

The Cross-Section of Positively Weighted Portfolios

Daniel Niedermayer* Heinz Zimmermann†

January 11, 2007

Abstract

This paper examines properties of mean-variance inefficient proxies with respect to producing a linear relation between expected returns and betas. The numerical results of a Monte Carlo simulation show that in the CAPM slightly inefficient, positively weighted proxies cause an almost perfect linear expected return - beta relation. Moreover, we show that a strong linearity among a predefined subset of assets exists. These implications are important for the interpretation of empirical tests as well as for asset pricing and for the improvement of proxies' benchmark properties. In contrast to current literature the results suggest that the CAPM's pricing error is small when slightly inefficient, positively weighted proxies are used.

Keywords: asset pricing, CAPM, Roll Critique, mean-variance analysis, short-sale constraint, market proxy

EFM-Classification: 310

*Credit Suisse, Academic Research Program and University of Basel, Department of Finance, Holbeinstrasse 12, 4051 Basel, Switzerland. Email: daniel.niedermayer@vwi.unibe.ch.

†University of Basel, Department of Finance, Holbeinstrasse 12, 4051 Basel, Switzerland. Email: heinz.zimmermann@unibas.ch.

It is well known from mean-variance portfolio analytics that if the market portfolio is efficient, it produces an exact linear relation between assets' expected returns and their betas (Roll (1977)). If it is inefficient, no simple statement about the cross-sectional slope or R^2 can be made. Roll and Ross (1994) demonstrate that even slightly inefficient benchmarks - those being close to the efficient frontier - may produce slopes of an OLS regression of expected returns on betas (and hence, R_{OLS}^2 -coefficients) of exactly zero.

Kandel and Stambaugh (1995) show that by repackaging the original set of assets it is possible to replicate any inefficient proxy's risk and expected return, while affecting the goodness-of-fit of the resulting OLS regression. They show that by such repackaging any R_{OLS}^2 in the open set between zero and one can be reached. However, they also propose a specific efficiency measure for the market portfolio which is directly related to the goodness-of-fit of a GLS regression - specifically: the R_{GLS}^2 value is equal to the squared relative efficiency measure and hence, is uniquely determined by the location (inefficiency) of the market proxy. More specifically, they show that slightly inefficient proxies produce high values of R_{GLS}^2 .

There are two problems with the Kandel and Stambaugh (1995) approach. First, the authors use the inverse of the assets' returns' covariance matrix instead of the residual covariance matrix in their GLS fit. The advantage is that this matrix is readily available, and moreover, it makes the estimated coefficients invariant to repackaging the original assets. However, note that the relationship between expected returns and betas is purely deterministic, so there is no natural choice to take the variance-covariance matrix (of the assets' returns) as weighting matrix in the GLS fit.

The second problem is related to the empirical applicability: Inverting a covariance matrix magnifies estimation errors. This suggests that the GLS

fit is economically and statistically problematic.

In this paper, we return to the OLS method and focus on positively weighted portfolios. We construct two scenarios. In the first scenario using empirical data we find that slightly inefficient positively weighted proxies produce with a high probability low values of R_{OLS}^2 . For the second scenario we construct a theoretical case where the Capital Asset Pricing Model is true, thus, where the market portfolio contains only positive, non-zero weights. In this setting we show that slightly inefficient positively weighted proxies generate high values of R_{OLS}^2 and therefore a strong cross-sectional expected return - beta relation.

As the cross-sectional relation seems to be relevant only in the theoretical setting we finally investigate the cause of the weak cross-sectional relation in the empirical scenario. We show that by a straight forward adjustment of the R_{OLS}^2 measure there is a strong empirical expected return - beta relation when using positively weighted inefficient proxies.

We therefore conclude that for slightly inefficient, positively weighted proxies the cross-sectional expected return - beta relation is not sensitive to the proxy's choice.

The implications of this paper are beyond portfolio theory. In capital market equilibrium the market portfolio must be strictly positively weighted. Our finding that positively weighted proxies close to the market portfolio reveal a close cross-sectional expected return-beta relationship¹ has strong implications for interpreting empirical CAPM tests and the results from cross-sectional regressions, such as Fama and French (1992).

The rest of the paper is organized as follows: The first section describes the numerical algorithm used in the Monte Carlo simulation and the results

¹i.e. assets lie close to the security market line if the market proxy is slightly inefficient.

of two tests. The first test shows the resulting R_{OLS}^2 distributions associated with inefficient, positively weighted proxies for monthly returns of the Dow Jones Industrial Average between 2001 and 2006. The second test is theoretical and constructs a case similar to Grauer (1999) where the CAPM is true - thus, where the market portfolio contains only positive, non-zero weights. The second section extends the R_{OLS}^2 measure in a way that allows for a better extraction of information about expected return - beta linearity when proxy weights are non-negative. The final section emphasizes the major implication for portfolio theory and the CAPM.

1 Portfolio Inefficiency and the Cross-Section of Expected Returns

If a portfolio p is mean-variance efficient, it will produce an exact linear relation between assets' expected returns (μ_i) and the corresponding betas ($\beta_{p,i}$) (see Roll (1977)). If a portfolio is inefficient, Roll and Ross (1994) find that the slope of a cross-sectional expected return - beta regression is very sensitive to the proxy's choice. Moreover, the analytics in Roll and Ross (1994) reveals that the cross-sectional slope is not uniquely determined by proxies with identical expected returns and variance. Roll and Ross (1994) therefore conclude that "[t]he nonexclusivity of our sets makes it impossible to determine the cross-sectional mean-beta relation simply by plotting the position of the proxy in the mean-variance plane. We wish this were possible. It is not." We show in this paper that more optimistic results are achieved if restrictions against short selling are imposed. This paper shows that when including restrictions against short sales proxies close to the efficient frontier will generally produce a higher cross-sectional expected return - beta relation

than proxies far away from the efficient frontier.

Kandel and Stambaugh (1995) provide an analytical solution to the problem raised by Roll and Ross (1994). They show that if the proxy portfolio is efficient, repackaging the original set of assets² has no effect on $R_{OLS}^2 (= 1)$. If a proxy is inefficient, repackaging the original set of assets and replicating a proxy's expected return and standard deviation, the R_{OLS}^2 measure can reach any value in the interval $(0, 1)$.³

As a consequence of these results the expected return - beta linearity can break down entirely when a proxy is only slightly inefficient. Therefore – supporting the critique of Roll (1977) – imposing linearity in empirical tests can hardly be justified.

Kandel and Stambaugh (1995) propose a GLS procedure that overcomes the problem of the non-uniqueness of the OLS measure. They are able to show that the value of R_{GLS}^2 is uniquely determined by the proxy's location in the expected return - standard deviation plane. Proxies close to the minimum variance frontier produce an R_{GLS}^2 close to 1. Furthermore, using the covariance matrix Σ in the GLS regression makes the coefficients invariant with respect to repackaging.⁴

However, the major difficulty regarding the inefficiency measure based on GLS is the choice of the residual covariance matrix. It should be noted that the assets' positions in the (β_p, μ) plane are deterministic and therefore no statistical estimation but rather a geometric fit takes place. The residuals ε

² Formally, this is done by multiplying the assets' random return vector with an $(n \times n)$ matrix (where n is the number of assets).

³In fact, R_{OLS}^2 can even be zero if the proxy is sufficiently inefficient, i.e. if it lies inside the boundaries shown by Roll and Ross (1994). If the OLS fit's slope (such as analyzed by Roll and Ross (1994)) is zero, the corresponding R_{OLS}^2 is zero as well. More interestingly, except for the global minimum variance portfolio the reverse relation is also true.

⁴However, as shown by Cochrane (2001, p.295), there exist other weighting matrices with this property.

are therefore not stochastic but represent the assets' deterministic deviations from the estimated linear fit. In fact, the use of Σ instead of $E(\epsilon\epsilon')$ as the artificial covariance matrix in the GLS fit is justified by the analytical tractability of the resulting equations.⁵

However, the GLS method is not only problematic from a methodological point of view. Since the theoretical value of the covariance matrix is unknown, the described procedure will cause computational problems when applying it to empirical data. When inverting a sample covariance matrix estimation errors are magnified.

In order to circumvent the difficulties related to the GLS methodology we will return to the simple cross-sectional OLS framework, but examine positively weighted proxies only. As shown in the following this approach combines the advantages of the OLS and GLS methods of Kandel and Stambaugh (1995).

2 The Goodness-of-fit of Positively Weighted Proxies

Market proxies (mostly stock indices) used in empirical tests are positively weighted by construction. This is also true when market proxies are used to test market equilibrium where the market portfolio is a proxy for the aggregate wealth with the components being in net positive supply. Therefore, we analyze the cross-sectional properties of positively weighted proxies⁶ and, in order to avoid the difficulties related to the GLS method, we return to the simple OLS framework. We introduce the following notation:

⁵Compare equation (22) in Kandel and Stambaugh (1995, p.167).

⁶This is in line with a critique pointed out by Grauer (1999) whereas the studies of Roll and Ross (1994) and Kandel and Stambaugh (1995) do not explicitly investigate positively weighted proxies.

n : number of assets,

$\boldsymbol{\mu}$: an n vector containing all assets' expected returns,

$\boldsymbol{\Sigma}$: the positive definite ($n \times n$) covariance matrix,

\boldsymbol{w}_q : an n vector describing the weights in a portfolio q ,

$\mathbf{1}$: an n vector containing ones, $[1 \ 1 \ \dots \ 1]'$,

μ_q : the expected return of portfolio q ,

σ_q : a portfolio q 's standard deviation ($= \sqrt{\boldsymbol{w}'_q \boldsymbol{\Sigma} \boldsymbol{w}_q}$).

In unconstrained optimization the weight vector of a minimum variance portfolio is uniquely determined by its expected return level. However, for an inefficient portfolio q with (σ_q, μ_q) , the number of weight vectors satisfying

$$\boldsymbol{w}'_q \mathbf{1} = 1 \tag{1}$$

$$\boldsymbol{w}'_q \boldsymbol{\mu} = \mu_q \tag{2}$$

$$\boldsymbol{w}'_q \boldsymbol{\Sigma} \boldsymbol{w}_q = \sigma_q^2 \tag{3}$$

can be infinite as long as the number of assets exceeds 3 (the number of conditions). Kandel and Stambaugh (1995) show that some of the solutions \boldsymbol{w}_q will yield high and some will yield low values of R_{OLS}^2 . As argued before, we require non-negative solutions of \boldsymbol{w}_q . Therefore, we study whether the Kandel and Stambaugh (1995) results can be extended to the case without short sales by investigating positively weighted proxies' R_{OLS}^2 conditional on proxies' position in the (σ, μ) plane. We apply a Monte Carlo simulation and examine the distributions of R_{OLS}^2 for given values of μ_q and σ_q .

For tackling this problem it is useful to recapitulate some basic properties of portfolios. For this purpose five different cases of portfolio locations in the (σ, μ) plane are distinguished; see figure 1.

Consider a universe of n (> 3) assets. Eq. (1) together with (2) form an $n - 2$ dimensional hyperplane. Geometrically, the set of solutions \mathbb{S} of (1), (2) and (3) is given by the intersection of the $n - 2$ dimensional hyperplane of the constraints (1) and (2) and the ellipsoid given by (3). Define further the feasible region (\mathbb{F}) where all weights are non-negative

$$\mathbf{w} \geq 0. \quad (4)$$

Consider the sets \mathbb{S} and $\mathbb{S}_{\mathbb{F}} = \mathbb{S} \cap \mathbb{F}$. Moving from point t in figure 1 towards the right, one can distinguish five cases:

t: $\mathbb{S} = \{\}$: the $n - 2$ dimensional hyperplane and the ellipsoid do not have common points. No solution exists for this (σ, μ) location.

R: \mathbb{S} consists of a single point: the $n - 2$ dimensional hyperplane is tangential to the ellipsoid. This is the (unconstrained) minimum variance portfolio. For the case shown in figure 1 this point lies, however, outside of the feasible region \mathbb{F} , $\mathbb{S}_{\mathbb{F}} = \{\}$.

r: Increasing σ_q further, \mathbb{S} will become an extended region, which still lies completely outside \mathbb{F} . Here, all portfolios will contain at least one negatively weighted asset.

Q: $\mathbb{S}_{\mathbb{F}}$ consists of a single point: the constrained minimum variance portfolio. In this portfolio all weights are non-negative.

q: $\mathbb{S}_{\mathbb{F}}$ becomes an extended region: there are infinite many solutions to

equations (1), (2) and (3), satisfying the inequality (4).⁷

In contrast to Kandel and Stambaugh (1995) who do not differentiate between locations r and q (respectively, Q), the focus of our interest are locations Q and q .

Knowing that for each location in the (σ, μ) plane such as q in figure 1 there exists an infinite number of feasible weight vectors, it is useful to analyze a large number of feasible \mathbf{w} vectors and to examine the resulting distribution of R_{OLS}^2 values. More specifically, we will pick out solutions with uniform probability density from the set $\mathbb{S}_{\mathbb{F}}$. The yielding histogram reveals how the location of positively weighted inefficient proxies affects the distribution of R_{OLS}^2 . The method for determining the distribution of an R_{OLS}^2 belonging to a specific (σ, μ) location is shown in the next section.

3 The Algorithm

The aim of this algorithm is to pick out feasible (non-negative) weight vectors uniformly from $\mathbb{S}_{\mathbb{F}}(\sigma_q, \mu_q)$, where Q is the constrained minimum variance portfolio and q has the same expected return as Q ($\mathbf{w}'_q \boldsymbol{\mu} = \mu_q = \mu_Q$).

Due to the positivity constraint the set $\mathbb{S}_{\mathbb{F}}(\sigma_q, \mu_q)$ is in general not trivial. Therefore, it is practically impossible to directly generate independent, uniformly distributed weight vectors in this set. We achieve this goal by a random walk in $\mathbb{S}_{\mathbb{F}}(\sigma_q, \mu_q)$.⁸ A solution to this problem is shown in appendix

⁷Although there are two more cases when moving rightwards in figure 1 we ignore them in our analysis as they are irrelevant for our purposes.

⁸A problem encountered in this procedure is that if we move beyond some value of σ_q one meets, however, a new problem – the set $\mathbb{S}_{\mathbb{F}}(\sigma_q, \mu_q)$ becomes disconnected. Starting from a given point of $\mathbb{S}_{\mathbb{F}}$ the random walk will remain in the corresponding subset of $\mathbb{S}_{\mathbb{F}}$, i.e. the procedure is not ergodic. To overcome this problem we let the value of σ_q change during the random walk in some interval $\sigma_q \in [\sigma_1, \sigma_2]$ so that at the lower end of the interval $\mathbb{S}_{\mathbb{F}}$ becomes a connected region. For the interval's upper part this specification of

A.

The random walk is implemented by a Metropolis algorithm. It has to satisfy two basic properties: the probability for a new trial weight w' from the actual weight w has to be symmetric, $p_t(w \rightarrow w') = p_t(w' \rightarrow w)$, and p_t has to be ergodic – any state w can be reached from any other state in some number of steps. These properties ensure that the resulting distribution will be uniform in $\mathbb{S}_{\mathbb{F}}(\sigma, \mu)$.

The Results

We run the algorithm for two scenarios. The *empirical scenario* is based on data from the Dow Jones Industrial Average. In the second scenario we change the expected returns used in the first scenario in a way such that the *unconstrained* minimum variance portfolio is strictly positively weighted. We call this scenario *the CAPM scenario*.

Empirical Scenario. Monthly returns from 30 stocks from the Dow Jones Industrial Average between January 2001 and February 2006 are used as input. The corresponding constrained and unconstrained minimum variance frontiers are plotted in figure 2a.

When running the simulation at $\mu_Q = 0.015$ and iterating 5 million times, the random walk moves horizontally on a distance illustrated by the dotted lines in figure 2a. The resulting histogram drawn from a simulation of 5 million weight vectors is shown in figure 3a.

Note that in figure 2a the unique weight vector describing the constrained minimum variance portfolio Q yields an R_{OLS}^2 of 0.1279. Moving rightwards

the algorithm connects all previously non-ergodic regions of $\mathbb{S}_{\mathbb{F}}(\sigma, \mu)$. The exact procedure is described in appendix A.

$\mathbb{S}_{\mathbb{F}}$ becomes an extended region where an infinite number of proxies have the same expected return and variance. However, in contrast to the findings of Kandel and Stambaugh (1995), none of these proxies produces values of R_{OLS}^2 close to 0 or 1. As shown in figure 3a the values of R_{OLS}^2 remain within a small range when being close to the constrained minimum variance portfolio Q . These results suggest that extreme values of R_{OLS}^2 are either impossible or very unlikely.

Another feature apparent in figure 3a is that for $\mu_Q = 0.015$ all distributions lie in a range between 0.02 and 0.13. Notice that for this sample the range of possible R_{OLS}^2 values cannot be strongly raised by an appropriate choice of μ_Q . As shown in the following section, the range of R_{OLS}^2 values that belong to an expected return level μ_Q depends on the number of non-zero weighted assets, k , of the respective constrained minimum variance portfolio. Because k cannot exceed 10 in the Dow Jones sample⁹ containing 30 assets, no expected return levels exist where high values of R_{OLS}^2 are likely to occur.

As the range in which the R_{OLS}^2 histograms lie strongly depends on k , we shall examine a scenario with a high value of k , more specifically, where $k = n$.

*The CAPM scenario.*¹⁰ In a world where the CAPM is true there must be at least one strictly positively weighted market portfolio M . Market equilibrium means for this case that k equals n . This property influences the resulting R_{OLS}^2 distributions for positively weighted proxies with the

⁹This result can be explicitly shown by algorithms that compute all turning points of the efficient frontier (compare Markowitz, Todd, and Sharpe (2000) and Niedermayer and Niedermayer (2006)).

¹⁰In this scenario we consider the CAPM as being true.

same expected return as M .

For constructing this scenario we use the covariance matrix from the Dow Jones sample and define a vector of non-negative weight $\mathbf{w}_M \geq \mathbf{0}$ describing the ‘market portfolio’ M with $\mathbf{w}'_M \boldsymbol{\mu} = \mu_Q = 0.015$ and chose arbitrarily $\mu_O = 0.005$ as the orthogonal portfolio’s expected return. When shifting all assets’ expected returns (μ_i) to the security market line through $(0, \mu_O)$ and $(1, \mathbf{w}'_M \boldsymbol{\mu})$ in the (β_M, μ) plane, k will equal n for the market portfolio.¹¹ In such a setting we run the simulation with $\mathbf{w}'_M \boldsymbol{\mu} = \mu_Q = 0.015$ and 5 million iterations; the results are illustrated in figure 3b.

Not surprisingly, the value of R_{OLS}^2 is 1 when looking at the market portfolio. When moving rightwards, R_{OLS}^2 falls gradually. Most importantly, for proxies close to the market portfolio, R_{OLS}^2 remains high.

Before discussing the implications of this result it will be shown in the following section that even for the empirical scenario with a low k/n ratio, there is evidence for strong linearity among a subset of assets when inefficient positively weighted proxies are used.

4 Adjusting the Goodness-of fit Measure

An explanation for the low values of R_{OLS}^2 in figure 3a can be found when looking at figures 4a and 4b. Although portfolio q_1 is located on the CMVF but not on the unconstrained frontier, R_{OLS}^2 will not be 1 even if there is an exact linear relation among the non-zero weighted assets in the (β_{q_1}, μ) plane.¹² The black dots in figure 4b show those assets’ (β_{q_1}, μ) locations whose weights in q_1 are not zero and the zero weighted assets in q_1 are marked

¹¹This procedure is similar to the one in Grauer (1999) which specifies a riskless asset.

¹²Assets with non-negative weights in a portfolio are said to be *in* the portfolio, otherwise *out* of the portfolio (see e.g. Sharpe (1991) or Markowitz, Todd, and Sharpe (2000)).

with a triangle in the (β_{q_1}, μ) plane. An OLS fit that weights each asset equally (shown by the dotted line in figure 4a) will produce a lower R_{OLS}^2 (here: 0.756) than a fit that only takes account of the non-zero weighted assets. In fact, the latter type of fit must produce an exact linear relation among the non-zero weighted assets in q_1 since q_1 is on the CMVF.

In order to understand the previous relationship, notice that the constrained optimization problem for n assets is equivalent to the unconstrained optimization for the reduced system where all assets with $\omega_{q,i} = 0$ are dropped. The quadratic optimization problem can be modified into a simple optimization problem without short selling restrictions by dropping all assets i with $\omega_{q,i} = 0$ from the original set of n assets. Crossing out the respective rows and columns from the covariance matrix and vector of expected returns, the resulting (unconstrained) optimal weight vector ω_q^* at an expected return level μ_q will lead to the same results as in the constrained case.¹³ Then the k positively weighted assets will be efficient in the modified set and the corresponding R_{OLS}^2 must be 1. Therefore, those k assets must be located on a straight line in the (β_{q_1}, μ) plane such as shown in figure 4b. Note further, that the $n - k$ unused assets – marked with a triangle – must lie below this line. This comes directly from their property of having a lower marginal utility than the assets *in* the portfolio.¹⁴

Moreover, when moving rightwards from q_1 to q_2 in figure 4a, we hope to obtain solutions where the non-zero weighted assets (black dots) do not

¹³ ω_q is marked with an asterisk since it will be of lower dimension than the original weight vector. Notice that $\omega_q' \Sigma$ and $\Sigma \omega_q$ correspond to the crossing-out of the covariance matrix' respective rows and columns and $\omega_q' \mu$ corresponds to the elimination of the unused assets in μ .

¹⁴For a detailed treatment of this topic see Sharpe (1991, p. 501f).

deviate too much from the fit as long as q_2 is close enough to the CMVF. Such a situation is plotted in figure 4c. Even if an OLS fit (dotted line) yields a relatively low R_{OLS}^2 (here: 0.708) the R^2 measure of a fit over the non-zero weighted assets will be still high (in this case 0.997).

In the following we correct for this distortion of R_{OLS}^2 caused by assets that are outside of the portfolio. The standard R_{OLS}^2 is derived from a linear regression model

$$\boldsymbol{\mu} = \gamma_1 \mathbf{1} + \gamma_2 \boldsymbol{\beta}_q + \boldsymbol{\varepsilon}, \quad (5)$$

by minimizing $\boldsymbol{\varepsilon}'\boldsymbol{\varepsilon}$, where γ_1, γ_2 are the intercept and slope coefficients of the linear fit, $\boldsymbol{\beta}_q$ is equal to $(\boldsymbol{\Sigma}\mathbf{w}_q)/(\mathbf{w}_q'\boldsymbol{\Sigma}\mathbf{w}_q)$ and $\boldsymbol{\varepsilon}$ is an n vector containing the deterministic deviations from the fitted line.

We introduce an $(n \times n)$ weighting matrix $\boldsymbol{\Omega}$ and minimize with respect to γ_1, γ_2 the expression

$$\boldsymbol{\varepsilon}'\boldsymbol{\Omega}\boldsymbol{\varepsilon} = (\mathbf{X}\boldsymbol{\gamma} - \boldsymbol{\mu})'\boldsymbol{\Omega}(\mathbf{X}\boldsymbol{\gamma} - \boldsymbol{\mu}) \quad \text{where } \boldsymbol{\gamma} = \begin{pmatrix} \gamma_1 \\ \gamma_2 \end{pmatrix} \quad \text{and } \mathbf{X} = [\mathbf{1} \quad \boldsymbol{\beta}_q].$$

The solution for γ_1 and γ_2 is

$$\boldsymbol{\gamma} = (\mathbf{X}'\boldsymbol{\Omega}\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\Omega}\boldsymbol{\mu}.$$

When setting $\boldsymbol{\Omega} = \mathbf{I}$ the fit corresponds to an OLS fit and when setting $\boldsymbol{\Omega} = \boldsymbol{\Sigma}^{-1}$ the fit describes exactly a GLS fit. As shown by Kandel and Stambaugh (1995) the latter case leads to an R_{PWA}^2 (here: equal to R_{GLS}^2) distribution that is focussed on exactly one value; the distribution's standard deviation is zero.

In order to focus the analysis only on the positive, non-zero weighted assets we specify $\boldsymbol{\Omega}$ in a way that ignores those assets of a portfolio q whose

weights are zero in the corresponding constrained minimum variance portfolio Q of the same expected return level μ_q . Therefore the weighting matrix $\mathbf{\Omega}$ contains zeros as off-diagonal elements and as diagonal elements a 1 if the corresponding asset was in portfolio Q and 0 otherwise.

Let our corrected goodness-of-fit value, R_{PWA}^2 (*portfolio weight adjusted R^2*) be defined as

$$R_{PWA}^2 \equiv 1 - \frac{\boldsymbol{\varepsilon}'\mathbf{\Omega}\boldsymbol{\varepsilon}}{\boldsymbol{\delta}'\mathbf{\Omega}\boldsymbol{\delta}}$$

where $\boldsymbol{\delta}$ corresponds to the residuals from a linear fit by a horizontal line.

Notice that the specification of $\mathbf{\Omega}$ depends on a proxy's expected return level. When examining proxy locations between two neighboring turning points¹⁵ the same weighting matrix $\mathbf{\Omega}$ is applied. Therefore, it is straightforward to introduce the term *turning point zone*.

A turning point zone describes a region where an R_{PWA}^2 measure examines the same subset of assets. Let t_1 and t_2 be two neighboring turning points on the CMVF. As all portfolios on the CMVF between t_1 and t_2 have zero weights at the same positions, the weighting matrix $\mathbf{\Omega}$ is the same for all portfolios with an expected return level in the open set of (μ_{t_1}, μ_{t_2}) .¹⁶ Thus, in order to interpret R_{PWA}^2 of an inefficient portfolio q , the turning point zone that belongs to q must be found. Then the subset of the examined assets is determined.

The pleasant feature of achieving goodness-of-fit measures of 1 when the

¹⁵Defined as constrained minimum variance portfolios in whose vicinity other constrained minimum variance portfolios contain a different number of non-zero weighted assets.

¹⁶Notice that a portfolio q with $\mu_q = \mu_{t_1}$ (or $\mu_q = \mu_{t_2}$) does not belong to the turning zone (t_1, t_2) but to the turning point 'zone' (t_1) (or (t_2)). Of course this is not a real zone but an expected return level. Since the case where a portfolio q is exactly on the expected return level of a turning point is the special case, we consider the term turning point zone still as adequate.

examined location is on the CMVF cannot be separated from two disadvantages. First, the subset of assets changes when looking at proxies in different turning point zones. In this case a direct comparison of R_{PWA}^2 should not be done. Second, inefficient portfolios inside certain turning point zones will describe the (μ, β_q) relation based on very few assets. In extreme cases a linear fit over two assets takes place and the value of R_{PWA}^2 is 1.¹⁷ Despite these inconveniences this way of defining R_{PWA}^2 makes it possible to use a simple OLS framework with results easy to interpret.

Applying this adjustment to the Dow Jones sample, we run a Monte Carlo simulation with 5 million iterations. Figure 5 shows the results for the expected return, $\mu_Q = 0.015$. Notice that the corresponding turning point zone describes a subset of the assets 6, 16, 21, 25 and 30 (CAT, JNJ, MO, PG and XOM). Figure 5 reveals that the R_{PWA}^2 that describes the relation among these assets is very high compared to the unadjusted R^2 . In fact, all distributions within 0.034 and 0.038 have a median of at least 0.95.

5 Conclusion

This paper numerically analyzes the relation between expected returns and betas when the proxy portfolio is restricted to positive weights. Two cases are examined: The empirical case is based on monthly returns of Dow Jones Industrial Average stocks from 2001 to 2006; the second, theoretical case is based on the assumption that the CAPM holds, i.e. the market portfolio is efficient and contains only strictly positively weighted assets.

¹⁷Notice, that in general it is not possible to obtain a subset belonging to a portfolio q 's location that contains only one asset. Such locations would be at the expected return level of the CMVF's highest or lowest point where the constrained minimum variance portfolio consists only of the respective asset. However, such a positively weighted portfolio q does not exist since it is not possible to find feasible solutions rightward from the CMVF's highest or lowest points.

Our key findings differ from those of Kandel and Stambaugh (1995). We show that slightly inefficient positively weighted proxies produce similar cross-sectional R_{OLS}^2 values as if closely located constrained minimum variance portfolios are used. As a consequence of this observation, in a scenario where the CAPM is true, positively weighted proxies close to the market portfolio will produce an almost linear relation between expected returns and betas. Most importantly, an almost linear relation is a pure consequence of mean variance mathematics and does not rely on an assumption about capital market equilibrium.

These results might help to interpret the findings of Fama and French (1992) who show that the cross-sectional average return - beta relationship has disappeared in the period between 1963 and 1990.¹⁸ According to our results one of the following three cases could explain this observation:

1. The market proxy used by Fama and French (1992) is located too far away from the CMVF.
2. The market proxy is located in a turning point zone (expected return range) where only few assets are strictly positively weighted in the corresponding constrained minimum variance portfolio.
3. ‘Bad luck’: One of the unlikely observations at the frequency distribution’s lower tail was realized.

These results show that if, in general, constrained minimum variance portfolios contain many non-zero weights - respectively, if the CAPM is (almost) true - the choice of the benchmark proxy is an absolutely crucial task

¹⁸Fama and French (1992) investigated the data of nonfinancial firms obtained from the Center for Research in Security Prices (CRSP).

for detecting a high cross-sectional expected return - beta relation. However, if all constrained minimum variance portfolios have a high number of zero weighted assets, none of the positively weighted proxies produces a high cross-sectional R_{OLS}^2 value. Moreover, it was shown that it can be misleading to search for a linear relationship among all assets. For example, in our Dow Jones sample, no strictly positively weighted minimum variance portfolio exists on the unconstrained frontier. Therefore, high R_{OLS}^2 values for positively weighted, inefficient proxies cannot be found. But mean-variance portfolio theory states more; portfolios on the constrained frontier produce an exact linear relation among those assets included in the portfolio. When assessing the implications of positively weighted proxies for the distributions of R_{OLS}^2 values, the right benchmark should be located on the constrained frontier. This is intuitively clear because unconstrained minimum variance portfolios often contain economically unrealistic negative weights, and therefore, the constrained and unconstrained frontiers differ.

We have demonstrated that for a given universe of assets, positively weighted proxies provide a stronger linear expected return - beta relationship among specific subsets of assets than among the entire set. Comparing an asset's performance with a benchmark and relying on the expected return - beta linearity is especially useful when the examined asset is *in* the constrained minimum variance portfolio that has the proxy's expected return. This type of adjustment substantially improves the benchmark properties of market proxies. Knowing that the set of analyzed portfolios consists of positively weighted proxies, and after adjusting the linearity measure as proposed in this paper, we find that the inefficiency of a proxy does not break down the expected return - beta relationship entirely. Thus, in a constrained portfolio setting the statement that 'if the market return is even slightly in-

efficient, covariances with the market return can generate arbitrary pricing errors' (Cochrane and Saa-Requejo (2000, p.83)) is not supported by our findings.

The scenario in figure 2b which represents the case where the CAPM is true shows that slightly inefficient proxies must still produce an almost linear relationship between expected returns and betas. Therefore, if there is a positively weighted efficient market portfolio such as stated by the CAPM, the notion of expected return - beta linearity is still valid even if slightly inefficient benchmark proxies are used.

A The Monte Carlo Algorithm

A Monte Carlo algorithm that covers an n -dimensional connected surface can be quite easily formulated. When searching in the vicinity of an initial solution by choosing the step size appropriately a sufficiently high acceptance rate¹⁹ for finding a new solution can be achieved. However, if the surface is disconnected, it becomes virtually impossible to find a solution on the disconnected region (especially with problems with high dimensionality). In such a case the algorithm is *not ergodic* and the stationary distribution is not unique (it depends on the starting point). In the following using an example with 4 assets we show that the problem of disconnected surfaces arises when σ becomes too large and we demonstrate a way of solving this difficulty.

Consider a universe of 4 assets. Solving the two linear constraints, $\mathbf{w}'\mathbf{1} = 1$ and $\boldsymbol{\mu}'\mathbf{w} = \mu_q$ for w_1 and w_2 and plugging $w_3 = 0$ and $w_4 = 0$ into the solution yields two lines in the (w_1, w_2) plane. As shown in figure 6, the

¹⁹However, a too small step size implying an acceptance rate close to 1 would be inappropriate because it would require too many steps to cover the whole feasible region.

four lines, $w_1 = 0$, $w_2 = 0$, $w_3 = 0$ and $w_4 = 0$ border the feasible region, i.e. weights in the shaded area belong to positively weighted portfolios with expected return level μ_q .

Moreover, those values of (w_1, w_2) that produce a certain σ and satisfy the two linear constraints plot an ellipse in the (w_1, w_2) plane. The part of the ellipse that lies within the feasible region represents solutions of $\mathbf{w}'\Sigma\mathbf{w} = \sigma^2$, $\mathbf{w}'\mathbf{1} = 1$ and $\mathbf{w}'\boldsymbol{\mu} = \mu_q$ where $\mathbf{w} \geq 0$. If σ is chosen to be sufficiently small (e.g. $\sigma = \sigma_a$ in figure 6), the solutions of this problem lie on a connected curve. However, when increasing σ over a certain value σ^* , the curve becomes disconnected leading to the discussed problem of non-ergodicity.

In order to overcome this difficulty, we allow σ to change. As on page 7 we define $\mathbb{S}_{\mathbb{F}}(\sigma, \mu)$ as the set of feasible weights belonging to a (σ, μ) position,

$$\mathbb{S}_{\mathbb{F}}(\sigma_q, \mu_q) = \{ \mathbf{w} \mid \mathbf{w}'\mathbf{1} = 1, \mathbf{w}'\boldsymbol{\mu} = \mu_q, \mathbf{w}'\Sigma\mathbf{w} = \sigma_q^2, \mathbf{w} \geq 0 \} .$$

When letting σ change during the random walk the newly obtained set is given by an expected return value μ_q

$$\mathbb{S}_{\mathbb{F}}(\mu_q) = \bigcup_{\sigma} \mathbb{S}_{\mathbb{F}}(\sigma, \mu_q) \quad \text{with } \sigma \in [\sigma_a, \sigma_b] ,$$

where the surface of $\mathbb{S}_{\mathbb{F}}(\sigma_a, \mu_q)$ is connected.

Generating a random vector $\boldsymbol{\delta}\mathbf{w}^{(i)}$ with properties $\mathbf{1}'\boldsymbol{\delta}\mathbf{w}^{(i)} = 0$, $\boldsymbol{\mu}'\boldsymbol{\delta}\mathbf{w}^{(i)} = 0$ allows for moving randomly from $\mathbf{w}^{(i)}$ to $\mathbf{w}^{(i+1)} = \mathbf{w}^{(i)} + \boldsymbol{\delta}\mathbf{w}^{(i)}$ where $\mathbf{w}^{(i)}$, $\mathbf{w}^{(i+1)} \in \mathbb{S}_{\mathbb{F}}(\mu_q)$.²⁰

Moreover, randomly generating all elements of $\boldsymbol{\delta}\mathbf{w}^{(i)}$ within a symmetric interval $[-\epsilon, \epsilon]$ so that $P(\boldsymbol{\delta}\mathbf{w}^{(i)}) = P(-\boldsymbol{\delta}\mathbf{w}^{(i)})$ ensures that the transition

²⁰Therefore, $|\boldsymbol{\delta}\mathbf{w}^{(i)}|$ must be sufficiently small.

probabilities $W(w \rightarrow w')$ satisfy the *detailed balance condition* for the uniform stationary distribution, i.e. $W(w \rightarrow w') = W(w' \rightarrow w)$.²¹ This condition and ergodicity are sufficient for the stationary distribution to be unique and uniform.

When running this algorithm, one additional problem arises. For small and large values of σ the surfaces describing set $\mathbb{S}_{\mathbb{F}}(\sigma, \mu_q)$ are too small and thus, these σ values are rarely visited in the Markov chain. As a consequence, one obtains too low statistics to determine the distribution of R_{OLS}^2 for these values of σ . One can overcome this problem by artificially changing the probability distribution: instead of a uniform distribution $P(\mathbf{w}) = \text{const}$ we can specify a distribution with probability density

$$P(w) \propto f(\sigma(w))$$

with an appropriately chosen weight function $f(\sigma)$. This will influence the relative probabilities to have different σ values but it will not change the distribution *within* a fixed value of σ , in particular, the R_{OLS}^2 distributions we are interested in.

This modified distribution can be obtained by introducing an extra acceptance probability in the Metropolis algorithm: one accepts the change $w \rightarrow w'$ (where w' is feasible) with a probability of

$$W_{acc}(w \rightarrow w') = \min \left\{ 1, \frac{f(\sigma(w'))}{f(\sigma(w))} \right\}.$$

This satisfies the detailed balance with the stationary distribution $P(w) \propto f(\sigma(w))$, i.e.

$$P(w)W_{acc}(w \rightarrow w') = P(w')W_{acc}(w' \rightarrow w).$$

²¹For convenience, we will use w for $\mathbf{w}^{(i)}$ and w' for $\mathbf{w}^{(i+1)}$ in the following.

At this point, the remaining step is to specify the weight function $f(\sigma)$. An appropriate choice of the weight function can be obtained by the following strategy: one divides the range $[\sigma_a, \sigma_b]$ into many bins and simulates the system (without the weight function $f(\sigma)$) allowing to move only within one bin. Then one chooses $f(\sigma)$ in a way to obtain an approximately flat distribution within this bin. Repeating this in all bins one can determine a continuous function $f(\sigma)$ for the whole range $[\sigma_a, \sigma_b]$.

We have chosen the following parametrization:

$$f(\sigma) = A_i e^{\gamma_i(\sigma - \sigma_i)} \quad \text{for } \sigma \in [\sigma_i, \sigma_{i+1}]$$

$$\text{where } i = 0, 1 \dots N - 1 \quad \text{and } \sigma_0 = \sigma_a, \sigma_N = \sigma_b.$$

The amplitudes A_i are chosen to achieve continuity of $f(\sigma)$:

$$A_{i+1} = A_i e^{\gamma_i(\sigma_{i+1} - \sigma_i)},$$

$$A_0 = 1,$$

where the γ_i 's are determined by the simulation within the corresponding bin.

Finally, simulating the system in the range $[\sigma_a, \sigma_b]$ will allow for covering the whole range in σ . This procedure avoids the problem of disconnected regions at larger values of σ (as in figure 6 where $\sigma > \sigma^*$). During the simulation the system can move down to lower values of σ (i.e. $\sigma < \sigma^*$) where the region is connected and move up into another disconnected region. By this way we ensure ergodicity and maintain a uniform distribution in $\mathbb{S}_{\mathbb{F}}(\sigma, \mu_q)$.

References

COCHRANE, J. H. (2001): *Asset Pricing*. Princeton University Press.

- COCHRANE, J. H., AND J. SAA-REQUEJO (2000): “Beyond Arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets,” *Journal of Political Economy*, 108(1), 79–119.
- FAMA, E. F., AND K. R. FRENCH (1992): “The Cross-Section of Expected Stock Returns,” *Journal of Finance*, 47(2), 427–465.
- GRAUER, R. R. (1999): “On the Cross-Sectional Relation between Expected Returns, Betas, and Size,” *Journal of Finance*, 54(2), 773–789.
- KANDEL, S., AND R. F. STAMBAUGH (1995): “Portfolio Inefficiency and the Cross-Section of Expected Returns,” *Journal of Finance*, 50(1), 157–184.
- MARKOWITZ, H. M., G. P. TODD, AND W. SHARPE (2000): *Mean Variance Analysis in Portfolio Choice and Capital Markets*. John Wiley & Sons.
- NIEDERMAYER, A. F., AND D. NIEDERMAYER (2006): “Applying Markowitz’s Critical Line Algorithm,” *University of Bern Working Paper*.
- ROLL, R. (1977): “A Critique of the Asset Pricing Theory’s Tests; Part I,” *Journal of Financial Economics*, 4(2), 129–176.
- ROLL, R., AND S. A. ROSS (1994): “On the Cross-Sectional Relation between Expected Returns and Betas,” *Journal of Finance*, 49(1), 101–121.
- SHARPE, W. F. (1991): “Capital Asset Prices with and without Negative Holdings,” *Journal of Finance*, 46(2), 489–509.

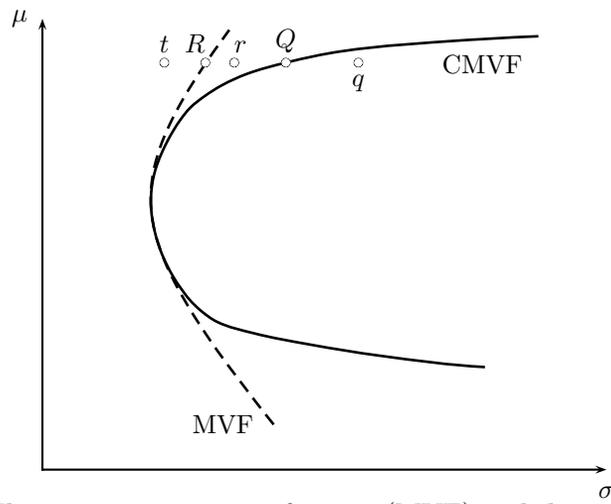


Figure 1: The minimum variance frontier (MVF) and the constrained minimum variance frontier (CMVF). Five cases can be distinguished when moving rightwards from point t .

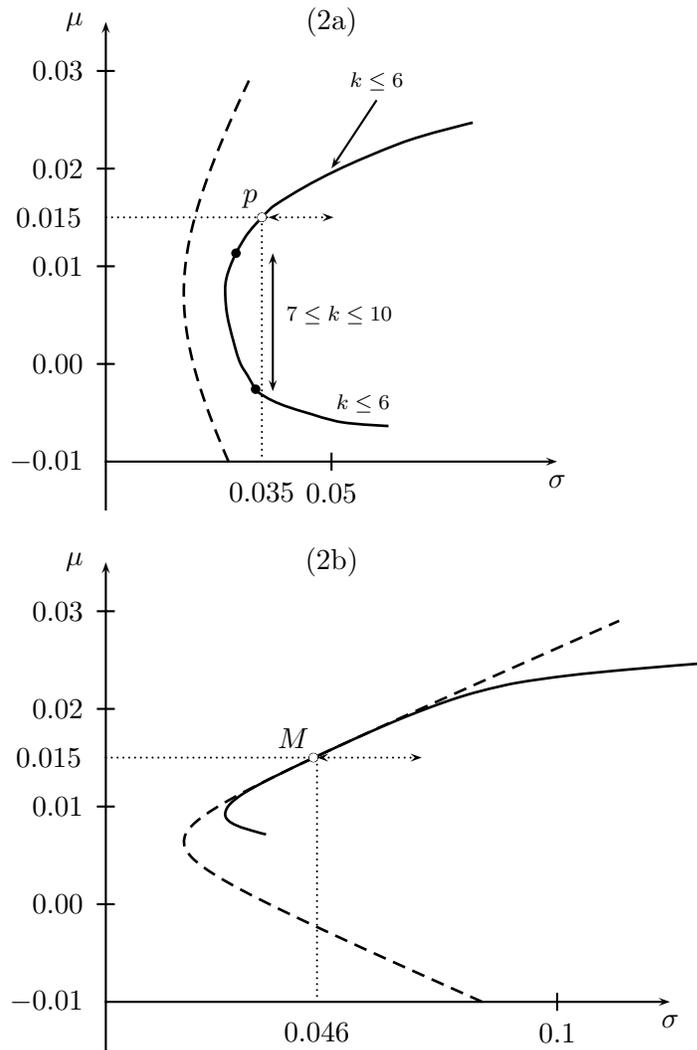


Figure 2: Panel (2a) shows monthly data from the Dow Jones Industrial Average from 2001 to 2006. Panel (2b) depicts a scenario where the CAPM is true. Here, the market proxy M does not contain negative weights.

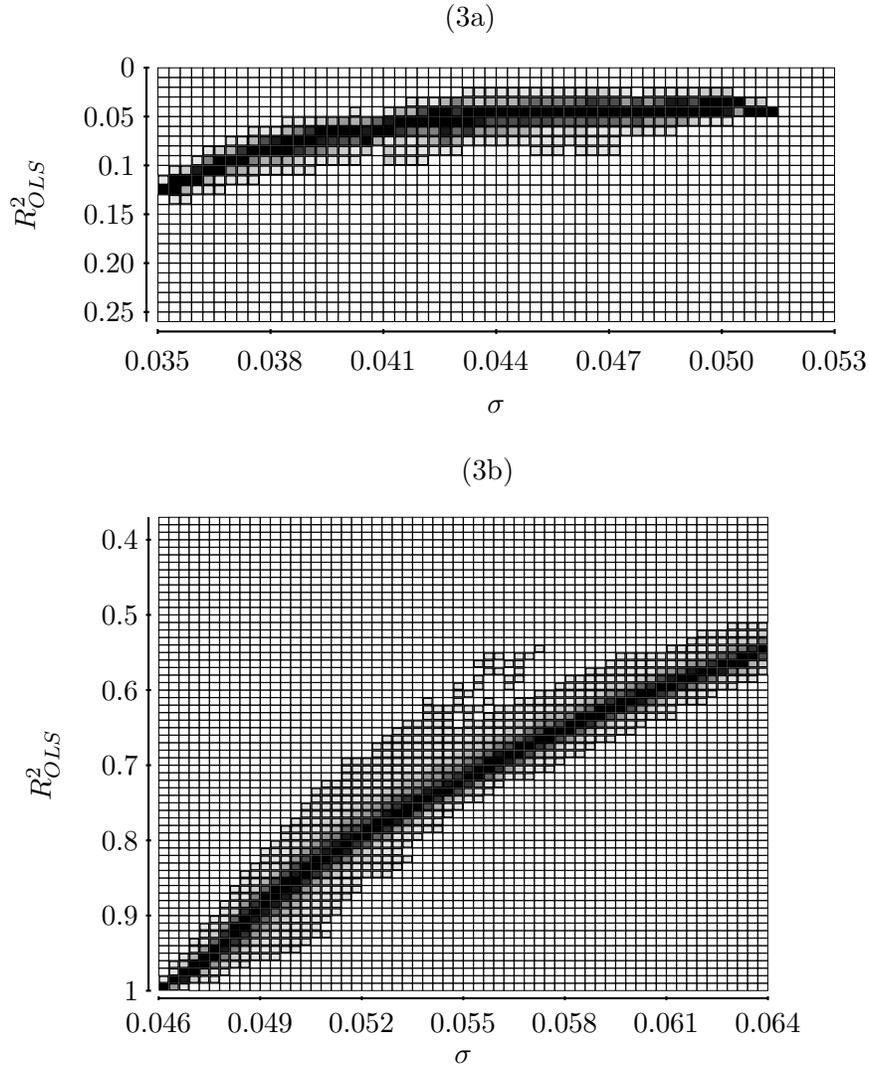


Figure 3: Figure 3a: R_{OLS}^2 histogram for the Dow Jones Industrial Average where $\mu_p = 0.015$. The values of σ are given by the distance shown by the dotted arrow in figure 2a. The number of iterations for the simulation was set to 5 million. The black shaded squares contain densities of over 90% of the distribution's maximum density. The brightest squares contain densities below 10% of the distribution's maximum density. Figure 3b: Shows the histograms that belong to the dotted arrow in figure 2b. where the proxy has only non-negative weights.

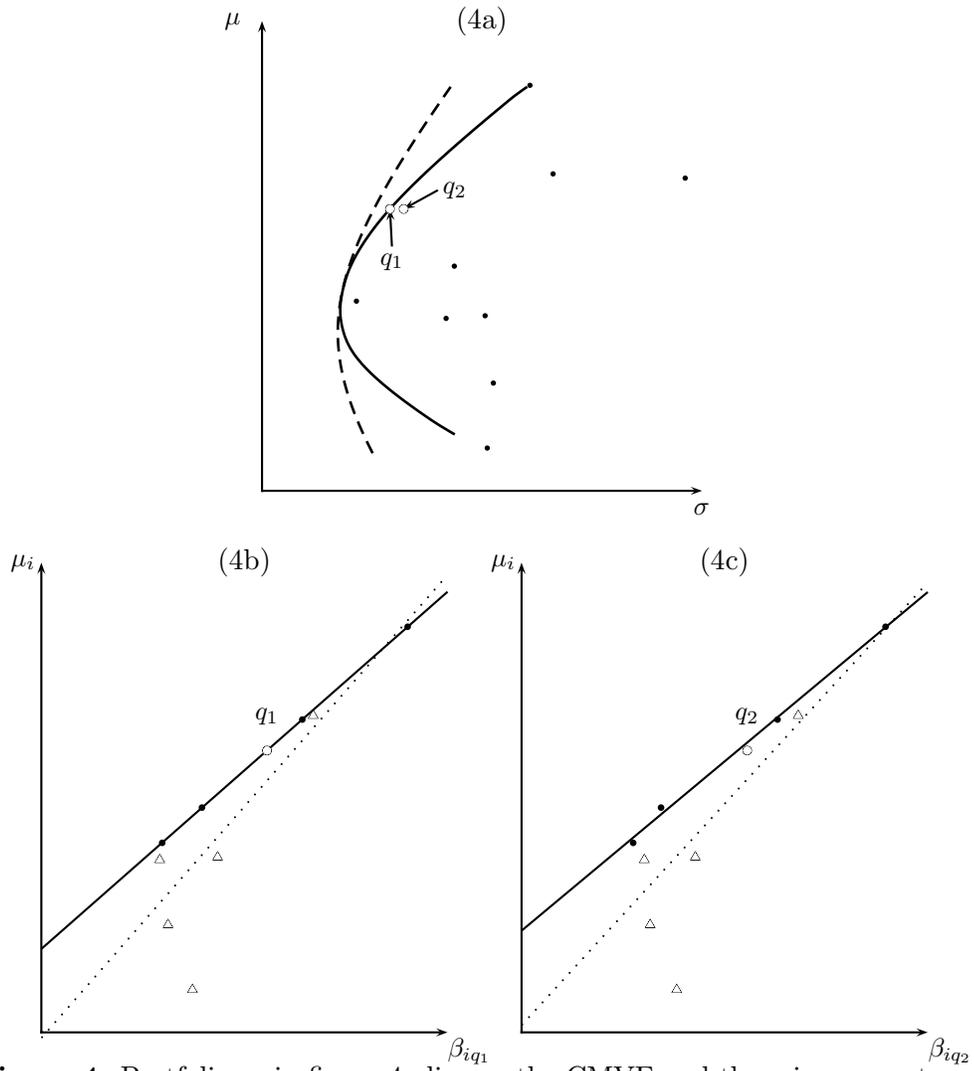


Figure 4: Portfolio q_1 in figure 4a lies on the CMVF and there is an exact linear relation between the non-zero weighted assets in portfolio q_1 (as shown in figure 4b). Portfolio q_2 in figure 4a has the same expected return level as q_1 but is inefficient. However, the linear relation among the originally examined assets is still very strong (see figure 4c).

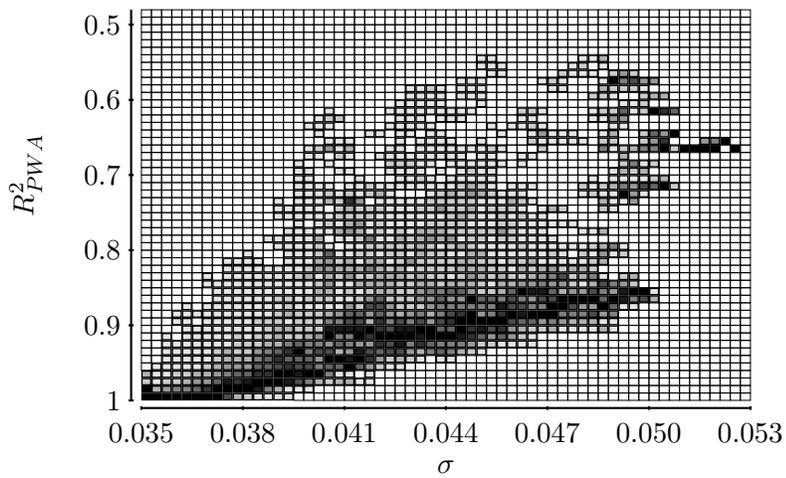


Figure 5: Distributions of R^2_{PWA} for the same $\mu_q = 0.015$ and for different σ_q are plotted horizontally.

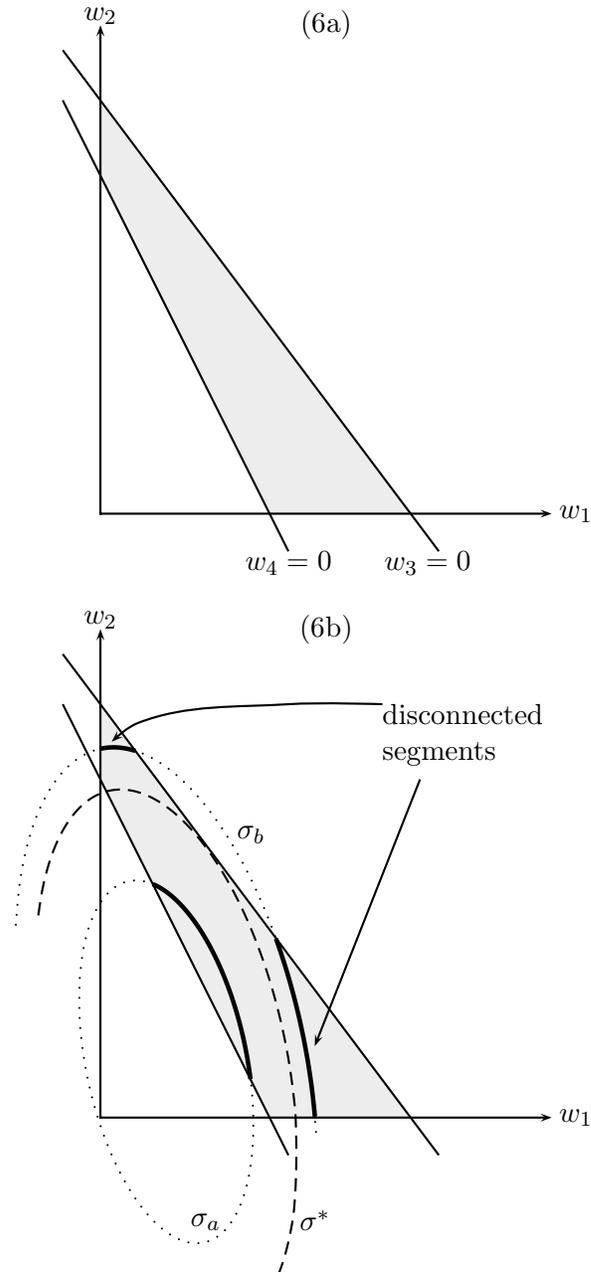


Figure 6: Example with 4 assets; as shown in panel (a) the feasible region where $w \geq 0$ and where both linear constraints are satisfied is shown by the shaded area. Panel (b) shows ellipses describing the solutions for $w' \Sigma w = \sigma^2$, $w' \mathbf{1} = 0$, $w' \mu = \mu_q$ for different values of σ . If $\sigma < \sigma^*$, the solutions form a connected line (i.e. segment of an ellipse). If $\sigma > \sigma^*$ the solutions are on disconnected line segments.