-Central Wishart distribution. We therefore conduct a second 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 $(\mu^{\text{sam}}, \sigma^{\text{sam}})$ and (μ^*, σ^*). Thus, we are assuming that the CAPM holds with the parameters (μ^*, σ^*), and we ask, given these parameters, how likely it is to sample parameters such as $(\mu^{\text{sam}}, \sigma^{\text{sam}})$ with T observations. 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. The detailed steps of the bootstrap procedure are provided in the appendix.

The distance between the sample parameter set $(\mu^{\text{sam}}, \sigma^{\text{sam}})$ and the parameter set found with the optimization in Equation (3), (μ^*, σ^*), is 0.06. Out of 10,000 resampled sets of T observations, ALL had a distance larger than this value. Figure 2 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 while 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.

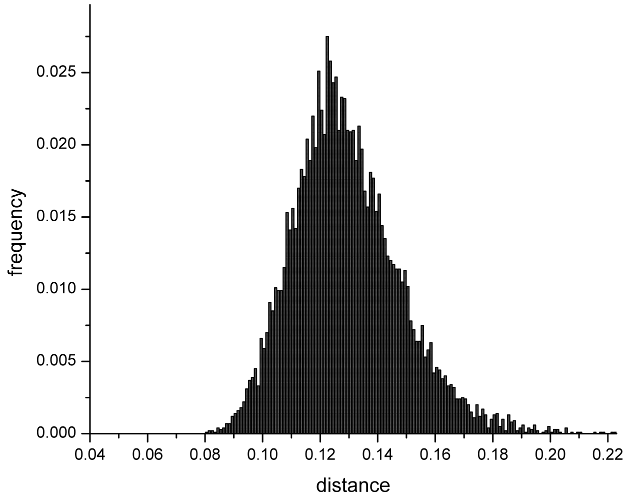

Figure 2 Probability distribution of Euclidean distance between bootstrapped and optimally adjusted parameters.

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.

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 12 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.

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.

4 Discussion 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 1. 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, $\mu_i^* - \mu_i^{\text{sam}}$, as a function of the sample average return, μ_i^{sam}. The figure reveals that high sample returns tend to get negative corrections

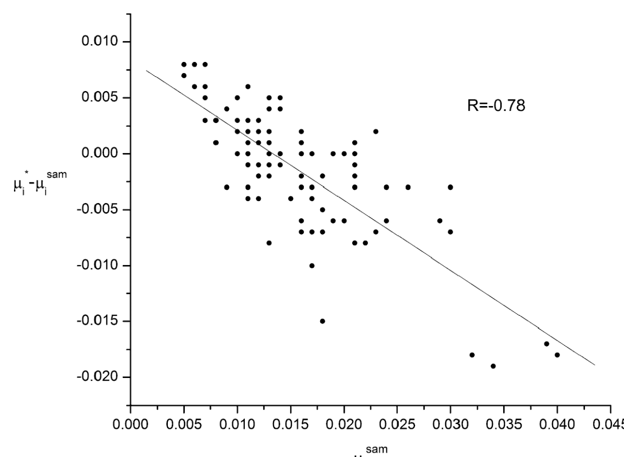

Figure 3 The correction to the expected returns and a function of the sample average return.

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.

($\mu_i^* < \mu_i^{\text{sam}}$), 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 case where all expected returns are identical, the frontier becomes completely flat—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.

There is excellent intuition behind such a result when one recalls two facts: (a) the efficient frontier itself is the result of an optimization; it gives the *minimum* variance for each level of mean return, and (b) 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 1.

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 *in*efficient, 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.

5 Asset pricing implications

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 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.[4] 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 ($\mu^{\text{sam}}, \sigma^{\text{sam}}$). The true and sample betas of stock i are given respectively by:

$$\beta_i^* = \frac{\sum_{j=1}^{N} x_{mj}\sigma_i^*\sigma_j^*\rho_{ij}}{x_m' C x_m} \tag{4a}$$

$$\beta_i^{\text{sam}} = \frac{\sum_{j=1}^{N} x_{mj}\sigma_i^{\text{sam}}\sigma_j^{\text{sam}}\rho_{ij}}{x_m' C^{\text{sam}} x_m}, \tag{4b}$$

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 be estimate? In other

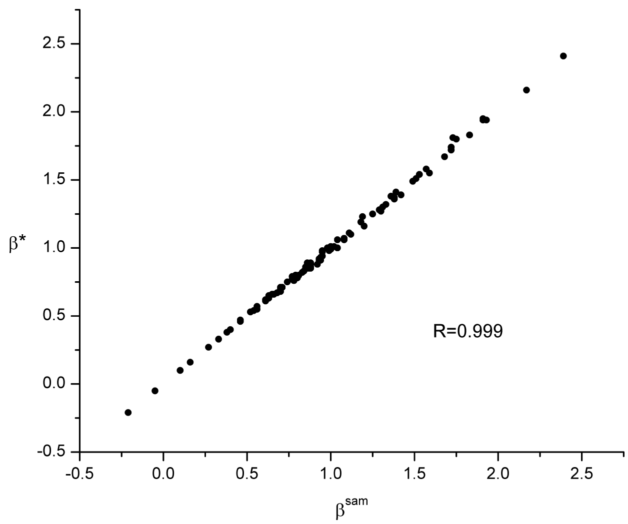

Figure 4 The relation between sample betas and the "true" betas.

The "true" parameters are those that solve the optimization problem (Equation (3)) and satisfy the CAPM: (μ^*, σ^*). The sample parameters are $(\mu^{\text{sam}}, \sigma^{\text{sam}})$. The true and sample betas are given by Equation (4). The figure shows that the sample betas are very close to the true betas, and thus yield excellent estimates of the expected returns.

words, how close are β_i^{sam} and β_i^*? The answer is shown in Figure 4, where the parameter set (μ^*, σ^*) employed is the solution to the optimization problem in Equation (3). 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 (4a) and (4b) 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 Figure 1). As for the covariances in the numerator, note that $\sigma_j^* \approx \sigma_j^{\text{sam}}$, and in addition, the deviations tend to cancel each other out in the summation, as in some cases $\sigma_j^* > \sigma_j^{\text{sam}}$, while in others $\sigma_j^* < \sigma_j^{\text{sam}}$ (see Column 7 in Table 1).[5]

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

$$\mu_i^* = r_z + \beta_i^*(\mu_m - r_z), \tag{5}$$

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 $\beta_i^{\text{sam}} \approx \beta_i^*$, employing the SML with the sample beta, as is commonly done in practice, provides an excellent estimate for the true expected return (assuming $r_f = r_z$)[6]:

$$\begin{aligned}\mu_i^* - \lfloor r_f + \beta_i^{\text{sam}}(\mu_m^{\text{sam}} - r_f)\rfloor \\ = \beta_i^*(\mu_m^* - r_f) - \beta_i^{\text{sam}}(\mu_m^{\text{sam}} - r_f) \approx 0.\end{aligned} \tag{6}$$

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

From a practical perspective, since sample betas are quite close to betas that have been adjusted to render the market proxy mean variance efficient, improved estimates of expected returns can be obtained from sample betas alone. Sample mean returns should be ignored! To illustrate, in Figure 5, Panel A shows the cross-sectional relation between sample mean returns and sample betas for our 100 stocks while in Figure 5, 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.[7]

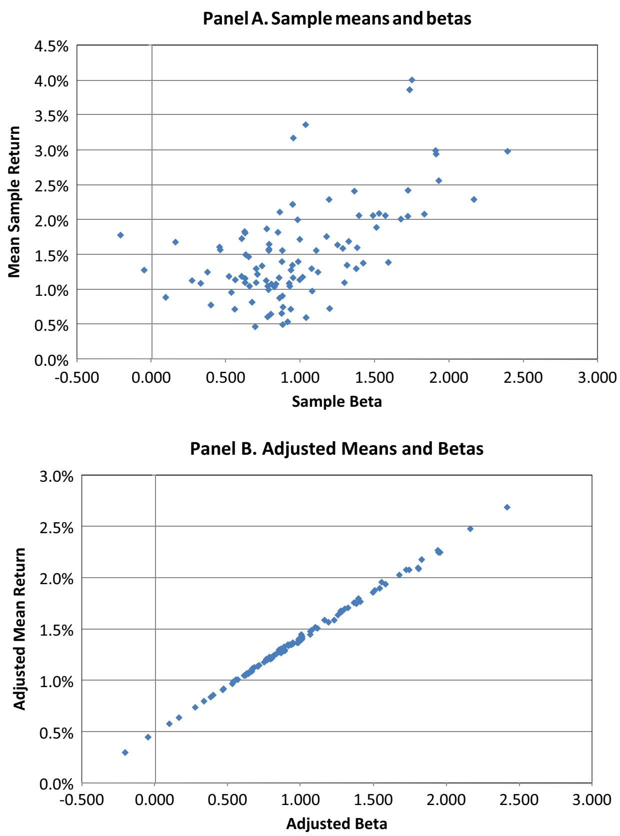

Figure 5 The securities market line scatter for sample versus adjusted means and betas.

Sample estimates of means and betas for our 100 stocks are plotted against each other in Panel A. Panel B plots the corresponding adjusted means and betas that are obtained from the solution to the optimization problem (Equation (3)).

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.[8] Plugging these numbers along with the sample beta (because it's close to the adjusted beta) into the usual CAPM formula delivers the improved estimate of expected return. Making the market index proxy mean variance efficiency 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.

6 Summary

The "common wisdom" in the financial community is that the empirical return parameters and market portfolio weights are incompatible with the CAPM theory. Almost any sample parameter set, as well as any standard shrinkage correction to the parameters, leads to an inefficient market proxy. This is the reason that most studies have concluded that the "CAPM is dead". Yet, our findings show that the corrections obtained by the "reverse optimization" method are both small, and at the same time they ensure the efficiency of the market proxy, and hence the validity of the CAPM risk–return relationship.

These findings suggest that the CAPM 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. We should also be careful to note that it is possible that other asset pricing models cannot be rejected when simultaneous corrections to the means and variances are taken into consideration. Such analysis is beyond the scope of the present paper. We focus on the CAPM, and our findings show that the 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 provides an excellent estimate of the true expected returns.

Appendix

Below are the step-by-step details of the bootstrap procedure:

(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 $b_i = \sigma^*/\sigma^{\text{sam}}$ and $a_i = \mu^* - b_i \mu^{\text{sam}}$. (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, 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:

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

We should note that one could employ various other more sophisticated distance measures (e.g., the distance D in Equation (2)). The results are robust to the distance measure employed. Obviously, we employ the same measure d for the distance between $(\mu^{\text{sam}}, \sigma^{\text{sam}})$ and (μ^*, σ^*) and between (μ^{BS}, σ^{BS}) and (μ^*, σ^*).

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

Notes

1 See, for example, Gibbons (1982), Jobson and Korkie (1982), Levy (1983), Shanken (1985), Kandel and Stambaugh (1987), Gibbons *et al.* (1989), Zhou (1991), and MacKinlay and Richardson (1991).

2 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).

3 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 $(\sigma^{\text{sam}})^2$ in the notation used in this paper), 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 $P(\chi^2_{119} > c_1) = 0.025$, and $P(\chi^2_{119} < c_2) = 0.025$. 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/σ^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 $0.790 < \sigma^2/s^2 < 1.319$.

4 This is, of course, one of the major debates in finance. See, for example, Reinganum (1981), Levy (1982), Lakonishok and Shapiro (1986), Fama and French (1992), and Roll and Ross (1994).

5 Figure 4 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 $\alpha = 0.97$ the β_i^*'s are still very close to the β_i^{sam}'s, with a correlation of 0.96.

6 In the optimization problem Equation (3) r_z is a free parameter. However, if one wishes to employ a specific value, e.g., the current T-bill rate, one can impose this value in the optimization. The expected return of the market proxy is typically estimated by the long-term historical sample average return.

7 The slight deviations from linearity in Figure 5, Panel B are caused by rounding error.

8 For most proxies, the sample means will be close to the adjusted means.

References

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

Black, F., Jensen, M. C., and Scholes, M. (1972). The Capital Asset Pricing Model: Some Empirical Tests, in *Studies in the Theory of Capital Markets*. M.C. Jensen, Ed., New York: Praeger, 79–121.

Fama, E. F. and French, K. R. (1992). "The Cross-Section of Expected Stock Returns," *Journal of Finance* **47**, 425–465.

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

Gibbons, M. R., Ross, S. A., and Shanken, J. (1989). "A Test of the Efficiency of a Given Portfolio," *Econometrica* **57**, 1121–1152.

Green, R. C. and Hollifield, B. (1992). "When Will Mean/Variance Efficient Portfolios be well Diversified?," *Journal of Finance* **47**, 1785–1809.

Jagannathan, R. and Ma, T. (2003). "Risk Reduction in Large Portfolios: A Role for Portfolio Weight Constraints," *Journal of Finance* **58**, 1651–1684.

Jobson, J. D. and Korkie, B. (1982). "Potential Performance and Tests of Portfolio Efficiency," *Journal of Financial Economics* **10**, 433–466.

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

Lakonishok, J. and Shapiro, A. C. (1986). "Systematic Risk, Total Risk and Size as Determinants of Stock Market Returns," *Journal of Banking and Finance* **10**, 115–132.

Levy, H. (1982). "A Test of the CAPM via a Confidence Level Approach," *Journal of Portfolio Management* **9**, 56–61.

Levy, H. (1983). "The Capital Asset Pricing Model: Theory and Empiricism," *Economic Journal* **93**, 145–165.

Levy, M. (2007). "Positive Portfolios are all around," Working Paper, Hebrew University.

Levy, M. and Roll, R. (2010). "The Market Portfolio may be Mean/Variance Efficient after all," *Review of Financial Studies* **23**, 2464–2491.

MacKinlay, A. C. and Richardson, M. P. (1991). "Using Generalized Method of Moments to Test Mean/Variance Efficiency," *Journal of Finance* **46**, 511–527.

Reinganum, M. R. (1981). "A New Empirical Perspective on the CAPM," *Journal of Financial and Quantitative Analysis* **16**, 439–462.

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**, 129–176.

Roll, R. and Ross, S. (1994). "On the Cross Sectional Relation between Expected Returns and Betas," *Journal of Finance* **49**, 101–121.

Shanken, J. (1985). "Multivariate Tests of the Zero-Beta CAPM," *Journal of Financial Economics* **14**, 327–348.

Sharpe, W. F. (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**, 237–268,

Zhou, G. (1991). "Small Sample Tests of Portfolio Efficiency," *Journal of Financial Economics* **30**, 165–191.

Keywords: Mean variance efficiency; CAPM; portfolio optimization; beta