Performance and Style Shifts in the Hedge Fund Industry

Ramin Baghai-Wadji [*] Stefan Klocker [†]

January 15, 2007

Abstract

We suggest a performance attribution model which is adapted to the requirements of hedge funds by a set of style-consistent benchmark indices obtained from a neural network based clustering procedure. We compare our approach with alternative models and analyze whether fund of hedge funds managers create added value. Employing factor loading-changes as a proxy for style changes, we show that style shifts are a characteristic feature of hedge funds and that style-consistency does not generally ameliorate performance. Finally, we demonstrate that while poor performers which change style can expect a subsequent performance improvement, the same does not hold for top performers.

Keywords: Performance Measurement, Hedge Funds, Neural Networks, Style Shifts

JEL classification: C45, G11, G15, G23

[*]Corresponding author. Finance PhD Programme, London Business School, Sussex Place, Regent's Park, London NW1 4SA, UK. E-mail: rbaghaiwadji.phd2006@london.edu. Tel.: +44 (0)20 7000 5050, ext. 3937. Fax: +44 (0)20 7000 7001.

[†]Corporate Finance Group, Department of Finance and Accounting, Vienna University of Economics and Business Administration, Nordbergstraße 15, A-1090 Vienna, Austria. E-mail: stefan.klocker@wu-wien.ac.at

1 Introduction

Models used to assess the performance of mutual funds work reasonably well to evaluate static buy-and-hold long only investments, which are common in the mutual fund universe, but are of limited value for the performance attribution of hedge funds. This is due to the unique features which distinguish hedge funds from ordinary mutual funds: The laxer regulatory framework they operate in (hedge funds are not subject to SEC regulations and the Investment Company Act of 1940) enables them to employ very flexible ("dynamic") trading strategies, including highly leveraged positions, short selling, and the use of derivative instruments. Furthermore, there are no reporting requirements to speak of (all information disclosed by hedge funds to database providers is entirely voluntary), and the fee structure differs from the one in the mutual fund business (for details refer to e.g. Lhabitant, 2002).

In the recent past, considerable effort went into devising performance measurement models specific to the hedge fund universe (see for example Fung and Hsieh, 1997, 2004; Lhabitant, 2001). However, researchers and practitioners are still far from accepting a particular model as the standard.

In this paper, we present a new approach for evaluating the performance of hedge funds which is based on a classical multi-factor model in the tradition of Sharpe's (1992) asset class factor model but uses neural network derived benchmark indices. In order to construct these hedge fund indices, we use an unsupervised neural network, the so-called Self-Organizing Map (SOM), to group hedge funds into homogenous categories. Kandel et al. (2004) have shown that the use of group-specific benchmarks, which they call "endogenous" benchmarks, markedly reduces the omitted risk factor problem normally encountered in performance attribution models. The classification of hedge funds is based on their monthly return histories, which can be assumed to reflect their actual investment style, and does not rely on the funds' self-classifications, which may be biased by fund managers' attempts to polish their historical performance (see e.g. Brown and Goetzmann, 1997). This allows us to create benchmark indices, which prove in the empirical section to capture the essence of dynamic trading strategies better than traditional models.

We use the neural network derived benchmark indices as well as several alternative models to evaluate a hedge fund's performance and assess a manager's investment skill via the regression's Alpha. Every performance attribution model implicitly claims to distinguish fund managers with superior investment skills (significantly positive Alpha) from their less talented counterparts. In order to compare the different performance attribution models along these lines, we analyze the evolution of Alphas over time. If a performance attribution model indeed captures the investment ability of a fund's manager this will lead to persistent model Alphas. While our SOM-based approach exhibits persistent Alpha estimates, all other models but one fail to have this desired property. The only model that exhibits similar predictive power for future Alphas also uses hedge fund indices as regressors, but is clearly outperformed by our SOM-based benchmark indices in terms of explanatory power. The finding of persistent Alphas does not only underline the usefulness of the hedge fund index-based performance attribution models, but also shows that investment skill is an important factor in the hedge fund business. We also use our SOM-based model to investigate whether managers of fund of funds are able to provide investors with an excess return through their superior selection of single strategy hedge funds despite the additional layer of fees.

Another goal of this paper is to use the neural network derived benchmark indices to address the question whether or not style shifts are a significant issue in the hedge fund universe, and to elucidate the link between style shifts and the performance of hedge funds. The issues of fund misclassification and style shifts have been discussed in many recent papers pertaining to the empirical and theoretical mutual fund literature.[1] However, these topics have not been addressed in much detail in the literature on hedge funds. In contrast to findings for mutual funds, our research suggests that style changes are at least as wide-spread after good performance as they are after bad performance. We also show that style-consistency does not, on average, lead to better performance. Furthermore, we also demonstrate that funds with bad performance which change style can expect a performance improvement in the period following the strategy change, whereas top performers which alter their strategy are more

[1] See, for example, diBartolomeo and Witkowski (1997); Gallo and Lockwood (1999); Kim et al. (2000); Lynch and Musto (2003). Cumming et al. (2005) discuss style drifts in venture capital.

likely not to benefit from the style shift.

The rest of the paper is organized as follows. In Section 2, we describe the data. Section 3 discusses and compares traditional performance measurement models used in the literature. In Section 4, we show how the neural network benchmark indices are constructed and discuss their out-of-sample performance. In Section 5, we compare our performance measurement model with its closest competitors. The link between fund performance and style shifts is examined in Section 6. Finally, Section 7 summarizes and concludes.

2 Data

Our paper is based on monthly return data from the CISDM (Center for International Securities and Derivatives Markets) hedge fund database, formerly known as the Managed Account Reports, Inc. (MAR) database. CISDM provides a summary of the self-declared investment strategy and style for each fund. This proprietary classification will be used as a starting point for the labeling of our SOM-derived clusters and benchmarks. The original dataset comprises a twelve year time period from May 1992 to April 2004 and contains 5,440 hedge funds. We only include funds with at least 36 monthly return observations in our sample in order to assure a sufficiently high degree of computational stability of the neural network as well as to obtain meaningful estimates in the regression analysis.[2] This eliminates 1,873 funds from our original data set. In a second step we exclude the fund of funds (FOF) category from the benchmark creation in order to ensure that the resulting benchmarks reflect the "pure" trading strategies, which reduces our sample by another 541 funds. Although these 541 funds of hedge funds are excluded from the benchmark creation, we analyze this *FOF Sample* with our performance attribution model in section 5.2. All of the above considered, this leaves us with a total sample of 3,026 single strategy hedge funds. Table 1 summarizes our data sample.

[2] The requirement that a fund must have a sufficiently long return history to be included in the sample can give rise to a so-called "multi-period sampling bias." However, according to Fung and Hsieh (2000), the resulting upward performance bias is negligibly small. Ackermann et al. (1999) even find that for their data sample the "multi-period sampling" requirement actually biases their statistics downwards. On the whole, the impact of a required return history of 36 monthly observations appears to be of limited significance.

Table 1: Hedge fund sample and summary statistics by self-declared investment strategy.

Self-Declared Investment Strategy	**Number of Funds**	**Mean Return** (in %)	**Std.Dev.** (in %)	**Skewness**	**Kurtosis**
Convertible Arbitrage	149	1.00	1.60	0.05	5.60
Distressed Securities	83	1.04	3.14	-0.05	5.55
Equity-Hedge	833	1.12	4.49	0.17	4.65
Emerging Markets	175	1.14	6.55	-0.04	5.35
Fixed Income	91	0.76	1.92	-1.23	8.61
Managed Futures	1180	0.91	5.37	0.46	4.39
Global Macro	116	0.96	4.56	0.17	4.56
Merger Arbitrage	128	0.86	1.82	-0.22	5.45
Market Neutral	46	0.74	2.34	0.03	4.17
Sector Financial	29	1.47	3.88	-0.39	5.89
Sector Healthcare	28	2.23	7.87	1.37	6.94
Short Selling	33	0.46	6.81	0.05	3.88
Sector Technology	61	1.73	11.06	0.38	4.11
Sector Multi Sector	32	1.15	4.78	0.33	4.86
Long Only	24	1.44	7.19	-0.08	4.29
Sector Energy	10	1.94	8.63	0.15	4.06
Sector Real Estate	8	0.98	2.94	0.44	4.38
Fund of Funds	541	0.78	2.17	-0.03	5.13
Total Funds in Sample	3567				

The numbers given correspond to the monthly category medians (e.g. "Mean Return" denotes the category median of the arithmetic means of the time series of the individual hedge funds).

We randomly split this sample into two non-overlapping sub-samples. In the first step we use sub-sample one (*SOM Sample*), which consists of 2,026 hedge funds, as a training set for the Self-Organizing Map and subsequently for the construction of our benchmarks. Next, we compare the explanatory power of our model with traditional performance attribution models using sub-sample two (*Regression Sample*), which consists of the remaining 1,000 single strategy hedge funds.

By construction, our data set does not suffer from survivorship bias, as 1,607 of the 3,026 single strategy funds are "non-surviving" hedge funds, i.e. funds which exhibit at least 36 observations but which have stopped reporting to the data base at some point in time during the period under observation.[3] We also mitigate another problem called either backfilling bias or instant history bias (see Fung and Hsieh, 2000). This bias is caused by the practice of backfilling the historical performance of hedge funds which are newly added to a database. Obviously funds with a good track record are more likely to disclose their historical returns and therefore the backfilling practice causes an upward bias. By requiring at least 36 monthly return observations the effects of the backfilling bias are markedly reduced, as the historical returns backfilled typically contain the last 12-15 months (cf. Fung and Hsieh, 2000).

3 Traditional Models for Measuring the Performance of Hedge Funds

In this section we discuss six traditional performance attribution models. All these models can be interpreted as special cases of the general Sharpe (1992) model, which is an asset class factor model of the following form:

$$R_{i,t} = a_i + \sum_{k=1}^{K} \beta_{i,k} F_{k,t} + \epsilon_{i,t} \tag{1}$$

where F_1 to F_K are the returns of K different asset classes and $\beta_{i,1}$ to $\beta_{i,K}$ represent the sensitivities of R_i, fund i's returns, to these K asset classes. The sensitivities $\beta_{i,k}$ can be interpreted as the exposure of the analyzed fund (or portfolio) to these asset classes. Models of this type attempt to separate a managed portfolio's returns into a component which can be easily replicated via the returns on standard asset classes and a remaining component, the so-called "Jensen's Alpha," which captures the manager's investment skill (cf. Jensen, 1968).

[3] A fund manager might cease reporting for various reasons, of which the most obvious is poor performance followed by a closure of the hedge fund. Alternatively, it could be argued that since hedge funds by their very nature try to exploit market inefficiencies, the capacity and possible size is limited (see, for example, Agarwal et al., 2003, who document decreasing returns to scale in the hedge fund industry). Therefore, a fund which has reached its capacity limit in terms of investment volume no longer needs to attract additional investors and is likely to cease reporting to the database. While the first scenario implies that poorly performing funds stop reporting, which would introduce a positive survivorship bias, the second scenario, which is more likely to occur in the case of funds performing well, would create exactly the opposite effect. The literature suggests that the first effect dominates and the survivorship bias amounts to 3% annually (see e.g. Fung and Hsieh, 2000).

Using the *Regression Sample* comprising 1,000 single strategy hedge funds, we analyze the applicability of the following six models for performance attribution in the hedge fund universe:

- Models traditionally used in conjunction with mutual funds
  - The Capital Asset Pricing Model (CAPM) independently suggested by Sharpe (1964), Lintner (1965) and Mossin (1966)
  - The three-factor model of Fama and French (1993)
  - The four-factor model of Carhart (1997)
- Models designed for hedge funds
  - The asset class factor model of Fung and Hsieh (1997)
  - The asset-based style factor model of Fung and Hsieh (2004)
  - The hedge fund index multi-factor model suggested by Lhabitant (2001)

The first three models are based on the CAPM and only include equity market specific risk factors. In contrast, the Fung and Hsieh (1997) model, which covers the equity, fixed income, currency and commodity markets, encompasses all major asset classes used by hedge fund managers. Fung and Hsieh (2004) go one step further by including regressors with a nonlinear exposure to standard asset classes. These so-called Asset-Based Style (ABS) factors attempt to explain returns of hedge fund strategies via observed market prices (see also Fung and Hsieh, 2002).[4] The multi-factor model of Lhabitant (2001) uses the strategy sub-indices of the CSFB/Tremont hedge fund index family as "risk factors."

[4] The ABS factors employ the contingent claims methodology pioneered by Glosten and Jagannathan (1994), whose approach has also been used in the context of hedge funds by Agarwal and Naik (2004).

Table 2: Comparison of the explanatory power (adjusted R^2) of various performance attribution models.

	CAPM	Fama/French	Carhart	FH 1997	FH 2004	Lhabitant	SOM
Minimum	-0.03	-0.08	-0.10	-0.17	-0.39	-0.52	-0.65
25% Quantile	0.00	0.01	0.01	0.05	0.06	0.12	0.15
Median	0.05	0.08	0.09	0.15	0.20	0.29	0.36
75% Quantile	0.19	0.28	0.31	0.30	0.33	0.47	0.55
Maximum	0.89	0.92	0.92	0.95	0.93	0.92	0.93

The table shows the regression performance of the CAPM, Fama and French (1993), Carhart (1997), Fung and Hsieh (1997), Fung and Hsieh (2004) and Lhabitant (2001) models, and the results for the SOM derived benchmarks. The numbers given are the adjusted R^2 values. The first and last rows show the adjusted R^2 of the worst and best regression results, respectively.

Table 2 presents an overview of the explanatory power achieved by the different performance attribution models for the *Regression Sample*. As the analyzed models do not have the same number of explanatory variables we use the adjusted R^2 statistic to compare them. The median adjusted R^2 of the funds analyzed shows that there are considerable differences between the models studied. In the case of the four-factor Carhart (1997) model, which is still the best performing of the three CAPM-based models, more than 50% of all funds display an adjusted R^2 below 0.09. As expected, the three models specifically designed for hedge funds do a better job with a median adjusted R^2 of 0.15 for the Fung and Hsieh (1997) asset class factor model, an adjusted R^2 of 0.20 for the Fung and Hsieh (2004) asset-based style factor model and an adjusted R^2 of 0.29 for the hedge fund index model of Lhabitant (2001). Overall it can be said that the CAPM-based models are inadequate for performance attribution in the hedge fund universe. However, surprisingly, even the asset class factor model of Fung and Hsieh (1997) and the asset-based style factor model of Fung and Hsieh (2004) do not perform exceptionally well.

This comparison of adjusted R^2s clearly shows that all asset class factor models using the Sharpe (1992) setup in combination with risk factors which feature a linear exposure to the standard asset classes, namely the CAPM, as well as the models suggested by Fama and French (1993), Carhart (1997), and Fung and Hsieh (1997), display a low explanatory power for single strategy hedge funds. This inadequacy is due to the unique features of hedge funds, like dynamic trading strategies and the extensive use of leverage. The consequences of dynamic trading strategies are alternating long and short positions within a certain asset class, often resulting in a regression coefficient for that asset class close to zero. In Sharpe's (1992) framework, this would imply that a particular hedge fund that invests its entire capital e.g. in the bond market and uses a dynamic trading strategy could be characterized as having no exposure to the bond market at all, if it alternates long and short positions within that asset class (cf. Fung and Hsieh, 1998).

In principle, there are two distinct ways of dealing with the dynamic nature of hedge fund investment strategies for performance measurement purposes. The first way is to use standard asset classes as regressors and to allow for time-varying factor loadings. This can be implemented using a linear regression model with rolling windows (see e.g. McGuire et al., 2005) or various types of generalized autoregressive conditional heteroscedasticity (GARCH) models, stochastic volatility (SV) models or Kalman filter-based approaches (for a recent comparison of alternative modeling techniques as applied to time-varying CAPM betas see Mergner and Bulla, 2005).

The second way of capturing the dynamic components of hedge fund investment strategies is to include regressors with non-linear exposures to standard asset classes, which then serve as proxies for the dynamic trading strategies. Using regressors that already incorporate the dynamic features of hedge fund investment styles opens up the possibility of using a simple linear regression framework for hedge fund performance measurement while still allowing for non-linearities in the analysis via the "dynamics" embedded in the regressors.

A first approach for identifying regressors which feature a non-linear exposure to the underlying asset classes is to use specifically constructed indices for each hedge fund style category. As mentioned above, Fung and Hsieh (2002) call this type of indices "Asset-Based Style (ABS) factors", which account for dynamic trading strategies by using primarily contingent claims. The Fung and Hsieh (2004) model analyzed above uses the ABS factor approach. Only a limited number of such ABS factors have been identified so far (see for example Fung and Hsieh, 2001, 2002; Mitchell and Pulvino, 2001). Furthermore, we conjecture that many of the discretionary investment strategies with the formidable set of investment possibilities open to hedge funds are hard if not impossible to explain in the way described by Fung and Hsieh (2002, 2004). The rather low average adjusted R^2 of the Asset-Based Style factor model indicates that the seven risk factors identified by Fung and Hsieh (2004) might not be sufficient to explain the risk in the returns of individual hedge funds and it is likely that some risk factors important for hedge funds are omitted. Fung and Hsieh (2004) point out themselves that for individual funds the construction of additional style-specific risk factors is unavoidable.

Another approach that remains in the linear regression framework while at the same time accounts for the dynamic nature of hedge fund trading strategies is to use hedge fund indices. This approach can mitigate the omitted risk factor problem, because hedge fund indices implicitly contain all the risk factors that the constituting funds are exposed to. While hedge fund indices cannot explicitly link hedge fund returns to specific risk factors, the indices can nonetheless act as useful proxies for those (unspecified) risk factors for all practical purposes such as performance attribution. Lhabitant (2001) for example suggests to use the CSFB/Tremont hedge fund index family. However, this particular approach suffers from several major drawbacks. First of all, the number of strategy categories employed (i.e. the number of indices used as regressors) is somewhat arbitrary and hence not optimal. Secondly, data vendors constructing hedge fund indices typically rely on the self-proclaimed style labels given to the funds by their managers.[5] Inevitably, the prevailing misclassification and style drifts in the hedge fund business markedly reduce the usefulness of standard hedge fund indices as proxies for the dynamic trading strategies of particular hedge fund investment styles.[6] It is also well documented that there is considerable heterogeneity within self-declared hedge fund styles, which can be readily seen by the low correlations between the monthly returns of hedge

[5] A notable exception are the hedge fund indices by Zurich Capital Markets which use both qualitative criteria as well as statistical clustering procedures for the construction of the indices. However, these indices only go back to 1998 and due to the way they are constructed they exhibit a severe selection bias (the indices are based on only 60 funds and the constituting funds are hand picked according to a number of constraints, such as reporting requirements, minimum years in existence and assets under management). For details see Amenc et al. (2003).

[6] For an analysis of misclassification and style drifts in the hedge fund universe, see Bares et al. (2001), Amenc and Martellini (2003), and Baghai-Wadji et al. (Forthcoming).

funds within a given style category.[7] Thirdly, when analyzing persistence in performance of hedge funds over time, the use of common hedge fund indices has the disadvantage that these indices are re-balanced regularly and that they do not represent the same composition during a given period of interest. Therefore, persistence estimates would likely be erroneous, as pointed out by Gregoriou et al. (2005). In the following section, we present an approach for constructing hedge fund benchmark indices which is able to take account of these concerns.

4 Neural Network-based Hedge Fund Indices

4.1 SOM Methodology

The Self-Organizing Map (SOM) is a widely-used tool for grouping and visualizing high-dimensional data; it is a single-layered unsupervised neural network which does not require any human intervention during the training process (Kohonen, 2000). The SOM maps high-dimensional input data into a lower dimensional output space (usually two-dimensional, hence the term "map") while preserving the inherent structure of the original data. Therefore, if two input vectors are similar, they will end up in the vicinity of each other on the map. In the present paper, the return time series of each hedge fund represents an input vector, the dimension of which is given by the number of monthly return observations. After the completion of the training process, hedge funds exhibiting similar return characteristics will be represented as homogenous clusters on the two-dimensional map.

The SOM classification proceeds in the following steps:

- **Step 1:** First the number of nodes, which are located on a lattice, has to be specified (in our case e.g. a quadratic lattice of 20x20 nodes). Each node i is represented by its reference vector $m_i(t)$. The dimension of the reference vector depends on the number of features characterizing the input. In our case the input is the set of individual hedge funds used for training the map and the features are the 144 monthly returns covering

[7] For example, Edwards and Caglayan (2001) document that the medians of the pairwise correlation coefficients range from 0.09 for the market-neutral to 0.53 for the short-selling category.

the period May 1992 to April 2004. The initial values of the reference vectors are randomized.

- **Step 2:** Next, the first input vector $x(1)$, which is chosen randomly (sampling without replacement), is presented to the map. The similarity between the input vector and each reference vector is computed. We use Euclidian distance as the measure of similarity. In general, let $x(t)$ denote the input vector randomly selected in training cycle t. The Euclidean distance between input vector $x(t)$ and each reference vector $m_i(t)$ is defined as $\sqrt{\sum_j (x_j(t) - m_{ij}(t))^2}$, where $x_j(t)$ denotes the j^{th} element of the input vector selected in training cycle t, i.e. the return of that particular fund in month j. The winning node is defined as the node with the smallest Euclidian distance with respect to a given input vector.

- **Step 3:** Once the winning node is determined learning starts. The winning node as well as its neighboring nodes within a given radius are updated. The radius starts with the user defined initial value and constantly decreases in the course of the training. In the updating process the reference vectors $m_i(t)$ of the nodes in the vicinity of the winning node (for example all nodes within a radius of 3 nodes) are adjusted towards the input vector $x(t)$. The adjustment is determined by the following formula:

$$m_i(t+1) = m_i(t) + \alpha(t) \cdot [x(t) - m_i(t)] \tag{2}$$

  where $\alpha(t)$ is the learning rate factor, with $0 < \alpha(t) < 1$. If $\alpha(t)$ were one, the adjusted nodes would be immediately set equal to the input vector.

- **Step 4:** The next input vector is chosen randomly and the process is repeated starting from Step 2 until a user defined number of training cycles is completed. In the case of 10,000 training cycles each of the approximately 2,000 *SOM Sample* hedge funds is presented about 5 times to the map. During the training process, the updating radius as well as the learning rate factor $\alpha(t)$ are linearly decreased after each training cycle (consisting of step 2 and 3).

- **Step 5:** Finally each fund is assigned to the node that features the smallest Euclidean distance between the fund's return vector and the node's reference vector. If two input vectors are similar in terms of the distance measure employed, they will ultimately be assigned either to the same node or at least to neighboring nodes on the map.

The SOM algorithm (Kohonen et al., 1995) requires the following user defined parameters and inputs for the training:

- The input vectors used for the training
- The dimension of the lattice, i.e. the number of nodes
- The training is divided into a rough and the fine tuning phase, which require each:
  - The number of training cycles
  - A starting value for the learning rate factor
  - A starting value for the radius defining the neighborhood

As has been shown by Cottrell et al. (2001) the SOM algorithm is not very sensitive to the chosen initial parameter values.

Since an input vector encompasses the whole return series of a particular hedge fund and therefore fully characterizes the fund's return history, we refrain from explicitly including return *moments* (such as mean and variance) in the input vectors used for training the SOM. As we use net returns for training the map the fee structure of a particular fund is implicitly contained in the input vector. Furthermore, we do not use any other fund characteristics either. diBartolomeo and Witkowski (1997) have pointed out that return-based classification methods have several advantages over methods which rely on other characteristics to classify funds. Returns are, in the end, what investors are interested in and they fully characterize the trading strategy of a particular hedge fund manager. A special feature of the SOM algorithm is its ability to cope with missing values in the input vectors (see Kohonen et al., 1995). This is essential for our application as few of the funds in our sample have a return time series covering the entire 12 year period from 1992 to 2004.

Finally the following considerations determined our choice of a 20x20 map dimension. As the correct number of hedge fund strategy clusters is not known a priori, we cannot set the number of nodes equal to the number of clusters in the data, which is a prerequisite for using the SOM to directly divide the set of hedge funds into a pre-specified number of clusters. This information would have also allowed us to use one of the traditional clustering algorithms. Therefore we have to rely on *clustering via visualization*, which is the typical application of Self-Organizing Maps in the field of finance (cf. Deboeck and Kohonen, 1998). Clustering via visualization requires a considerably higher number of nodes than the assumed number of clusters to avoid negative effects from the discretization of the SOM's output space (cf. Flexer, 2001). On the other hand, too many nodes will result in large empty sections on the map with several nodes remaining unused and Euclidean distances losing their discriminatory power. Given the number of input vectors in our sample (2,026 funds) and based on experiments with maps of different sizes, we chose to use 400 nodes.

4.2 Construction of the SOM-based strategy indices

Our benchmark construction consists of two steps. In the first stage we use the SOM algorithm in order to group hedge funds according to their innate return characteristics rather than rely on their self-proclaimed investment styles. In the second stage we use all the funds establishing a strategy cluster detected by the SOM to construct our benchmark indices.

The clustering procedure starts by training the Self-Organizing Map with the *SOM Sample* comprising 2,026 funds. On the resulting map the strategy clusters are determined with the following mechanical algorithm. Following Fung and Hsieh (1997) and Brown and Goetzmann (2003), the labeling of the distinct style categories is done according to the preponderance of managers of a given self-declared style in each cluster. We label a node with a specific style if funds having this self-declared trading style constitute the largest individual group and account for at least 40% of all funds assigned to that node. When the entire map is labeled according to this rule adjacent nodes having the same style label automatically evolve as clusters. This labeling / clustering procedure eliminates the subjectivity typically connected with clustering via visualization. Applied to our map this procedure detects eleven strategy clusters: Convertible arbitrage (CA), distressed securities (DS), emerging markets (EM), fixed income (FI), currency futures (FUC), diversified futures (FUD), merger arbitrage (MA), sector financial (SF), sector healthcare (SH), short selling (SS) and sector technology (ST).[8] An advantage of this clustering procedure is that we do not need to specify the number of clusters in advance.

The strategy clusters identified in the first step are then used to construct corresponding benchmark indices which serve as a proxy for the dynamic trading strategy. This is achieved by forming an equally weighted portfolio of all funds constituting a given trading strategy cluster. Each cluster also contains funds with self-declared styles which do not match the strategy label of the cluster. These funds are not excluded from the benchmark creation, because according to our SOM classification their returns correspond to the strategy that gives the cluster its name; This might be either due to two different trading style labels actually describing the same investment strategy, or can be caused by intentional or unintentional misdeclaration.

These points are in fact important for distinguishing our model from other "peer-group" approaches and eliminate most drawbacks typically connected with this type of approach; see Fung and Hsieh (2002) and the discussion of Lhabitant's (2001) model in section 3. Instead of relying on an ad hoc classification of hedge funds based on self-declared strategies, we identify and incorporate into our model specification only those style groups which also produce a sufficiently discernable and characteristic return pattern. In addition, our indices are also, to the largest possible extent, "purified" from data biases that standard hedge fund indices inherit from the databases underlying their construction (see section 2).

The cross-correlations between the benchmark indices are rather low, averaging around 0.26 in absolute terms. The correlation coefficient between the merger arbitrage and distressed securities benchmark indices, amounting to 0.62, is one of the highest. This can be explained

[8] Note that for equity hedge funds this procedure always resulted in the identification of multiple smaller scattered clusters, which were not connected to each other. For this reason it was not possible to locate a single homogenous equity hedge cluster on the map. Given the generally low correlations between funds in this self-declared category (see Edwards and Caglayan, 2001; Schneeweis et al., 2004), this result is hardly surprising.

by the digital nature of the underlying business (deal closure or not and bankruptcy or not) and by the fact that companies that are being taken over are often in a state of financial "distress." Overall the low correlations support the conjecture that the SOM derived strategy benchmarks represent unique return patterns and furthermore avoid multi-collinearity problems in the regression analysis performed in the following section.

4.3 Regression results for the SOM-based strategy indices

Using the benchmark indices, we specify the following multi-factor model:

$$\begin{aligned} R_{i,t} - R_{f,t} = \alpha_i &+ \beta_{i,1}(FUC_t - R_{f,t}) + \beta_{i,2}(FUD_t - R_{f,t}) + \beta_{i,3}(ST_t - R_{f,t}) \\ &+ \beta_{i,4}(SH_t - R_{f,t}) + \beta_{i,5}(SS_t - R_{f,t}) + \beta_{i,6}(DS_t - R_{f,t}) + \beta_{i,7}(MA_t - R_{f,t}) \\ &+ \beta_{i,8}(FI_t - R_{f,t}) + \beta_{i,9}(CA_t - R_{f,t}) + \beta_{i,10}(EM_t - R_{f,t}) + \beta_{i,11}(SF_t - R_{f,t}) + \varepsilon_{i,t} \end{aligned} \tag{3}$$

where $R_{i,t}$ is hedge fund i's return in month t, $R_{f,t}$ is the risk-free interest rate (we use the return on the one month Treasury Bill as a proxy), and the right-hand side explanatory variables are the returns of the SOM-based benchmarks (Convertible arbitrage (CA), distressed securities (DS), emerging markets (EM), fixed income (FI), currency futures (FUC), diversified futures (FUD), merger arbitrage (MA), sector financial (SF), sector healthcare (SH), short selling (SS), and sector technology (ST)).

The last column of table 2 gives the adjusted R^2 statistic that the SOM-based benchmarks achieve for the *Regression Sample* of 1,000 funds. As can be seen from the table, our benchmarks produce very satisfactory regression results, in particular when compared with the models discussed in section 3 (see table 2), which are commonly used in the literature and real-world applications (i.e. CAPM, Fama and French, 1993; Carhart, 1997; Fung and Hsieh, 1997, 2004; Lhabitant, 2001). The adjusted R^2 is above 0.35 in more than half of the regressions (compared to 0.29 for the CSFB/Tremont indices and 0.20 for the Fung and Hsieh, 2004, model) and above 0.55 in more than 25% of the regressions (compared to 0.47 for the

CSFB/Tremont indices and 0.33 for the Fung and Hsieh, 2004, model).

Since our benchmarks are exclusively constructed with funds which are not in the *Regression Sample*, whereas the CSFB/Tremont indices used by Lhabitant (2001) are very likely to contain several of the funds analyzed, it is noteworthy that our benchmarks achieve higher R^2s than the CSFB/Tremont hedge fund indices. One particular difference between the CSFB/Tremont hedge fund indices and our benchmarks is the fact that in the creation of our benchmarks all funds receive the same weight whereas the CSFB/Tremont indices are value-weighted. However, the most likely reason for the superior performance of the SOM-based indices is that they have been constructed to be style consistent (see section 4.2) and that the number of regressors included in our model specification is not arbitrary but based on the return characteristics of the hedge fund universe.

5 Performance analysis with neural network derived benchmarks

The main question of interest in performance measurement is whether or not an individual hedge fund manager can create additional value when the performance of his fund is compared with suitably chosen benchmarks. Strictly speaking, since we use hedge fund indices as regressors, the Alpha in equation (3) is not exactly equivalent to "Jensen's Alpha," but can be interpreted as the excess investment ability of an individual manager relative to his peer group of hedge fund managers (see also Kandel et al., 2004). Looking at the Alphas of individual funds can therefore elucidate whether one particular fund manager produces abnormal returns relative to other hedge funds.

5.1 Persistence of Alphas

The idea behind all performance attribution models is that the calculated Alpha reflects the investment skill of the manager. If that is the case and skill in fact matters, we would expect to find persistence in the calculated Alphas over time. It follows that the persistence of Alphas is a good criterion to assess the quality of different performance attribution models. Using this procedure we compare our SOM-based model with the three hedge fund performance attribution models introduced in section 3, as well as the Carhart (1997) model (which nests the CAPM and the Fama and French, 1993, model and in addition produces higher adjusted R^2s).

In order to compare the usefulness of the five models we split our twelve year sample period into four non-overlapping three year subperiods and calculate the Alphas for all funds in the *Regression Sample*. Next, we sort the 1000 funds into quartiles with respect to their Alpha in each individual subperiod. This information is used to calculate transition probabilities for the subsequent subperiod. In general, a fund can end up in one of the four quartiles or it can stop reporting its performance to the database. Overall, our *Regression Sample* displays a probability of funds no longer reporting in the subsequent three year period amounting to 22.7%.

Table 3 reports the transition probabilities for the subsequent period calculated with the Carhart (1997) (Panel A), the Fung and Hsieh (1997) (Panel B), the Fung and Hsieh (2004) (Panel C), the Lhabitant (2001) (Panel D), and our SOM-based model (Panel E). If the Alphas calculated with the different performance attribution models were pure white noise, we would expect funds which ended up in one particular quartile in period t to have an equal probability of being ranked in any of the four quartiles in the subsequent period. Given the drop out rate of 22.7% the ex ante probability of a fund ending up in any of the four quartiles would be 19.3% in that case. We can see that for all models the probability of a fund staying in the same quartile is well above this 19.3% threshold except in one case (see Panel B, Fung and Hsieh, 1997, model: funds ranked in the 4^{th} quartile). However, only in the case of the SOM-based model we can observe the desired result that funds which do not leave the database will most likely stay in the same performance quartile in the next subperiod, which underscores the usefulness of the SOM-based model in detecting manager talent. Moreover, we can see that the probability of funds exiting the database in the next subperiod significantly decreases from hedge funds ranked in the fourth quartile to funds ranked in the second quartile for all models but the Fung and Hsieh (1997) model. Hedge funds ranked in the first quartile, however, display a drop out probability which is always higher than the corresponding probability for funds ranked in the second quartile. A possible explanation is that funds of the first quartile that cease reporting are primarily hedge funds which do not need further publicity due to their good past performance.

In a second step we use the 146 funds of the *Regression Sample* for which Alphas can be calculated in all four three-year subperiods to analyze the transition probabilities over longer time horizons. Table 4 shows these long term transition probabilities for the five performance attribution models analyzed. It is obvious that this sample, which is conditioned on a long return history plus survival, has different properties than the full *Regression Sample* of 1,000 funds analyzed before. For example, according to the SOM-based model, the average Alpha of this four-period sample is 0.17% per month higher than the average of the entire *Regression Sample*, which indicates a 2% annual survivorship bias. Therefore, the resulting transition probabilities draw a slightly differentiated picture.

Table 3: Transition probabilities (in %) of hedge funds ranked according to their Alpha for the subsequent period.

PANEL A: Carhart (1997)						
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	Ceased reporting	$p[\chi^2]$
1^{st} Quartile	23.6	19.0	19.6	24.6	13.2	0.006
2^{nd} Quartile	15.8	26.3	22.7	24.9	10.4	0.000
3^{rd} Quartile	15.9	20.5	24.3	20.5	18.9	0.042
4^{th} Quartile	15.2	13.4	18.8	21.6	30.9	0.000
PANEL B: Fung and Hsieh (1997)						
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	Ceased reporting	$p[\chi^2]$
1^{st} Quartile	24.2	18.1	12.0	20.6	25.1	0.000
2^{nd} Quartile	14.6	22.8	23.7	19.0	19.9	0.033
3^{rd} Quartile	14.2	22.0	29.0	19.9	14.8	0.000
4^{th} Quartile	24.7	17.9	14.8	19.0	23.6	0.010
PANEL C: Fung and Hsieh (2004)						
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	Ceased reporting	$p[\chi^2]$
1^{st} Quartile	25.3	17.0	19.1	23.0	15.6	0.003
2^{nd} Quartile	13.4	24.5	27.2	23.6	11.3	0.000
3^{rd} Quartile	15.7	21.9	24.4	19.6	18.4	0.029
4^{th} Quartile	16.5	12.5	11.7	24.6	34.8	0.000
PANEL D: Lhabitant (2001) (CSFB/Tremont Indices)						
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	Ceased reporting	$p[\chi^2]$
1^{st} Quartile	21.3	16.7	16.7	21.7	23.6	0.197
2^{nd} Quartile	15.3	28.8	24.0	16.6	15.5	0.000
3^{rd} Quartile	17.7	18.5	23.1	19.0	21.8	0.177
4^{th} Quartile	17.6	14.8	16.9	23.0	27.8	0.019
PANEL E: SOM-based Benchmarks						
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	Ceased reporting	$p[\chi^2]$
1^{st} Quartile	22.0	17.2	18.5	19.2	23.1	0.556
2^{nd} Quartile	15.0	26.5	25.2	16.1	17.2	0.000
3^{rd} Quartile	17.7	18.6	21.9	19.0	22.8	0.677
4^{th} Quartile	21.6	15.4	13.9	21.6	27.5	0.001

Rows indicate the ranking in the first period, columns give the ranking in the subsequent period. Hence, the first row shows the transition probabilities of funds which were ranked in the first quartile in the first period. For example, in Panel A, 24.6% of the funds which were ranked in the first quartile in the first period ended up in the fourth quartile in the subsequent period; 13.2% of the funds which were ranked in the first quartile in the first period ceased reporting some time during the subsequent period. (Rows add up to 100%, slight discrepancies are due to rounding.) The column labeled $p[\chi^2]$ gives the p-value of the χ^2 Goodness-of-Fit test for the null hypothesis of a uniform distribution across quartiles.

Table 4: Long term transition probabilities (in %) of hedge funds ranked according to their Alpha for the following three subperiods.

PANEL A: Carhart (1997)					
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	$p[\chi^2]$
1^{st} Quartile	26.8	18.8	22.5	31.9	0.155
2^{nd} Quartile	29.7	23.2	24.6	22.5	0.622
3^{rd} Quartile	17.4	30.4	26.8	25.4	0.171
4^{th} Quartile	26.1	27.5	26.1	20.3	0.635
PANEL B: Fung and Hsieh (1997)					
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	$p[\chi^2]$
1^{st} Quartile	28.2	29.9	18.8	23.1	0.310
2^{nd} Quartile	23.1	19.7	26.5	30.8	0.366
3^{rd} Quartile	15.4	31.6	35.0	18.0	0.004
4^{th} Quartile	33.3	18.8	19.7	28.2	0.076
PANEL C: Fung and Hsieh (2004)					
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	$p[\chi^2]$
1^{st} Quartile	24.6	22.2	26.2	27.0	0.881
2^{nd} Quartile	30.1	20.3	22.8	26.8	0.431
3^{rd} Quartile	16.3	34.2	30.9	18.7	0.009
4^{th} Quartile	30.9	22.8	19.5	26.8	0.308
PANEL D: Lhabitant (2001) (CSFB/Tremont Indices)					
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	$p[\chi^2]$
1^{st} Quartile	31.5	21.6	17.1	29.7	0.104
2^{nd} Quartile	19.8	35.1	25.2	19.8	0.074
3^{rd} Quartile	20.4	27.8	29.6	22.2	0.472
4^{th} Quartile	28.8	16.2	26.1	28.8	0.188
PANEL E: SOM-based Benchmarks					
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile	$p[\chi^2]$
1^{st} Quartile	30.6	23.4	20.7	25.2	0.506
2^{nd} Quartile	21.3	25.0	36.1	17.6	0.040
3^{rd} Quartile	22.2	25.9	27.8	24.1	0.864
4^{th} Quartile	27.0	24.3	14.4	34.2	0.030

Rows indicate the ranking in first three-year subperiod, whereas columns give the ranking in the subsequent nine-year period. The column labeled $p[\chi^2]$ gives the p-value of the χ^2 Goodness-of-Fit test for the null hypothesis of a uniform distribution across quartiles.

Over this longer time period it becomes evident that the models of Carhart (1997), Fung and Hsieh (1997) and Fung and Hsieh (2004) do not provide a long term investor with useful information. In almost all cases it is more likely that a hedge fund will end up in another quartile than remain within the same (see Panel A, B and C of table 4). Only the two models relying on actual hedge fund data (Lhabitant, 2001, and the SOM-based model) provide the investor with valuable information over this long time horizon. While the SOM-based model does a good job at identifying poorly performing hedge funds (see 4^{th} row of Panel E in table 4), funds initially ranked in quartile two show a tendency to move into quartile three. In the case of the Lhabitant (2001) model the probability of the weakest performing funds of staying in the fourth quartile is as high as the probability of jumping into the first quartile. Overall the long-term transition probabilities clearly underscore the value of the SOM-based model for the investor. Except for the Lhabitant (2001) model all competing performance attribution models seem to have some predictive power for the subsequent three year subperiod (see table 3), which might be caused by some sort of momentum effect, but do not succeed in detecting a manager's investment skill, which can be seen from the poor results with respect to the long-term transition probabilities (see table 4).

The finding of persistent excess returns can have two possible reasons. The first explanation is that a manager's investment skill is important in the hedge fund business. Therefore an expert of a particular trading strategy is able to consistently outperform his peers and thereby creates a persistently positive Alpha. This is in contrast to the mutual fund universe where Carhart (1997), for example, finds hardly any evidence of investment skill. Another possible explanation (also see Carhart, 1997) for persistent Alphas is a misspecification of the performance attribution model. Since persistence is detected with the two models based on hedge fund indices, which do not use "synthetic" risk factors like the option-based ABS-factors but investable benchmarks[9] and also feature the highest R^2 values in the out-of-sample regressions, we deem a misspecification unlikely.

5.2 Do Fund of Funds create a significant Alpha?

Single strategy hedge funds are normally run by one or two managers. Investors, however, might be interested in hedge funds as superior return vehicles on the one hand, while at the same time striving to reduce portfolio risk by entrusting their money to a larger group of decision makers and thereby diversifying across hedge fund styles. This can be achieved by

[9] The exact replication of the hedge fund indices might fail due to the inclusion of closed funds.

either forming a private portfolio of hedge funds or by investing into a fund of hedge funds (FOF). Establishing a private portfolio of hedge funds requires a considerable amount of money due to the high minimum investment requirements of most hedge funds. Investing in a fund of funds means that the investor commits the portfolio decision to a FOF manager, who is likely to be better informed but will charge fees for his services.

We use the SOM-based model as well as the other performance attribution models presented in the previous section to analyze whether funds of hedge funds can make up for their additional fees[10] by entrusting the money to exceptionally gifted single strategy hedge fund managers. For this purpose we compare the average Alpha of all single strategy hedge funds in the *Regression Sample* with the average Alpha of the 541 funds of hedge funds in the *FOF Sample*. Table 5 presents the results of this analysis. We can see that according to all performance attribution models used, funds of hedge funds produce a lower average Alpha than the single strategy hedge funds, which is statistically significant at the 1% level in all cases. The Lhabitant (2001) and our SOM-based model even report an average Alpha for the FOF category that is significantly smaller than zero (see third line of Panel B in table 5). These figures clearly show that most fund of funds managers produce an excess return which is considerably lower than the fees they charge. Therefore an informed investor should rather buy a self-selected portfolio of single strategy hedge funds than rely on the ability of the average fund of funds manager.

[10] The 541 funds in our *FOF Sample* feature a management fee of 0.8% and a performance fee of 15.0%, on average.

Table 5: Alpha comparison between the *Regression Sample* and the *FOF Sample* in the period from May 1992 to April 2004.

PANEL A: *Regression Sample* (1000 hedge funds)					
Statistic	**Carhart (1997)**	**Fung/Hsieh (1997)**	**Fung/Hsieh (2004)**	**CSFB/Tremont**	**SOM Benchmarks**
average (monthly) Alpha in %	0.49	0.59	0.71	0.20	0.05
Alpha of the median fund in %	0.46	0.20	0.72	0.24	0.06
standard deviation of monthly Alphas (across funds, in %)	1.01	9.52	1.18	1.63	1.57
% of funds with Alpha > 0	76.80	56.20	84.00	60.10	53.80
% of funds with Alpha < 0	23.20	43.80	16.00	39.90	46.20
% of funds with Alpha significantly* > 0	27.20	7.20	41.30	12.80	7.80
% of funds with Alpha significantly* < 0	2.30	1.70	1.80	3.50	5.50
Adjusted R^2 of the median fund	0.086	0.152	0.198	0.295	0.356

PANEL B: *FOF Sample* (541 funds of hedge funds)					
Statistic	**Carhart (1997)**	**Fung/Hsieh (1997)**	**Fung/Hsieh (2004)**	**CSFB/Tremont**	**SOM Benchmarks**
average (monthly) Alpha in %	0.21	-0.25	0.55	-0.09	-0.16
$p[\mu_{\alpha_{FOF}} < \mu_{\alpha_{RS}}]$	1.0000	0.9930	0.9997	1.0000	0.9998
$p[\mu_{\alpha_{FOF}} < 0]$	0.0000	0.9352	0.0000	0.9994	1.0000
Alpha of the median fund in %	0.23	0.07	0.57	-0.02	-0.13
standard deviation of monthly Alphas (across funds, in %)	0.61	3.87	0.60	0.67	0.69
% of funds with Alpha > 0	75.42	53.42	90.94	47.13	35.86
% of funds with Alpha < 0	24.58	46.58	9.06	52.87	64.14
% of funds with Alpha significantly* > 0	35.86	8.50	63.77	9.98	5.36
% of funds with Alpha significantly* < 0	2.40	2.59	1.29	12.01	14.97
Adjusted R^2 of the median fund	0.289	0.307	0.300	0.587	0.536

Panel A shows the results for five different performance attribution models applied to the *Regression Sample*, whereas Panel B gives the results of the same models for the *FOF Sample*. The second and third line of Panel B show the p-values of a one sided t test on the equality of means. The second line has the alternative hypothesis of $\mu_{\alpha_{FOF}} < \mu_{\alpha_{RS}}$ and the third line assumes the alternative hypothesis of $\mu_{\alpha_{FOF}} < 0$.

* $\alpha = 0.05$

6 Style Shifts

In this section, we analyze the link between style shifts and performance. While there has been some research on the issue of misclassification and style consistency in the hedge fund universe (see, for example, Bares et al., 2001; Amenc and Martellini, 2003; Schneeweis et al., 2004; Baghai-Wadji et al., Forthcoming), the issue of hedge fund style shifts has not been studied extensively so far.

Lynch and Musto (2003) address style shifts in the context of *mutual funds*. They introduce a two-period model and derive hypotheses which they test using funds' changes in risk loadings on the four factors of the Carhart (1997) model as well as data on manager replacements as proxies for style shifts. Due to the differences between mutual funds and hedge funds it is conceivable that the findings of Lynch and Musto (2003) are not necessarily applicable to the hedge fund business. In fact, as will be shown in the following section, the conclusions we draw from our sample of hedge funds are very different. Our main conjectures are summarized in the six hypotheses below, which will undergo empirical investigation in section 6.1.

Hypothesis 1: *Style shifts are a characteristic feature of the hedge fund business.* Hedge fund managers have considerable leeway in finding profitable investment strategies. They are not subject to the stringent regulations that mutual fund managers have to adhere to. The fact that hedge funds are expected to deliver "absolute returns" (i.e. positive returns irrespective of market conditions) and are known to employ very flexible ("dynamic") trading strategies, lends support to our conjecture that style shifts are very common in the hedge fund business.

Hypothesis 2: *Poor performers are more likely to change their investment style than hedge funds with a better track record.* We suspect that hedge funds in the bottom performance quartile are more prone to style shifts than funds in higher performance quartiles. The intuition is that fund managers with a bad track record will try to abandon their investment strategy which has proven to be unsuccessful, while fund managers who were able to deliver better performance are more likely to retain their previous strategies.

Hypothesis 3: *Funds which are more style-consistent, i.e. do not exhibit excessive style* *changes over time, feature significantly better performance than funds which are less style-consistent.* Brown and Harlow (2004), who investigate the link between style-consistency and mutual fund performance, give three reasons in support of this argument: First, a more consistent investment style enables a manager to avoid asset allocation and security selection mistakes arising from attempts to "time" investment decisions. Second, style-consistency can be interpreted as a signalling device used to communicate superior investment skill to investors. Third, style-consistency leads to lower portfolio turnover and hence reduces transaction costs.

HYPOTHESIS 4: *Replacing a strategy in one period does not, on average, lead to a performance improvement in the following period.* In general, there is no reason to assume that replacing a strategy with a new one should result in a performance improvement.

HYPOTHESIS 5: *Replacing a losing strategy with a new one does, on average, lead to a performance improvement.* While, in general, replacing one strategy by another will not necessarily improve performance, it is sensible to assume that shifting from a losing strategy to a new one will in fact entail a performance improvement. We conjecture that there exists a linear relationship between style shifts and performance improvements; specifically, Hypothesis 5 implies that badly performing funds which change their investment strategy are expected to show a more pronounced performance improvement than bad performers which don't.

HYPOTHESIS 6: *Replacing a winning strategy with another strategy does not, on average, result in a performance improvement. In fact, it entails a worsening of performance in the following period.* Hedge funds which are ranked in the top performance quartile in a given period and change their strategy cannot be expected to improve their performance *through a change in strategy.* Once a winning investment strategy is implemented, the manager is better off keeping the strategy than replacing it with a new one: "If it isn't broken, don't fix it."

At first glance, Hypothesis 5 and 6 seem to be slightly contradictory. Why should changing a losing strategy entail a performance improvement and altering a winning strategy lead to a deterioration in performance? In order to support these conjectures, let's assume that a particular manager-strategy combination is likely to result in persistent performance (this assumption is in line with our findings on persistent performance discussed in section 5.1). In other words, if a manager retains his strategy in the next period, he will, with high probability (at least above 25%), remain in the same performance quartile in the following period since he already has experience with that particular investment strategy. However, the implementation of a new investment strategy leads to an unproven manager-strategy combination, the result of which can be thought of as the outcome of a random experiment.

To be specific, suppose that a manager generated a bottom (top) quartile Alpha in one period. If he retains his strategy, he will most likely remain in the bottom (top) performance quartile. However, if he changes the investment algorithm, then performance in the next period can be viewed as a random draw from the Alpha distribution, and the bottom (top) quartile hedge fund will have, a priori, a 75% chance of improving (worsening) its performance in the following period.

For bottom quartile funds, the pressure is on their managers to "do something" about the disappointing past performance, which will most likely result in a shift in the strategy profile. A top performer on the other hand could, in some cases, be forced to implement a new investment strategy so as to accommodate the increased inflow of investor resources.[11] He will therefore deviate from the strategy he has previously mastered and will apply a new and unproven investment algorithm. In doing so, he draws from the aforementioned Alpha distribution which will, most likely, result in a worsening of performance, as outlined above.

6.1 Empirical Appraisal

6.1.1 Definition of Variables

In order to test the hypotheses outlined above, we require a proxy for strategy change. Lynch and Musto (2003), for mutual funds, use the change in factor loadings from the Carhart (1997) model as well as the event of a manager replacement as proxies for a change in strategy. Chan et al. (2002) employ changes in factor loadings from the Fama and French (1993) model, as well as changes in portfolio characteristics to analyze style changes and the impact of style

[11] The possibilities of making money with a particular niche-strategy are limited, especially since there are decreasing returns to scale in the hedge fund industry (see, for example, Agarwal et al., 2003).

shifts on mutual fund performance. Our proxy for strategy change will be based on the SOM-benchmark model since data on manager replacements and specific portfolio characteristics are not available for hedge funds and, as has been shown in previous sections, the SOM indices are much better suited to explain hedge fund returns than the Fama and French (1993) and the Carhart (1997) models.

We measure a shift in a fund's investment style via the average absolute change in the fund's factor loadings (regression coefficients) from one subperiod to the next.[12] Clearly, loading-changes are imperfect estimates of shifts in investment styles. The reliability of the proxy depends on the overall capability of the underlying model to capture variations in returns over time. In this sense, as has been demonstrated in the preceding sections, the multi-factor model suggested in this paper does a satisfactory job, especially if compared to other models currently used in the literature and in practice. Furthermore, as has been pointed out by Lynch and Musto (2003), for factor loading changes to be a useful proxy for strategy changes, it is sufficient to make the reasonable assumption that changing a strategy, on average, results in larger loading-changes than maintaining a given strategy. Therefore, a style shift is said to take place whenever there is a "sufficiently" large change in the regression coefficients on the factors specified in the model from one period to the next.[13]

Specifically, a loading-change (LC) for fund i from subperiod t to subperiod $t{+}1$ is defined as follows:

$$LC_{i,t} = \frac{1}{11} \sum_{j=1}^{11} |\beta_{i,j,t+1} - \beta_{i,j,t}| \tag{4}$$

where $\beta_{i,j,t}$ is fund i's loading on factor j during subperiod t; Our model is based on 11 factors, so there are 11 coefficients in each regression.

As a proxy for the risk adjusted performance we use the Alpha estimates of the SOM-based performance measurement model. In each subperiod, we calculate the Alphas of the hedge

12 Our definition is similar, but not identical to the definition suggested by Lynch and Musto (2003).
13 Note that, in order to test Hypotheses 1-6, there is no need to specify a threshold value for the loading-change above which a fund is assumed to have shifted investment style.

funds in the *Regression Sample* and sort them into quartiles. Poor performance is defined as bottom-quartile performance, while top performance corresponds to Alphas in the uppermost quartile.

6.1.2 Results

The results discussed in this section are based on the following calculations: The twelve-year sample period is divided into four non-overlapping three-year subperiods. In order to investigate Hypotheses 1-6 it is necessary to restrict the *Regression Sample* of 1,000 hedge funds to those 844 funds which have at least 12 return observations in each of two or more consecutive three year subperiods. After subdividing the twelve-year sample period from May 1992 to April 2004 into four non-overlapping three-year subperiods, we regress each fund in the sample on our set of benchmark indices. For each subperiod and hedge fund, we obtain Alpha and coefficient estimates. The Alpha estimates are ranked into quartiles, and the coefficient estimates for each subperiod are used to calculate the proxies for strategy changes (LC) from one subperiod to the next.

We employ Chow tests in order to assess the validity of Hypothesis 1.[14] More specifically, we determine the incidence of significant coefficient changes in regression (3) from one subperiod to the next, which we interpret as "style shifts". Using the sample of 844 funds described above, we find that on average 18.0% (5% significance level) of the funds significantly change their trading style from one three-year subperiod to the next. In the balanced sample of 146 funds (described in section 5.1), 40.4% (5% significance level) of the funds experience at least one significant style change during the 12 year sample period. Overall, these test results lend support to the conjecture outlined in Hypothesis 1: Style shifts are a characteristic feature of

[14] We assume the breakpoints to be at the boundaries of the three-year subperiods specified above. Since the sample size for each fund around a breakpoint is rather small, the following results are based on the F-statistic due to its better finite sample properties. The results based on the LR-statistic resulted in a considerably larger incidence of style shifts among the funds in our sample.

the hedge fund universe.[15]

In order to examine Hypothesis 2, we test whether or not all funds in our sample shift styles to a similar extent: Table 6 reports the output for the t-tests of the hypothesis of a same (two-sided test) or higher (one-sided test) average value of the style shift measure LC for the bottom compared to the upper three performance quartiles. From these results it follows that the average LC in the bottom-performance quartile is significantly bigger than the average LC of hedge funds in the 50%- and 75%-performance quartiles. However, the difference in the average LC measures between bottom and top performers is not significant. In other words, our sample suggests that bottom quartile performers feature more prominent style shift behavior than mediocre performers, but there is no significant difference between the degree of style changes of bottom and top performing hedge funds in our sample. The latter finding is rather counter-intuitive. It is also in contrast to findings by Lynch and Musto (2003) who, for *mutual funds*, document a greater incidence of large loading changes for bottom quartile performers than for mutual funds in any of the upper three performance quartiles. Chan et al. (2002) also document that style shifts in mutual funds occur almost exclusively after bad performance. However, it should be kept in mind that hedge funds by their very nature try to exploit short-lived market inefficiencies, which limits the amount of money that can be invested profitably in a particular niche-strategy. Therefore successful managers who are reluctant to close their fund might be forced to invest in new styles.

[15] Due to the nature of the SOM-based benchmark indices our methodology produces a very conservative estimate of the number of funds changing their trading style: A given index fully captures the salient investment features of a given trading style. The exposure of a given style to certain asset classes might change over time, but the index already takes into account that investment strategies are adapted to changing market environments. The methodology used in this paper implies that a style shift is only assumed to occur if an individual hedge fund deviates from the investment strategy represented by a given index, which is captured by a significant change in a fund's exposure to that index.

Table 6: T-tests of Hypothesis 2.

PANEL A: comparison of average LC of bottom performance quartile vs 50%-quantile.		
	LC (bottom quartile)	**LC (50%-quantile)**
mean / variance of LC	1.57 / 3.47	0.75 / 0.44
number of observations	351	349
t-statistic	7.75	
p-value one-sided test / two-sided test	0.000 / 0.000	
PANEL B: comparison of average LC of bottom performance quartile vs 75%-quantile.		
	LC (bottom quartile)	**LC (75%-quantile)**
mean / variance of LC	1.57 / 3.46	0.69 / 0.62
number of observations	351	350
t-statistic	8.10	
p-value one-sided test / two-sided test	0.000 / 0.000	
PANEL C: comparison of average LC of bottom performance quartile vs top-quartile.		
	LC (bottom quartile)	**LC (top quartile)**
mean / variance of LC	1.57 / 3.46	1.50 / 2.28
number of observations	351	352
t-statistic	0.53	
p-value one-sided test / two-sided test	0.298 / 0.596	

This table reports the output for the t-tests for Hypothesis 2. Funds are ranked into quartiles according to their average Alpha over the life of the fund. Then a t-test is used to investigate the hypothesis of an equal (two-sided test) or higher (one-sided test) average value of the style shift measure LC for the bottom performance quartile compared to the 50%-performance quantile (Panel A), the 75%-performance quantile (Panel B), and the top performance quartile (Panel C). In analogy to the Alphas, the LC values used for the tests are the averages of the LC measures obtained for each subperiod for each fund.

Hypothesis 3 suggests that funds which are more style-consistent are better performers than funds with more pronounced style shifts. We test this conjecture in the following way: First, we calculate average loading changes for all funds in our sample and rank the funds into quartiles according to the loading-change proxy LC. Funds which are sorted into the bottom quartile along the loading-change dimension are defined as "style consistent." Using a t-test, we then compare the average performance of funds in the bottom LC quartile (i.e. style consistent funds) with the performance of funds in the top LC quartile (i.e. the 25% least style-consistent hedge funds in our sample). As can be seen from Table 7, there is no significant difference in performance between style-consistent hedge funds and funds with style shifts in high-gear, and hence no empirical support for Hypothesis 3. This is in contrast to findings of Brown and Harlow (2004), who, using a different methodology, report that style-consistent *mutual* funds perform better, on average, than less consistent mutual funds.

Table 7: T-tests of Hypothesis 3.

	bottom quartile LC	**top quartile LC**
mean / variance of Alpha	0.0011 / 0.0000	0.0031 / 0.0052
number of observations	211	211
t-statistic	-0.39	
p-value one-sided test / two-sided test	0.348 / 0.696	

This table reports the output for the t-test for Hypothesis 3. First, we calculate average loading changes for all funds in our sample and rank the funds into quartiles according to the loading-change proxy LC. Funds which are sorted into the bottom quartile along the loading-change dimension are defined as "style consistent." Using a t-test, we then compare the average performance of funds in the bottom LC quartile (i.e. style consistent funds) with the performance of funds in the top LC quartile (i.e. the 25% least style-consistent hedge funds in our sample). There is no significant difference in performance between style-consistent hedge funds and funds with style shifts in high-gear. T-tests comparing the performance of bottom LC quartile funds with the performance of funds in the 50%- and 75%- LC quantiles lead to the same conclusion but are not reproduced here.

In order to test Hypothesis 4, we regress, for all hedge funds in our sample, the change in Alpha (our proxy for risk-adjusted performance of a hedge fund) from subperiod t to subperiod $t{+}1$ on the style shift measure LC_t corresponding to a loading change from period t to subperiod $t{+}1$:

$$(\alpha_{i,t+1} - \alpha_{i,t}) = a_t + b_t LC_{i,t} + e_{i,t} \tag{5}$$

According to Hypothesis 4, there is no reason to assume that, in general, replacing a strategy with a new one should result in a performance improvement in the following period. As can be seen from Table 8, the regression results for all subperiods are consistent with Hypothesis 4. The low R^2 values and the insignificant coefficient estimates indicate that for the whole sample of hedge funds across all performance quartiles, the existence of style shifts alone cannot explain differences in performance from one subperiod to the next.

Table 8: Regression results for Hypothesis 4.

	regression (1) change in Alpha (1992-1998)	**regression (2)** change in Alpha (1995-2001)	**regression (3)** change in Alpha (1998-2004)
loading change (LC) 1992-1998	0.0105		
	(0.0078)		
LC 1995-2001		0.0132	
		(0.0097)	
LC 1998-2004			0.00967
			(0.027)
Constant	-0.0113	-0.0156*	-0.00980
	(0.0095)	(0.0094)	(0.020)
Observations	379	509	514
R^2	0.06	0.08	0.02

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

This table reports the results from regressions of the change in Alpha (our proxy for risk-adjusted performance of a hedge fund) from subperiod t to subperiod $t+1$ on the style shift measure LC_t corresponding to a loading change from period t to subperiod $t+1$. Robust (White) estimates of standard errors are given in parentheses. For example, the first column in the table shows the results from regressing the change in Alpha from subperiod 1992-1995 to subperiod 1995-1998 on the loading-change proxy LC obtained from the change in factor loadings between subperiod 1992-1995 and subperiod 1995-1998.

In Hypothesis 5, we conjecture that replacing a losing strategy with a new one does, on average, lead to a performance improvement. In order to test this, as for Hypothesis 4, we regress the change in Alpha from one subperiod to another on the style shift measure LC corresponding to the same subperiods. However, in the tests for Hypothesis 5, we restrict our sample in each subperiod to those funds which are in the bottom performance quartile as measured by the regressions' Alphas.

If Hypothesis 5 holds, we would expect a significantly positive regression coefficient on the style shift proxy. The reason is that we conjectured that a change in strategy, which manifests itself in a larger value of the style shift proxy LC, leads to a performance improvement for bottom quartile funds. Table 9 indicates that there is strong support for our theory: The regression coefficients for all sub-periods analyzed are significantly (for $\alpha = 0.01$) different from zero and positive. With all the necessary caveats, this result implies that for bottom-quartile performers, the larger the extent of the style shift, the more substantial the performance improvement will be in the subsequent period.

Table 9: Regression results for Hypothesis 5.

	regression (1) change in Alpha (1992-1998)	**regression (2)** change in Alpha (1995-2001)	**regression (3)** change in Alpha (1998-2004)
loading change (LC) 1992-1998	0.0333***		
	(0.0059)		
LC 1995-2001		0.0336***	
		(0.011)	
LC 1998-2004			0.0559***
			(0.019)
Constant	-0.0222**	-0.00385	-0.0311*
	(0.0098)	(0.015)	(0.018)
Observations	95	127	129
R^2	0.61	0.62	0.61

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

This table reports the results from regressions of the change in Alpha from subperiod t to subperiod *t+1* on the style shift measure LC_t corresponding to a loading change from period t to subperiod *t+1*. Only loading- and Alpha-changes from funds in the bottom performance quartile are used in the regressions. Robust (White) estimates of standard errors are given in parentheses.

Hypothesis 6 suggests that replacing a winning strategy with another strategy does not, on average, result in a significant performance improvement. In fact, as has been argued above, we expect a worsening of performance following the style change. In order to test this conjecture, we proceed in a similar fashion as in the tests for Hypotheses 4 and 5, namely by regressing Alpha changes of top performers on the loading-change variable LC. Overall, the regression results documented in Table 10 support Hypothesis 6: The negative coefficient estimates indicate that style shifts implemented by top-quartile performers result in a subsequent performance deterioration. As a caveat it should be noted that the relation is somewhat weaker for subperiod 1995-2001, however while not statistically significant, the negative sign of the coefficient is still in line with Hypothesis 6.

Table 10: Regression results for Hypothesis 6.

	regression (1) change in Alpha (1992-1998)	**regression (2)** change in Alpha (1995-2001)	**regression (3)** change in Alpha (1998-2004)
loading change (LC) 1992-1998	-0.0169***		
	(0.0025)		
LC 1995-2001		-0.0120	
		(0.019)	
LC 1998-2004			-0.0480***
			(0.016)
Constant	-0.00144	-0.0254	0.0177
	(0.0039)	(0.023)	(0.015)
Observations	95	128	129
R^2	0.34	0.07	0.49

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

This table reports the results from regressions of the change in Alpha from subperiod t to subperiod $t+1$ on the style shift measure LC_t corresponding to a loading change from period t to subperiod $t+1$. Only loading- and Alpha-changes from funds in the top performance quartile are used in the regressions. Robust (White) estimates of standard errors are given in parentheses.

7 Conclusion

In this paper, we present a new approach for evaluating the performance of hedge funds, which relies on a classical multi-factor model in the tradition of Sharpe's (1992) asset class factor model, but features neural network derived hedge fund benchmarks. We use the Self-Organizing Map (SOM) algorithm, which is a neural network particularly suitable for cluster analysis, to group hedge funds into homogenous style-consistent categories. By classifying the hedge funds based on their monthly return histories, we are able to identify homogenous hedge fund classes, which are not influenced by the well-known problem of faulty self-classification. We use all hedge funds contained in a specific investment style cluster to construct benchmark indices, which capture the features of the underlying dynamic trading strategies. These SOM-based benchmark indices allow us to use a simple linear regression framework for performance measurement of hedge funds while still allowing for non-linearities in the analysis via the "dynamics" embedded in the regressors. Moreover, these SOM-based hedge fund indices mitigate the problem of omitted risk factors encountered in other models traditionally used for hedge fund performance evaluation.

We show that the explanatory power of our SOM-based model exceeds all other analyzed performance attribution models in the case of our *Regression Sample* comprising 1,000 funds. We also investigate the evolution of Alphas over time, in order to determine whether the performance attribution models analyzed are capable of distinguishing skilled from untalented managers. Over our twelve year sample period only the SOM-based approach and the model of Lhabitant (2001), which also uses hedge fund indices, are able to detect funds with persistent excess returns. This persistence of Alphas indicates that our SOM-based performance attribution model is able to distinguish excessive risk taking from true investment skill and can therefore contribute valuable information to the portfolio selection process of individual investors. Moreover fund of hedge funds managers could use our model for the selection of single strategy hedge funds. Our analysis of the performance of funds of hedge funds suggests that fund of hedge funds managers still have to improve considerably in this discipline. All tested performance measurement models lead to the conclusion that the average fund of hedge funds does not create a sufficiently high excess return to compensate for the additional layer of fees.

Finally, we examine the link between style shifts and performance in the hedge fund universe. Previous research reported that in the mutual fund industry, style changes are more common among poorly performing funds and that style-consistent funds produce higher returns than funds with more pronounced style shifts (see, for example, Lynch and Musto, 2003; Brown and Harlow, 2004). Our findings suggest that hedge funds are a different breed of investment vehicle in this respect. The empirical tests performed lend support to the hypothesis that style shifts are a characteristic feature in the hedge fund business. In contrast to findings for mutual funds, however, our research suggests that style changes are at least as wide-spread after good performance as they are after bad performance. We also show that style-consistency does not, on average, lead to better performance compared to funds which display more pronounced shifts in investment style over time. Furthermore, we also demonstrate that funds with bad performance which change style can expect a performance improvement in the period following the strategy change, whereas top performers which alter their strategy are more likely not to benefit from the style shift.

Acknowledgements

The authors would like to thank Alois Geyer, Alex Michaelides, and Josef Zechner, seminar participants at the University of Vienna (Rust, 2005), as well as participants of the German Finance Association 2006 meeting in Oestrich-Winkel, Germany, for helpful comments. Moreover we thank the OeNB (Oesterreichische Nationalbank) for providing access to the CISDM hedge fund database.

References

Ackermann, C., R. McEnally, and D. Ravenscraft (1999): "The performance of hedge funds: Risk, return, and incentives," *Journal of Finance*, 54, 833–874.

Agarwal, V., N. Daniel, and N. Naik (2003): "Flows, Performance, and Managerial Incentives in Hedge Funds," *EFA 2003 Annual Conference Paper No. 501, August 2003*.

Agarwal, V. and N. Naik (2004): "Risk and Portfolio Decisions involving Hedge Funds," *Review of Financial Studies*, 17(1), 63–99.

Amenc, N., S. El Bied, and L. Martellini (2003): "Predictability in Hedge Fund Returns," *Financial Analysts Journal*, 59, 32–46.

Amenc, N. and L. Martellini (2003): "The Alpha and Omega of Hedge Fund Performance Measurement," *Working Paper*, EDHEC.

Baghai-Wadji, R., R. El-Berry, S. Klocker, and M. Schwaiger (Forthcoming): "Changing Investment Styles - Style Creep and Style Gaming the Hedge Fund Industry," *Intelligent Systems in Accounting, Finance & Management*.

Bares, P., R. Gibson, and S. Gyger (2001): "Style Consistency and Survival Probability in the Hedge Funds' Industry," *Working Paper*, Swiss Institute of Banking.

Brown, K. and W. Harlow (2004): "Staying the Course: Mutual Fund Investment Style Consistency and Performance Persistence," *Working Paper*, University of Texas.

Brown, S. and W. Goetzmann (1997): "Mutual fund styles," *Journal of Financial Economics*, 43(3), 373–399.

——— (2003): "Hedge Funds with Style," *Journal of Portfolio Management*, 29(2), 101–113.

Carhart, M. (1997): "On persistence of mutual fund performance," *Journal of Finance*, 52, 57–82.

Chan, L., H.-L. Chen, and J. Lakonishok (2002): "On Mutual Fund Investment Styles," *Review of Financial Studies*, 15(5), 1407–1437.

Cottrell, M., E. De Bodt, and M. Verleysen (2001): "A statistical tool to assess the reliability of self-organising maps," in *Advances in Self-Organising Maps*, ed. by N. Allinson, H. Yin, L. Allinson, and J. Slack, Springer, 7–14.

Cumming, D., G. Fleming, and A. Schwienbacher (2005): "Style drift in private equity," *Working Paper*, University of Amsterdam.

Deboeck, G. and T. Kohonen (1998): *Visual Explorations in Finance with Self-Organizing Maps*, Springer, first ed.

diBartolomeo, D. and E. Witkowski (1997): "Mutual Fund Misclassification: Evidence Based on Style Analysis," *Financial Analysts Journal*, 53(5), 32–43.

Edwards, F. and M. Caglayan (2001): "Hedge Fund Performance and Manager Skill," *The Journal of Futures Markets*, 21, 1003–1028.

Fama, E. and K. French (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33, 3–56.

Flexer, A. (2001): "On the use of self-organizing maps for clustering and visualization," *Intelligent Data Analysis*, 5, 373–384.

Fung, W. and D. Hsieh (1997): "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds," *The Review of Financial Studies*, 10, 275–302.

——— (1998): "Performance Attribution and Style Analysis: From Mutual Funds to Hedge Funds," *Working Paper*, Duke University.

——— (2000): "Performance Characteristics of Hedge Funds and Commodity Funds: Natural vs. Spurious Biases," *Journal of Financial and Quantitative Analysis*, 35, 291–307.

——— (2001): "The risk in hedge fund strategies: theory and evidence from trendfollowers," *The Review of Financial Studies*, 14, 313–341.

——— (2002): "Asset-Based Style Factors for Hedge Funds," *Financial Analysts Journal*, 58, 6–27.

——— (2004): "Hedge Fund Bechmarks: A Risk-Based Approach," *Financial Analysts Journal*, 60, 56–80.

Gallo, J. and L. Lockwood (1999): "Fund management changes and equity style shifts," *Financial Analysts Journal*, 55(5), 44–52.

Glosten, L. and R. Jagannathan (1994): "A contingent claim approach to performance evaluation," *Journal of Empirical Finance*, 1(2), 133–160.

Gregoriou, G. N., K. Sedzro, and J. Zhu (2005): "Hedge fund performance appraisal using data envelopment analysis," *European Journal of Operational Research*, 164, 555–571.

Jensen, M. (1968): "The Performance of Mutual Funds in the Period 1945–1964," *Journal of Finance*, 63, 389–416.

Kandel, E., S. Kandel, and R. Wermers (2004): "Endogenous Benchmarks," *Working Paper*, University of Maryland.

Kim, M., R. Shukla, and M. Tomas (2000): "Mutual fund objective misclassification," *Journal of Economics and Business,*, 52(4), 309–323.

Kohonen, T. (2000): *Self-organizing maps*, Springer, third ed.

Kohonen, T., J. Hynninen, J. Kangas, and J. Laaksonen (1995): "SOM_PAK, The self-organizing map program package," *Helsinki University of Technology, Laboratory of Computer and Information Science.*

Lhabitant, F.-S. (2001): "Hedge funds investing: A quantitative look inside the black box," *Journal of Financial Transformation*, 1, 82–90.

——— (2002): *Hedge Funds: Myths and Limits*, John Wiley & Sons.

Lintner, J. (1965): "The Valuation of Risk Assets and The Selection of Risky Investments in Stock Portfolios and Capital Budgets," *Review of Economics & Statistics*, 47, 13–37.

Lynch, A. and D. Musto (2003): "How Investors Interpret Past Fund Returns," *The Journal of Finance*, 58, 2033–2058.

McGuire, P., E. Remolona, and K. Tsatsaronis (2005): "Time-varying exposures and leverage in hedge funds," *BIS Quarterly Review (March*, 2005), 59–72.

Mergner, S. and J. Bulla (2005): "Time-varying Beta Risk of Pan-European Industry Portfolios: A Comparison of Alternative Modeling Techniques," *Working Paper*, Georg–August–University Gottingen.

Mitchell, M. and T. Pulvino (2001): "Characteristics of Risk and Return in Risk Arbitrage," *Journal of Finance*, 56, 2135–2175.

Mossin, J. (1966): "Equilibrium in a Capital Asset Market," *Econometrica*, 34, 768–783.

Schneeweis, T., H. Kazemi, and V. Karavas (2004): "Hedge Fund Style Consistency and Performance," *Working Paper*, University of Massachusetts.

Sharpe, W. (1964): "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance*, 19, 425–442.

——— (1992): "Asset Allocation: Management Style and Performance Measurement," *Journal of Portfolio Management*, 18, 7–19.