Available online at www.sciencedirect.com
SCIENCE DIRECT®

Journal of Empirical Finance 12 (2005) 556–575

Journal of EMPIRICAL FINANCE

www.elsevier.com/locate/econbase

Evaluating the importance of missing risk factors using the optimal orthogonal portfolio approach

Hossein Asgharian, Björn Hansson *

Department of Economics, Lund University, Box 7082 S-22007 Lund, Sweden

Accepted 13 September 2004
Available online 25 August 2005

Abstract

We apply the orthogonal portfolio approach to analyse the importance of risk factors potentially missing from the CAPM. We generalize the approach proposed by MacKinlay and Pastor (2000) [MacKinlay, A.C., Pastor, L., 2000. Asset pricing models: implications for expected returns and portfolio selection. Review of Financial Studies 13, 883–916] by estimating the Sharpe ratio of the optimal orthogonal portfolio. Our result, based on US industry portfolios for the period 1927–2002, reveals important risk factors missing from the CAPM during periods with high market volatility. We show that a priori fixing the Sharpe ratio, which is an assumption used by MacKinlay and Pastor (2000) [MacKinlay, A.C., Pastor, L., 2000. Asset pricing models: implications for expected returns and portfolio selection. Review of Financial Studies 13, 883–916], may produce less plausible estimates of the expected returns.

JEL classification: G11; G12
Keywords: Orthogonal portfolio; Investment strategy; Factor pricing; Simulated annealing

1. Introduction

Factor-based asset pricing models have been common in finance since the work by Merton (1973) and Ross (1976). However, these theoretical contributions give only

* Corresponding author. Tel.: +46 462228668; fax: +46 462224118.
E-mail addresses: Hossein.Asgharian@nek.lu.se (H. Asgharian), Bjorn.Hansson@nek.lu.se (B. Hansson).

0927-5398/$ - see front matter
doi:10.1016/j.jempfin.2004.09.002

general directions for the empirical identification of the factors. There are different strategies for building factor models when the factors are unobserved. On the one hand, there are the statistical approaches that start from the sample covariance matrix, e.g., the asymptotic principal components approach (see Connor and Korajczyk, 1986) and the factor analysis approach (see Lehman and Modest, 1988). On the other hand, one can start from grouping assets on the basis of certain characteristics that are related to differences in the average returns (see Fama and French, 1993). In a new approach, MacKinlay and Pastor (2000) use information both from the covariance matrix and average returns to define a factor model. They assume a factor-based asset pricing model, but not all factors are known and the so-called "optimal orthogonal portfolio" represents the unknown factors.[1] The factor model is estimated by relating the asset mispricing and the residual covariance matrix. It is well known that the covariance matrix can be estimated more accurately than the expected returns (see Merton, 1980). Thus, the imposed link between asset mispricing and the residual covariance matrix may result in better estimates of the expected returns. MacKinlay and Pastor (2000) use these estimates, henceforth denoted as the "restricted" estimates, to determine the tangency portfolio within a mean-variance framework. They focus on the usefulness of the suggested factor models for portfolio management purposes but not on the adequacy of these models from an asset pricing perspective.

We have a twofold purpose in our use of the orthogonal portfolio approach: first, we investigate, from an asset pricing angle, the time-varying importance of potential risk factors missing from the CAPM; then we evaluate the economic importance of these factors from the viewpoint of portfolio management.

Our approach generalises the orthogonal portfolio approach proposed by MacKinlay and Pastor (2000) in two ways. First, they consider the squared Sharpe ratio of the orthogonal factor as given, while in our analysis it is endogenously determined and can vary over time—hence, we let the model decide the importance of the orthogonal factor in each estimation period. Second, MacKinlay and Pastor (2000) assume an identical idiosyncratic risk of the test assets, but our approach allows the test assets to have different idiosyncratic risk. In addition, we examine a different set of test assets, namely, US industrial portfolios over the period 1927 to 2002, whereas MacKinlay and Pastor (2000) use size portfolios and a number of individual stocks from the Dow Jones Industrial Index.

We use two alternative models: a model without any observed factor, denoted the "no-factor model", where the orthogonal factor represents all factors, and a model with the market index as the observed factor, denoted the "one-factor model", where the orthogonal portfolio represents all potential factors left out by the market index. A comparison of the factor risk premium from the first model with the mean excess return of the market index may uncover the existence of some latent factor or factors in addition to the market index. The second model estimates the contribution of the latent factors directly and separates the effect of the missing factors from the market index. Consequently, the sample mean excess

[1] The optimal orthogonal portfolio has its starting point in Roll (1980) and has been used by MacKinlay (1995) to discuss the ability of the conventional F-test of asset pricing models.

return can be divided into three components: one is related to the market risk; one is due to unobserved risk factors that are represented by the optimal orthogonal portfolio; finally, the remaining component may be considered the result of market inefficiencies or overreactions.

We first analyse the importance of the orthogonal portfolio as an agent of the priced factors in an asset pricing context. The focus is then on the estimated expected returns, and we look in particular at the effects of fixing the squared Sharpe ratio of the orthogonal portfolio. Finally, we evaluate the orthogonal portfolio approach from a portfolio perspective, and in this exercise, we use several alternative models besides the orthogonal approach. We use the out-of-sample Sharpe ratio of the tangency portfolio for each model as the evaluation metric.

The results show the advantage of using the optimal orthogonal portfolio both as a sole factor and as a complement to the market index. For the case with no observed factor, the orthogonal portfolio always has a positive risk premium, which is a necessary characteristic of the true "market portfolio" if CAPM is the correct model.[2] For the case with one observed factor, the orthogonal portfolio captures the latent factors active during periods of high market volatility, i.e., the 1930s and the 1970s, and these factors are not represented by the market index. This is evident from the relatively high Sharpe ratio of the orthogonal portfolio as well as the significant contribution of the orthogonal factor in explaining mean excess returns of several industries during these periods. After considering the effect of the latent priced factors together with the market index, the non-risk-based component is seldom significant. On the whole, we interpret these results as support for an exact two-factor pricing model, which converges to CAPM during the periods when the orthogonal portfolio is unimportant.

Our results show that MacKinlay and Pastor's (2000) presumption of a fixed squared Sharpe ratio of the orthogonal portfolio has significant empirical consequences for the estimation of expected returns and may lead to inaccurate estimates of the expected industry returns, both with respect to time variation and level.

The estimated expected returns from the models that assume identical idiosyncratic risk of the test assets do not differ from the models that allow for different idiosyncratic risk. This supports MacKinlay and Pastor's (2000) conjecture that the potential differences in residual variances are not important for assets belonging to the same asset class.

The performances of the optimal portfolios constructed according to different models are assessed by their out-of-sample Sharpe ratios. Our results indicate that all models exploiting a risk and return relationship, except the one-factor restricted models, have essentially the same performance and outperform the model based on the sample moments, which is largely in accordance with the results of MacKinlay and Pastor (2000). The poor performance of the one-factor restricted model, relative to the other factor models, reflects the non-permanent nature of its orthogonal portfolio rather than the insufficiency of the market index as stated by MacKinlay and Pastor (2000). Moreover, the models with an

[2] This is not necessarily so if consumption CAPM is the true model.

endogenous Sharpe ratio, despite their in-sample advantage for estimation of expected returns, have no important effects on the out-of-sample Sharpe ratios.

The outline of the paper is as follows: Section 2 covers the definition of the factor models and the applied econometric methods; Section 3 analyses the empirical results; and Section 4 concludes the study.

2. Econometric models

2.1. An exact factor model with an orthogonal portfolio

Consider the following linear K-factor pricing model for N assets:

$$\begin{aligned}
&\mathbf{R}_t = \boldsymbol{\alpha} + \boldsymbol{\beta}\mathbf{F}_t + \boldsymbol{\varepsilon}_t,\\
&E[\mathbf{R}_t] = \boldsymbol{\mu}\\
&E[(\mathbf{R}_t - \boldsymbol{\mu})(\mathbf{R}_t - \boldsymbol{\mu})'] = \mathbf{V}\\
&E[\boldsymbol{\varepsilon}_t] = 0, E[\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t'] = \boldsymbol{\Sigma}\\
&E[\mathbf{F}_t] = \boldsymbol{\mu}_F, E[(\mathbf{F}_t - \boldsymbol{\mu}_F)(\mathbf{F}_t - \boldsymbol{\mu}_F)'] = \boldsymbol{\Omega}_F\\
&\mathrm{Cov}[F_{kt}, \varepsilon_{it}] = 0 \text{ for } \forall\ k \text{ and } i,
\end{aligned} \tag{1}$$

where $\mathbf{R}_t$ is an $(N \times 1)$ vector of excess returns for N assets, $\boldsymbol{\alpha}$ is an $(N \times 1)$ vector of model mispricing, i.e., the part of the excess return that is not captured by the factors, $\boldsymbol{\beta}$ is an $(N \times K)$ matrix of the sensitivities of N assets to the K factors and $\mathbf{F}$ is a $(K \times 1)$ vector of excess returns for the K factors.

In an exact factor pricing model, expected returns are exclusively explained by the included risk factors, i.e., there are no non-risk-based components in expected returns, and all the elements of the vector $\boldsymbol{\alpha}$ are zero. In this case, the portfolio with the maximum Sharpe ratio (the tangency portfolio in the mean-standard deviation space) can be constructed as a linear combination of the K known factors. On the other hand, if the K factors are not sufficient to explain the expected returns, then some elements of $\boldsymbol{\alpha}$ will differ from zero. In this case, an exact factor pricing model can be achieved by adding an optimal orthogonal portfolio, h, as a complement factor to the K known factors (see MacKinlay, 1995):

$$\begin{aligned}
&\mathbf{R}_t = \boldsymbol{\beta}_h R_{ht} + \boldsymbol{\beta}\mathbf{F}_t + \mathbf{u}_t,\\
&E[\mathbf{u}_t\mathbf{u}_t'] = \boldsymbol{\Phi}\\
&E[R_{ht}] = \mu_h\\
&E\left[(R_{ht} - \mu_h)^2\right] = \sigma_h^2,\\
&\mathrm{Cov}[R_{ht}, F_{kt}] = 0 \text{ for } \forall\ k \text{ and } t,
\end{aligned} \tag{2}$$

where R_{ht} is the excess return on the optimal orthogonal portfolio, h, and $\boldsymbol{\beta}_h$ is an $(N \times 1)$ vector of the sensitivities of N assets to the factor h. Since h is orthogonal to the K known factors, it can be added to the factor model without changing the sensitivity matrix $\boldsymbol{\beta}$. This construction allows us to relate the mispricing vector, $\boldsymbol{\alpha}$, to the residual covariance matrix, $\boldsymbol{\Sigma}$. Combining Eq. (1) with Eq. (2) gives

$$\boldsymbol{\alpha} = \boldsymbol{\beta}_h \mu_h$$

$$\boldsymbol{\Sigma} = \boldsymbol{\beta}_h \boldsymbol{\beta}_h' \sigma_h^2 + \boldsymbol{\Phi} = \boldsymbol{\alpha}\boldsymbol{\alpha}' \frac{\sigma_h^2}{\mu_h^2} + \boldsymbol{\Phi}. \tag{3}$$

MacKinlay and Pastor (2000) use this link to provide a joint estimate of $\boldsymbol{\alpha}$ and $\boldsymbol{\Sigma}$, which are then used to estimate the mean and the covariance matrix of the returns, $\boldsymbol{\mu}$ and $\mathbf{V}$, as:

$$\hat{\boldsymbol{\mu}} = \hat{\boldsymbol{\alpha}} + \hat{\mathbf{B}}\hat{\boldsymbol{\mu}}_F$$

$$\hat{\mathbf{V}} = \hat{\boldsymbol{\beta}}\hat{\boldsymbol{\Omega}}_F\hat{\boldsymbol{\beta}}' + \hat{\boldsymbol{\Sigma}}. \tag{4}$$

The means and the covariance matrix of the known factors, $\hat{\mu}_F$ and $\hat{\Omega}_F$, are estimated from the sample, but $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\Sigma}$ are estimated by maximising the following log likelihood function:

$$\ln L(\boldsymbol{\alpha}, \boldsymbol{\beta}, \theta_h, \boldsymbol{\Phi}) \propto -\frac{T}{2}\det(\boldsymbol{\alpha}\boldsymbol{\alpha}'\theta_h + \boldsymbol{\Phi}) - \frac{1}{2} \times \sum_{t=1}^{T} (\mathbf{R}_t - \boldsymbol{\alpha} - \boldsymbol{\beta}\mathbf{F}_t)'(\boldsymbol{\alpha}\boldsymbol{\alpha}'\theta_h + \boldsymbol{\Phi})^{-1}(\mathbf{R}_t - \boldsymbol{\alpha} - \boldsymbol{\beta}\mathbf{F}_t), \tag{5}$$

where θ_h is $(\sigma_h^2)/(\mu_h^2)$, the inverse squared Sharpe ratio of the orthogonal portfolio, which is an unknown parameter and should be estimated. The higher the theta, the more important the contribution of the observed factor in explaining expected return. The matrix $\boldsymbol{\Phi}$ is restricted to be diagonal, which is justified by the assumption that all the covariations among asset returns are explained by the K known factors, plus the optimal orthogonal portfolio. MacKinlay and Pastor (2000) restrict $\boldsymbol{\Phi}$ to be equal to $\sigma^2\mathbf{I}$, i.e., the idiosyncratic risks among assets are equal. This restriction is very strong, but it has the advantage of decreasing the number of estimated parameters. Therefore, we examine the model both with and without this restriction.

2.2. An exact factor model with only observed factors

An alternative way to estimate the mean and the covariance matrix is to assume that the known factor model is sufficient for explaining expected returns, this model is denoted an "exact factor model". The vector of the mean returns is now defined by

$$\hat{\boldsymbol{\mu}} = \hat{\boldsymbol{\beta}}\hat{\boldsymbol{\mu}}_F, \tag{6}$$

where $\hat{\boldsymbol{\beta}}$ is the estimated sensitivities resulting from N time series regressions of the model in Eq. (1). The estimated intercepts are considered to be due to sampling error or other unsystematic sources, and they are not supposed to persist over time. Hence, these intercepts are neglected in estimation of the expected returns. Since the factor structure is assumed to take into account the entire cross-covariances of the returns, it follows that the variance–covariance matrix of the residuals will be diagonal. The covariance matrix of the returns is

$$\hat{\mathbf{V}} = \hat{\boldsymbol{\beta}}\hat{\boldsymbol{\Omega}}_F\hat{\boldsymbol{\beta}}' + \hat{\mathbf{D}}, \tag{7}$$

where $\hat{\mathbf{D}}$ is a $(N \times N)$ diagonal matrix of the estimated residual variances for N time series regressions of the model in Eq. (1). As before, $\hat{\boldsymbol{\mu}}_F$ and $\hat{\boldsymbol{\Omega}}_F$ are both estimated from the sample.

2.3. Optimal portfolio

We use the estimated mean vectors and the covariance matrices to construct the optimal mean-variance portfolio, i.e., the portfolio with the maximum Sharpe ratio (MSR). The portfolio weights can be defined as

$$\mathbf{w}_{\mathrm{msr}} = \left(\boldsymbol{\iota}'\hat{\mathbf{V}}^{-1}\hat{\boldsymbol{\mu}}\right)^{-1}\hat{\mathbf{V}}^{-1}\hat{\boldsymbol{\mu}}, \tag{8}$$

where $\boldsymbol{\iota}$ is a $(N \times 1)$ vector of ones, and $\hat{\boldsymbol{\mu}}$ and $\hat{\mathbf{V}}$ are the estimated mean vector and covariance matrix, respectively. For the case with one observed factor and one optimal orthogonal portfolio, $\hat{\boldsymbol{\mu}}$ is defined in Eq. (4) and

$$\hat{\mathbf{V}} = \hat{\boldsymbol{\alpha}}\hat{\boldsymbol{\alpha}}'\hat{\theta}_h + \hat{\boldsymbol{\beta}}\hat{\boldsymbol{\Omega}}_F\hat{\boldsymbol{\beta}}' + \hat{\boldsymbol{\Phi}}. \tag{9}$$

In the case with no observed factor, $\boldsymbol{\mu}$ is equal to $\boldsymbol{\alpha}$ and $\mathbf{V}$ is equal to $\boldsymbol{\Sigma}$, and using the relationship between $\boldsymbol{\Sigma}$ and $\boldsymbol{\Phi}$ from Eq. (3), the weights can be defined in terms of $\boldsymbol{\alpha}$ and $\boldsymbol{\Phi}$ as

$$\mathbf{w}_{\mathrm{msr}}\left(\boldsymbol{\iota}'\hat{\boldsymbol{\Phi}}^{-1}\hat{\boldsymbol{\alpha}}\right)^{-1}\hat{\boldsymbol{\Phi}}^{-1}\hat{\boldsymbol{\alpha}}. \tag{10}$$

Since $\boldsymbol{\alpha}$ is linked to the covariance matrix $\boldsymbol{\Sigma}$, the residual covariance matrix $\boldsymbol{\Phi}$ replaces the total covariance matrix $\mathbf{V}$ in the calculation of the weights. With the strong assumption of the equality of the residual variances, the matrix $\boldsymbol{\Phi}$ will be proportional to an identity matrix and Eq. (10) will be reduced to

$$\mathbf{w}_{\mathrm{msr}} = \left(\boldsymbol{\iota}'\hat{\boldsymbol{\alpha}}\right)^{-1}\hat{\boldsymbol{\alpha}}. \tag{11}$$

In fact, we will get this simple expression for any estimator of the mean as long as the return covariance matrix is proportional to the identity matrix.

If the restricted factor model includes one observed factor portfolio and its optimal orthogonal portfolio h, then the tangency portfolio can always be constructed as a linear combination of these two portfolios. More specifically, these two portfolios form a frontier that has one common portfolio with the frontier constructed from all assets, and this common portfolio is the portfolio with the maximum Sharpe ratio, i.e., the tangency portfolio.

2.4. Selected models

We use the following four restricted factor models that impose a link between the mispricing and the residual covariance matrix:

i. Model with no observed factor with different diagonal elements in $\mathbf{\Phi}$, NFD.
ii. Model with no observed factor and the matrix $\mathbf{\Phi}$ is assumed to be proportional to an identity matrix, NFI.
iii. Model with one observed factor represented by the value-weighted market index with different diagonal elements in $\mathbf{\Phi}$, OFD.
iv. Model with one observed factor represented by the value-weighted market index and the matrix $\mathbf{\Phi}$ is assumed to be proportional to an identity matrix, OFI.

We also use an exact factor model, OFE where the market portfolio represented by the value-weighted market index is the only factor. In addition, we use the sample means and sample covariance matrix, Sample, as the unrestricted estimates of the return moments.

For portfolio selection analysis, we use two other approaches in addition to the models above to define the weights of the MSR portfolios:

i. Sample means and a covariance matrix that is proportional to an identity matrix, Sample-I. In this case, the covariance matrix in Eq. (8) is replaced by the identity matrix for the construction of the MSR portfolio. If there are no non-risk-based components in average return, i.e., the restricted estimate of the expected return from the orthogonal portfolio approach with no observed factor is equal to the realised mean return from the Sample, then this model is equivalent to the optimal portfolio according to NFI (see Eq. (11)).
ii. The equally weighted portfolio, EW, is equivalent to an optimal portfolio according to the NFD with an additional assumption that the estimated expected returns are proportional to the diagonal elements of $\mathbf{\Phi}$ (see Eq. (10)).

2.5. Estimation method

The estimation of the parameters of the restricted model, particularly the one-factor model, requires a solution of a difficult and non-linear likelihood function. As mentioned by MacKinlay and Pastor (2000), it is particularly difficult to estimate the model with theta as an unknown parameter. They tackle this problem by fixing θ_h a priori. To further reduce the number of parameters, they assume $\mathbf{\Phi}$ to be proportional to an identity matrix. These simplifications, especially the restriction on θ_h, may result in incorrect inferences. We use therefore the simulated annealing approach (see Goffe et al., 1994) to estimate the likelihood function in Eq. (5). Simulated annealing is a random-search optimisation technique based on a Metropolis Monte Carlo simulation, which is developed to deal with highly nonlinear problems. The advantages of this approach compared to the conventional methods are that it is very robust and fails quite seldom even for very complicated problems (see Corana et al., 1987; Goffe et al., 1994). Thus, in contrast to MacKinlay and Pastor (2000), we can let the model determine the importance of the unobserved factor, the orthogonal portfolio, and allow for different idiosyncratic risk of the test assets.

Due to complexity of the likelihood function a bootstrap procedure is used for statistical tests. We resample the original data with replacement and reestimate the model for the "new" data. The procedure is repeated 200 times and we use the empirical confidence intervals for inference.

All the parameter estimates are based on a 10-year moving window. The parameters for each model are then used to estimate the expected industry returns and the return covariance matrix, in order to calculate the optimal portfolio weights for that model. These weights are applied during the following month to form the out-of-sample returns of the optimal portfolios. The estimates are updated each month. Finally, the Sharpe ratios of the entire time-series of these returns are used to compare the models. We use Jobson and Korkie (1981) for the significance test of the Sharpe ratios as well as for testing the equality of two Sharpe ratios.

3. Analysis

We first examine the expected industry returns estimated by the different models. In particular, we compare the no-factor restricted models, where an orthogonal portfolio represents all latent risk factors, and the one-factor restricted model, where the market index is an observed factor and an orthogonal portfolio represents all other latent risk factors. Secondly, we analyse the effects of fixing the Sharpe ratio of the orthogonal portfolio, which is the approach taken by MacKinlay and Pastor (2000). Finally, we assess the power of the models by comparing the out-of-sample Sharpe ratios—there are now several competitors besides the orthogonal approach.

The data cover the period 1927–2002 and consist of monthly returns for US industry portfolios, the market index and the risk-free rate. We use 10 value-weighted industrial portfolios as test assets. The industrial portfolios represent the following branches: non-durables (NoDur), durables (Durbl), oil (Oil), chemical (Chems), manufacturing (Manuf), telecoms (Telcm), utilities (Utils), shops (Shops), money (Money) and others (Other). All data are from the data library on Kenneth French's home page. Each stock on NYSE, AMEX and NASDAQ is assigned to an industry portfolio at the end of June of year t based on its four-digit SIC code at that time. The returns are then computed from July of t to June of $t+1$. Market return is the value-weighted return on all NYSE, AMEX and NASDAQ stocks (from CRSP). The risk-free rate is the 1-month Treasury bill rate.

In Table 1, the descriptive statistics of the industrial portfolios for the entire period show that all value-weighted industry portfolios have average returns that are significantly different from zero at the 1% level. The monthly Sharpe ratios range from 0.074 for Other to 0.137 for NoDur and the value-weighted market index has a ratio of 0.112.

Table 1
Summary statistics of the industrial portfolios

	Mean	Standard deviation	*t*-statistics	Sharpe ratio
NoDur	0.007	0.049	4.141**	0.137
Durbl	0.008	0.074	3.340**	0.111
Oil	0.007	0.060	3.589**	0.119
Chems	0.008	0.059	3.960**	0.131
Manuf	0.007	0.071	2.923**	0.097
Telcm	0.005	0.048	3.372**	0.112
Utils	0.005	0.058	2.828**	0.094
Shops	0.007	0.063	3.310**	0.110
Money	0.008	0.069	3.300**	0.109
Other	0.005	0.069	2.242*	0.074
Market	0.006	0.055	3.375**	0.112

The table shows the summary statistics of the monthly excess returns on the value-weighted industrial portfolios and the market index for the period January 1927–December 2002. The *t*-values marked with one asterisk are significant at the 5% level and those with two asterisks are significant at the 1% level.
The industries are non-durables (NoDur), durables (Durbl), oil (Oil), chemical (Chems); manufacturing (Manuf), telecoms (Telcm), utilities (Utils), shops (Shops), money (Money) and others (Other). The industry portfolios are formed from all NYSE, AMEX, and NASDAQ stocks. Market is the return on the value-weighted return on all NYSE, AMEX, and NASDAQ stocks minus the 1-month Treasury bill rate.

3.1. Expected returns

The estimated expected excess industry returns from the restricted models are presented in Fig. 1.[3] It is evident that the estimated expected excess returns vary over time, with all models exhibiting particularly low values from 1972 to 1982. We can easily see that the models with strong restrictions on the matrix $\mathbf{\Phi}$, NFI and OFI, do not differ from the models that allow for different diagonal elements of $\mathbf{\Phi}$, NFD and OFD. This supports the conjecture by MacKinlay and Pastor (2000) that the potential differences in residual variances across the industries are not important if the assets belong to the same asset class. In the analysis below, therefore, for the sake of space we will not present the results for the models with different idiosyncratic risk, i.e., NFD and OFD.

Fig. 1 shows that the expected excess return for manufacturing is always positive for NFI, while OFI exhibits negative estimates for the 1970s. This pattern is also true for all the other industries. The negative expected industry excess return for OFI is related to the negative mean excess returns of the market index in the 1970s (see the factor risk premium for OFE in the first diagram in Fig. 2) and the importance of the market index in the OFI model. The first diagram in Fig. 2 also shows that the estimated risk premium of the optimal orthogonal portfolio in NFI is always positive, which is the desired characteristic of a true “market portfolio” according to CAPM.[4] The second diagram in Fig. 2 compares the Sharpe ratios

[3] We have shown only the results for manufacturing but all other industries exhibit the same pattern.

[4] If there is only one missing factor, then the optimal orthogonal portfolio should by construction be an exact representation of this factor. However, if there are several missing factors, then the orthogonal portfolio must be an exact manifestation of all these factors. Hence, from our analysis with NFI we cannot exactly differentiate between the following two explanations: the orthogonal portfolio is a better representative of the market portfolio, or it represents the market portfolio plus some other missing factors.

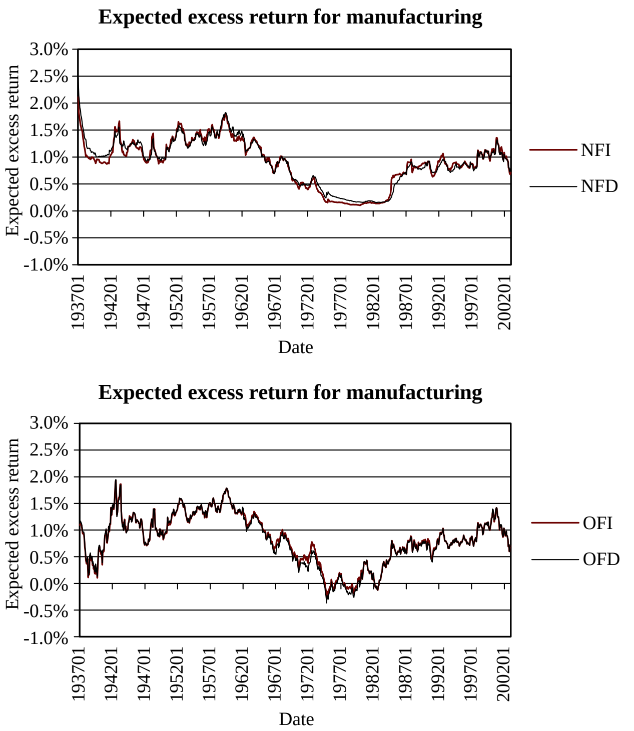

Fig. 1. The expected industry excess returns estimated by the restricted models. The figure plots the estimates of the expected excess industry returns for manufacturing from the no-factor and the one-factor restricted models, respectively. The models are estimated both with and without the strong assumption of equality of the diagonal elements of the residual covariance matrix. The models are no-factor with identical diagonal elements (NFI), one-factor with identical diagonal elements (OFI), no-factor allowing for different diagonal elements (NFD) and one-factor allowing for different diagonal elements (OFD).

of the factor portfolios for NFI and OFE. The Sharpe ratio for NFI is always positive and is very close to but mostly above the OFE. The difference between the factor risk premiums of the two models might be due to the fact that the model with only an orthogonal portfolio (NFI) picks up systematic risk factors that are not represented by the market index.

To analyse the importance of the orthogonal portfolio for OFI, we study the implied Sharpe ratio of the orthogonal portfolio as well as the component in the expected return that is related to the orthogonal portfolio. Fig. 3 plots the implied Sharpe ratio, which is the inverse of the square root of the estimated theta, over time. The results show that there is an obvious time variation in the implied Sharpe ratio: there are some distinctive periods when the ratio is quite high and other periods when it is almost zero.

To measure the overall importance of the orthogonal portfolio over time, it is not necessarily instructive to analyse the implied Sharpe ratio of the orthogonal portfolio in isolation, since the entire portfolio frontier may change across years. We start instead from the following relationship between the tangency portfolio and the factor portfolios: $s_q^2 = s_m^2 + s_h^2$, where s_q^2 is the squared Sharpe ratio of the tangency portfolio, s_m^2 is the

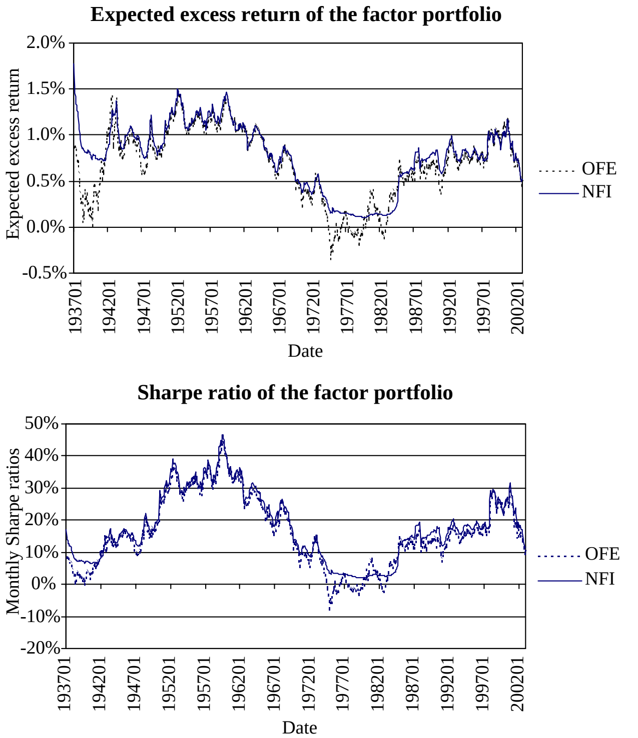

Fig. 2. The factor risk premium estimated by the OFE model and the NFI model. The figure plots the estimated risk premium and the in-sample Sharpe ratio for the factor portfolios of the model with only the market index (OFE) and the model with no observed factor (NFI). The factor risk premium is the expected excess return of the tangency portfolio given by each model. For OFE, this portfolio corresponds to the value-weighted market index. For NFI, this portfolio corresponds to the optimal orthogonal portfolio. All the estimations are based on 10-year moving windows.

squared Sharpe ratio of the observed factor and s_h^2 is the squared Sharpe ratio of the orthogonal factor. Our metric is now the ratio s_h^2/s_q^2 over time, which is presented in the second diagram of Fig. 3. The result shows that the orthogonal factor is only important, both from an absolute and a relative point of view, during certain periods, which implies that during these periods, the orthogonal portfolio represents risk factors that are left out by the market index. In fact, a closer inspection of the data shows that there are approximately three periods during which the orthogonal portfolio is "important" for consecutive months.[5] Roughly speaking, there are three sequences of 10-year periods: starting with 1928–1937 and ending with 1933–1943; starting with 1962–1971 and ending with 1969–1978; and starting with 1981–1990 and ending with 1982–1991. Hence, there are long periods when the market index is the sole important factor in an exact factor pricing model.

Our results show that the orthogonal portfolio is important above all during periods characterised by high market volatility, in particular the 1930s and the 1970s (see, for

[5] We define "importance" as having a value that is above the average of the ratio s_h^2/s_q^2.

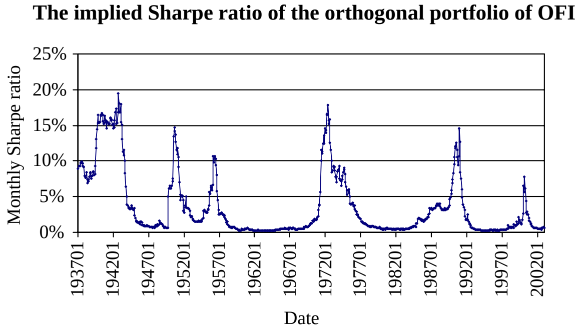

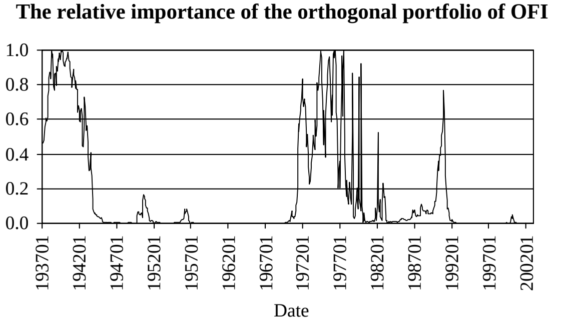

Fig. 3. The importance of the orthogonal portfolio for OFI. The first diagram plots the implied Sharpe ratio of the optimal orthogonal portfolio for the model with one observed factor (OFI). The implied Sharpe ratio is the inverse of the square root of the estimated theta. The second diagram plots the ratio of the squared Sharpe ratio of the orthogonal portfolio to the squared Sharpe ratio of the tangency portfolio. This ratio is a measure of the relative importance of the optimal orthogonal portfolio. All the estimates are based on 10-year moving windows.

example, the variance of NFI in Fig. 6, which is a good approximation of the variance of the market index). This leads us to interpret the orthogonal portfolio as representing a latent factor or factors that is or are mostly active during distress periods and these factors convey risks that are not captured by the market index. The industry betas against the orthogonal factor reflect their relative distress/strength (see Fama and French, 1996, for a discussion of a relative-distress premium). Our result shows that the industry risk premiums related to the orthogonal portfolio, i.e., their estimated alphas from OFI, change over time. These premiums are mostly significant during the periods when the ratio s_h^2/s_q^2 is large. Some industries, for example, Manufacturing, have positive alphas while some other industries, such as Utilities, have negative alphas, meaning that the latter industries are hedges towards the latent factors (see bootstrapped confidence intervals in Fig. 4).

The non-risk-based component is by construction the difference between the average excess return and the estimated risk premium from OFI. Our result shows that the non-risk-based component is seldom significant, but for some industries, it is significant during shorter periods. The best exponent of this pattern is Utility (see Fig. 4).

Our results show that the orthogonal portfolio for OFI is not important all the time and that the non-risk-based component is mostly unimportant. Hence, in general, we have an exact two-factor pricing model with the market index and the orthogonal portfolio as the

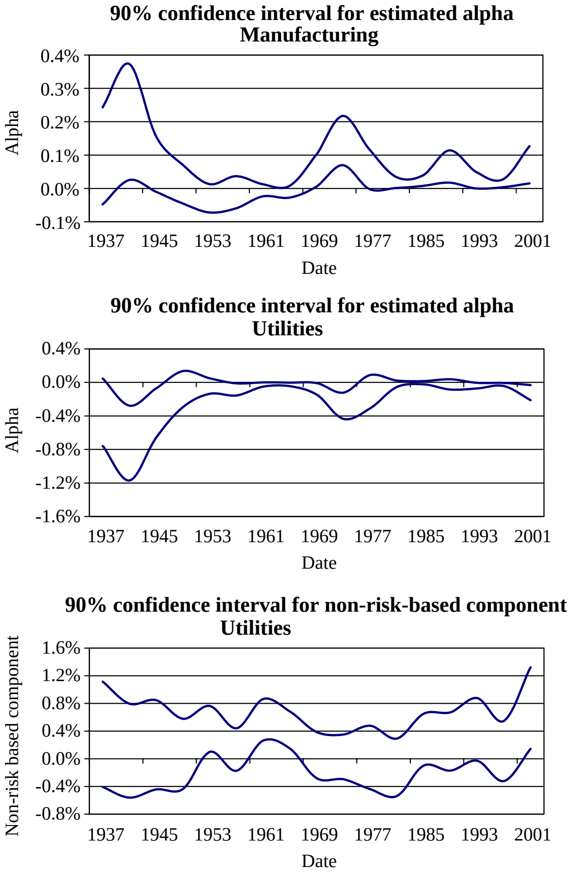

Fig. 4. Confidence intervals for alpha from the OFI model and the non-risk-based component. The figure illustrates the 90% bootstrapped confidence interval for the estimated alpha by the OFI and the non-risk-based component for selected industries. The non-risk-based component is defined as the difference between the sample mean excess return and the expected excess industry return given by the OFI model.

factors.[6] Furthermore, we can look upon CAPM as a sufficient asset pricing model during those periods when the orthogonal portfolio is unimportant, which implies that the market index is an adequate representative of the market portfolio.

3.2. Expected returns and fixed theta

We analyse the effects of fixing theta by comparing the estimates from NFI with endogenous theta to estimates with fixed theta. Following MacKinlay and Pastor (2000), we use two values for theta: one corresponds to a Sharpe ratio of 0.1732 on a monthly

[6] Note that it might be slightly imprecise to call it a “two-factor model” since the orthogonal portfolio might represent several priced factors.

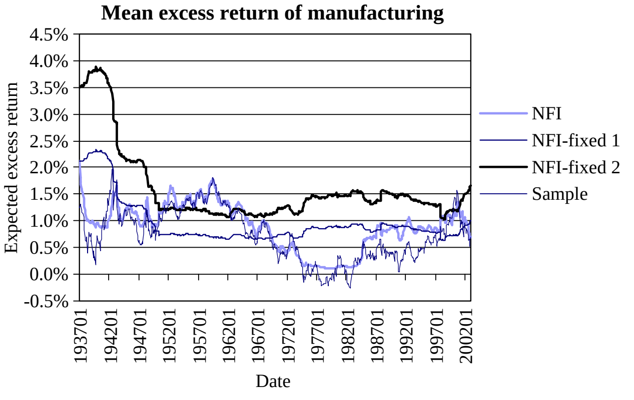

Fig. 5. The stability of the estimated expected industry excess returns. The figure compares the sample mean excess return for the manufacturing industry with the expected excess returns estimated by the model with no observed factor (NFI) for fixed and endogenous theta. We use two values for theta: one corresponds to a Sharpe ratio of 0.1732 on a monthly basis (NFI-fixed 1) and the second is 0.2887 (NFI-fixed 2). All the estimations are based on 10-year moving windows.

basis, which is equal to its average value in our original specification, and the second is 0.2887. As expected, the results show that the level of the estimates is sensitive to the choice of theta (see Fig. 5).[7] MacKinlay and Pastor (2000) state that the estimates of the no-factor model (NFI) are much less variable than the Sample estimates. However, this is not fully supported by our findings with endogenous theta (see Fig. 5), which show that there is a substantial time variation in the estimates of expected excess return from NFI as well as from Sample. For fixed theta, there is a shift in the level at the beginning of the 1950s, but from then onwards, the estimate of expected return exhibits very little variation.[8]

To get a closer look at what is happening when theta is fixed in NFI, we examine the relation between the estimated variance and the expected return of the factor portfolio (Fig. 6). The estimated variance is the same for NFI and two NFI-fixed models, and we therefore illustrate the variance for one model only. As we see in Fig. 6, the estimated expected excess returns of the factor portfolios differ across models. The estimates from models with fixed theta have an almost perfect positive correlation with the estimated variances, since with fixed theta a high estimated variance is necessarily connected with high estimates of expected return in order to keep the Sharpe ratio constant. Hence, in periods with high realised volatility, independently of whether the realised return is high or low, the estimated expected future excess returns from the model with fixed theta will be high. This pattern is particularly evident for estimates based on periods with high realised

[7] For the sake of space, we have only illustrated the result for manufacturing industry, but the same pattern is valid for the other industries.

[8] The average standard deviations in estimated industry returns after 1950 are 0.44% for Sample, 0.37% for NFI and 0.13 for NFI-fixed 1.

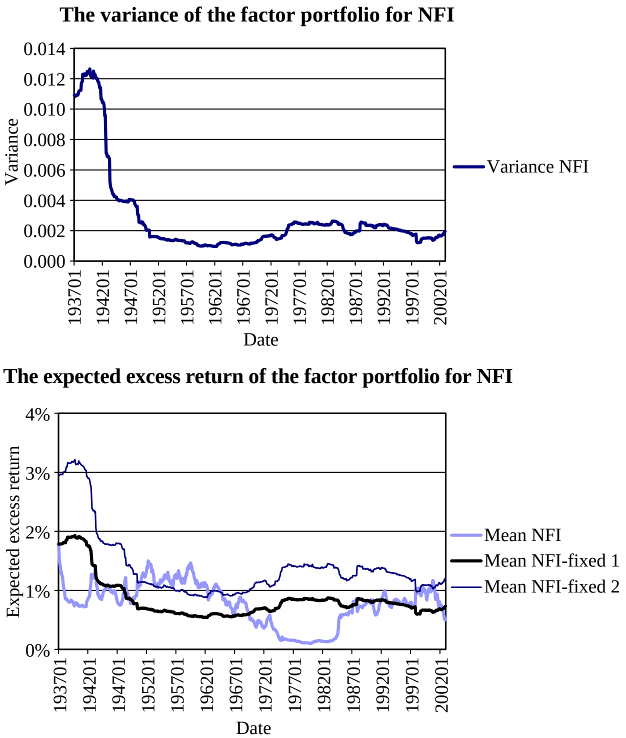

Fig. 6. Optimal portfolio of the NFI model with and without fixed theta. The figure compares the expected excess return and the variance of the optimal portfolio based on the model with no observed factor (NFI) for fixed and endogenous theta. We use two values for theta: one corresponds to a Sharpe ratio of 0.1732 on a monthly basis (NFI-fixed 1) and the second is 0.2887 (NFI-fixed 2). The first diagram plots the estimated variances for NFI with endogenous theta over time and the second diagram plots the estimated expected excess returns for NFI both with fixed and endogenous theta. All the estimations are based on 10-year moving windows.

volatility, which can be seen in Fig. 5 for the expected excess return of manufacturing (compare Sample with the two NFI-fixed models).[9]

The estimation of the variance is not sensitive to our choice of model, and therefore, the lower the value of the fixed theta, the higher the estimated expected return. Hypothetically, if the Sharpe ratio of the optimal portfolio is known a priori, then it might be beneficial to use a fixed theta for estimation, since based on the sample data the covariance matrix may be estimated more correctly than the expected return. In this case, it might be better to rely entirely on the realised covariance matrix and ignore all the information in the realised mean return. However, the results of the models based on fixed theta are not convincing due to the ad hoc character of the choice of theta and the ignorance regarding its possible time variation.

[9] For Sample in Fig. 4, the estimate is the average of the realised return of the manufacturing industry in a preceding 10-year period.

To analyse the effect of fixing theta for the one-factor restricted model, we focus on the estimated alphas, i.e., the part of the industry mean excess return that is explained by the optimal orthogonal portfolio. OFI with an endogenous theta is compared to a model where theta is fixed (OFI-fixed). We follow MacKinlay and Pastor (2000) and use two alternative values for theta in the fixed models, corresponding to a monthly Sharpe ratio of 0.0289 and 0.0577, respectively, with the former value approximately equal to the average value of the estimated thetas from OFI. We see from Fig. 7 that OFI-fixed with a low Sharpe ratio is more stable than the other models. However, our results with an endogenous theta underline the time-varying nature of the Sharpe ratio. As could be anticipated, the level of the estimated alpha with fixed theta is sensitive to the chosen level of theta. Since the true theta is unknown a priori, the arbitrary choice of the level of fixed theta leads to ambiguity regarding to the fraction of the expected return that is related to the missing risk factors, i.e., the estimated alpha.

3.3. Out-of-sample Sharpe ratios

Finally, we examine the out-of-sample results for portfolios constructed according to the weights of the MSR portfolios from different models.

There are years when the expected excess return of the minimum variance portfolio is negative, which leads to a tangency portfolio with a negative mean return. For these cases, we follow MacKinlay and Pastor (2000) and multiply all means by -1, which implies that the weights are not changed since the frontier is simply mirrored in the abscissa. The motivation for this is that firms with the worst performance in down-markets have high betas, and they will consequently perform best in up markets. We use this method for all the models.

There is a strong case for an outlier analysis of the out-of-sample portfolio returns that are due to extreme portfolio weights. The occurrences of extreme weights are closely connected to periods when the expected excess return of the MVP portfolio is close to zero. This is a problem that is inherent in the MV methodology, since in the case of a MVP

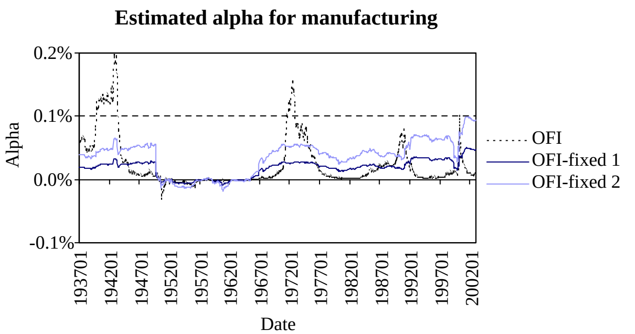

Fig. 7. Alpha estimated by the OFI model with and without fixing theta. The figure compares the alpha for the manufacturing industry given by the restricted model for one observed factor (OFI) with and without fixed theta. The alpha from the restricted model OFI is equivalent to the part of the industry mean excess return that is explained by optimal orthogonal portfolio. All the estimates are based on 10-year moving windows.

Table 2
Out-of-sample Sharpe ratios of the optimal portfolios

	1937–2002			1937–1958	1959–1980	1981–2002
	Mean	Standard deviation	SR	SR	SR	SR
EW	0.006	0.044	0.142**	0.194**	0.086	0.141**
Sample	0.014	0.376	0.038	0.117	0.019	0.087
Sample-I	0.008	0.065	0.120**	0.194**	0.051	0.137*
OFE	0.006	0.047	0.133**	0.194**	0.082	0.118
NFI	0.006	0.046	0.135**	0.190**	0.080	0.128*
OFI	0.012	0.217	0.056	0.075	0.033	0.124*
OFI-fixed 1	0.005	0.081	0.063	0.064	0.042	0.096
OFI-fixed 2	0.000	0.125	0.000	−0.017	−0.014	0.039

The table shows the out-of-sample statistics and Sharpe ratios, SR, of portfolios constructed from the different models for the entire sample, 1937–2002. The Sharpe ratios are also reported for three sub-samples. The portfolios are updated each month. The Sharpe ratios with one asterisk are significant at the 5% level, and those with two asterisks are significant at the 1% level.
The different portfolios are equally weighted (EW), sample (Sample), sample with an identity covariance matrix (Sample-I), one-factor exact (OFE), no-factor identity (NFI), one-factor identity (OFI) and one-factor identity with fixed thetas (OFI-fixed 1 and OFI-fix 2). We use two values for theta: one corresponds to a Sharpe ratio of 0.1732 on a monthly basis (OFI-fixed 1) and the second is 0.2887 (OFI-fixed 2).

portfolio with an expected excess return exactly equal to zero, there is no tangency portfolio. However, in an empirical work, the estimated expected excess return is never exactly zero and the tangency portfolio is therefore very extreme with respect to expected return, standard deviation and weights. We have chosen to substitute the outlier portfolios corresponding to the most extreme weights from each model by holding an equally weighted portfolio, which is a strategy that any investor can imitate ex ante.[10] The affected models are Sample, OFI and OFI-fixed. There are two outliers for Sample, one outlier for OFI and four outliers for OFI-fixed. This problem cannot exist for the OFE since the market index is identical to the tangency portfolio, and for NFI, the weights depend only on the proportions between estimated returns (see Eq. (11)). In fact, the weights for the entire sample and across all the industries vary between 2.5% and 16% for NFI and between 1% and 36% for OFE.

For each month, from January 1937 to December 2002, we calculate the portfolio returns for each model. To compare the out-of-sample results, we use the returns for the entire period as well as for three subperiods to estimate the mean, standard deviation and Sharpe ratio for each model. Note that NFI with fixed thetas are excluded from this race, since fixing theta has no effect on the relative relations between expected returns. Therefore, there are no implications for the weights of the tangency portfolio and the out-of-sample result is identical to NFI.

Table 2 presents the statistics of the out-of-sample portfolio returns and standard deviations, and the Sharpe ratios with their test statistics. Four out of eight models have Sharpe ratios that are significantly different from zero: EW, Sample-I, OFE and NFI. The

[10] MacKinlay and Pastor (2000) place restrictions on the weights, but this method will not separate portfolio weights that are truly optimal from those weights that are driven by the expected return of the MVP portfolio being close to zero.

results are essentially the same for the subperiods, except that no model produces a significant Sharpe ratio for the middle period that includes the bad years during the 1970s. Sample and OFI have much higher means and standard deviations than the other models, but their Sharpe ratios are not significant. However, OFI is significant for the last subperiod, which is a period when the orthogonal factor is not important and the result is therefore close to OFE and NFI.

The no-factor model NFI is significantly better than the one-factor model OFI, which is in accordance with the results in MacKinlay and Pastor (2000).[11] Their explanation is that the unobserved portfolio is a better representation of the "market portfolio" than the value-weighted portfolio used in OFI. However, our results show that OFE, with the value-weighted market portfolio as the sole factor, performs significantly better than OFI (except for the last subperiod) and OFI-fixed. Hence, it is more likely that in the one-factor model, the orthogonal portfolio represents a non-permanent component.

OFI-fixed with a low prespecified theta significantly underperforms OFI-fixed with a high theta, but is not significantly below OFI with an endogenous theta. The reason is that the lower the theta (i.e., the higher the Sharpe ratio), the more important the optimal orthogonal portfolio and the larger the effects on the portfolio weights. As discussed above, this factor is not permanent, and thus, the out-of-sample results of the portfolio based on this factor will be poor. In fact, from the viewpoint of the out-of-sample portfolio performance, the best choice for fixed theta would be infinity (a zero Sharpe ratio for the orthogonal factor), where OFI converges to OFE.

In the last subperiod when the latent factor is not important, the Sharpe ratio for OFI with endogenous theta is significant and converges to the OFE. In contrast, OFI-fixed, particularly the one with a low theta, incorrectly amplifies the importance of the orthogonal factor in the last subperiod, which results in a relatively poor performance.

NFI has a similar out-of-sample performance as OFE, despite the advantage of NFI in the estimation of the expected industry returns. The explanation is straightforward: since these two models result in a similar ranking of the expected returns (due to similarity in relative betas[12]) the weights of the corresponding MSR portfolio will be approximately the same.

The Sample-I is significantly better than the Sample for all periods, which supports the result of MacKinlay and Pastor (2000). In contrast to their findings, our restricted no-factor model NFI is better than Sample-I for the whole period, but the difference is not significant.

The EW strategy has the highest Sharpe ratio and is significantly better than nearly all the other models. This may be explained by the fact that EW is closely related to the family of the restricted models, but with the additional restriction that the mean is proportional to the residual risk (see Section 2.4).

[11] For the inference we use a 10% level. The results of the significance test for all the pair-wise comparisons are available on request. Due to the low power of this test, the results should be interpreted with great caution (see Jobson and Korkie, 1981).

[12] The cross-sectional correlation between industry betas against these two factors is almost equal to one for each month.

Overall, models based on one factor (NFI and OFE) and the related models (Sample-I and EW) exhibit very minor differences over all the periods and outperform the model without restriction (Sample), which is the same result as MacKinlay and Pastor (2000). Hence, there is a clear benefit in using the identity covariance matrix (NFI and Sample-I) in portfolio selection.

4. Conclusion

In this paper, we have used the orthogonal portfolio approach to examine the importance of potential risk factors missing from the CAPM in a sample of US industrial portfolios from 1927 to 2002. We generalised the empirical investigation of MacKinlay and Pastor (2000) by endogenously determining the Sharpe ratio of the orthogonal portfolio as well as letting idiosyncratic risk vary across the test assets.

From an asset pricing perspective, the orthogonal portfolio is useful for detecting missing priced factors. In the no-factor model (NFI), the orthogonal portfolio always has a positive risk premium, which is a necessary characteristic of the true "market portfolio" if CAPM holds. In the one-factor model (OFI), the market index is augmented by an orthogonal portfolio representing a factor that is important for the estimation of expected return during some periods—the distress periods of the 1930s and the 1970s, which were characterised by high volatility. During these periods, the importance of the orthogonal portfolio increases relative to the market index, and it makes a significant contribution to the estimated industry risk premiums. Finally, the non-risk-based component is seldom significant after considering the effect of the latent priced factors. Thus, most of the time we have an exact two-factor pricing model with the market index and the orthogonal portfolio as the factors, and when the latter is unimportant CAPM is a sufficient asset pricing model.

We find that fixing the Sharpe ratio of the orthogonal portfolio a priori leads to less plausible estimates of the expected returns. In particular, for the no-factor model, a fixed Sharpe ratio introduces by construction a proportional relationship between the mean and the variance of the tangency portfolio. As a result, the estimated expected returns are not connected to the realised mean returns of the estimation window, and the estimates differ substantially from the estimates from all other models. For the case of the one-factor model, fixing the Sharpe ratio of the orthogonal portfolio obscures the time variation in the orthogonal portfolio and its relative importance.

We also examined models that allow for differences in the idiosyncratic risk. However, our results show that there are almost no differences between these models and the models in which the idiosyncratic risk is identical over the test assets. This result might be due to the fact that the test assets belong to the same asset class, which corroborates the hypothesis of MacKinlay and Pastor (2000) that the idiosyncratic risk is only important for different asset classes.

The out-of-sample Sharpe ratios show that all restricted models outperform Sample. Thus, it proves the general importance of exploiting the relationship between return and risk. A general comparison shows that the restricted no-factor model (NFI) is better than the restricted one-factor model (OFI). These results are in accordance with the results in

MacKinlay and Pastor (2000). They explain the latter result by saying that the unobserved portfolio is a better representation of the "market portfolio" than the value-weighted portfolio used in the one-factor models. Our results show that a more likely explanation is that the orthogonal portfolio in the one-factor model picks up a component in the data that is not stable over time. Hence, the instability of the orthogonal portfolio implies that the use of an endogenous theta is not essential out-of-sample. In addition, our results demonstrate the advantage of using the identity covariance matrix in portfolio allocation, which is in accordance with the findings of MacKinlay and Pastor (2000).

We conclude that the orthogonal portfolio approach is particularly useful as a tool for estimating expected return in sample, since it reveals the existence of time-varying priced risk factors missing from CAPM. However, we find that simpler models are equally as good for achieving high out-of-sample Sharpe ratios as the more complicated models.

Acknowledgements

We are very grateful for the valuable comments by an anonymous referee. The financial support from Jan Wallander and Tom Hedelius Foundation is highly appreciated.

References

Connor, G., Korajczyk, R.A., 1986. Performance measurement within the arbitrage pricing theory: a new framework for analysis. Journal of Financial Economics 15, 373–394.

Corana, A., Marchesi, M., Martini, C., Ridella, S., 1987. Minimizing multimodal functions of continous variables with the 'simulated annealing algorithm'. ACM Transactions on Mathematical Software 13, 262–280.

Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3–56.

Fama, E.F., French, K.R., 1996. Multifactor explanations of asset pricing anomalies. Journal of Finance 51, 55–84.

Goffe, W.L., Ferrier, G.D., Rogers, J., 1994. Global optimization of statistical functions with simulated annealing. Journal of Econometrics 60, 65–99.

Jobson, J.D., Korkie, B.M., 1981. Performance hypothesis testing with the Sharpe and Treynor measures. Journal of Finance 36, 889–908.

Lehman, B., Modest, D., 1988. The empirical foundations of the arbitrage pricing theory. Journal of Financial Economics 21, 213–254.

MacKinlay, A.C., 1995. Multifactor models do not explain deviations from the CAPM. Journal of Financial Economics 38, 3–28.

MacKinlay, A.C., Pastor, L., 2000. Asset pricing models: implications for expected returns and portfolio selection. Review of Financial Studies 13, 883–916.

Merton, R.C., 1973. An intertemporal capital asset pricing model. Econometrica 41, 867–887.

Merton, R.C., 1980. On estimating the expected return on the market: an explanatory investigation. Journal of Financial Economics 8, 323–361.

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

Ross, S., 1976. The arbitrage theory of capital asset pricing. Journal of Economic Theory 13, 41–360.